WORKING NOTE AC57

Alan Creak
22 September 1986

SOME NOTES ON THE COURSE
ROBOTICS AND REAL-TIME CONTROL
OFFERED IN 1986

HOW IT ALL STARTED

For some years, the Computer Science Department has listed among its courses the entry "07.473 : Robotics and Real-time Control", but until 1986 the course had never been presented. My attention was drawn to the course when Steven Lomas, then a stage 3 student, inquired about the possibility of its being available in 1986, with a view to including it in his coursework for MSc.

I made a guarded reply. Although I had some knowledge of real-time computing, mainly through a long-standing project on designing a computer language for real-time work, I could certainly lay no claim to expertise in the field. At the same time, I had already come to believe that the Computer Science department should show some activity in the area of real-time computing, as it is unquestionably a topic of increasing importance in the spectrum of computer applications.

There was one possible course of action which could perhaps satisfy everyone. I began to think in terms of a reading course for 1986, in which one or two students and I would explore the literature on real-time control computing, and, in effect, write my lecture notes for a "real" course to be offered in 1987. I was confident that my knowledge and experience, while a shaky base for a full lecture course, were adequate to set directions of work for a less formal study. I said as much to Steven, and also mentioned that interest from a few other students would help me to make a more persuasive case to the department for putting on some sort of course in 1986.

Steven turned out to be an effective missionary - if anything, too effective. Within a few weeks, more than 10 stage 3 students had expressed interest in attending a Master's course on a real-time control topic. It was, of course, possible that not all the students would decide to continue for a Master's degree; but the response certainly suggested that an informal reading course of the sort I had envisaged would not do. It would have to be a full lecture course or nothing.

I could doubtless invent two or three convincing reasons for deciding to offer a full course. It *was* a gap in the department's offering; it *is* an important aspect of computer use; it *is* sufficiently distinct from other topics to merit a separate course; its intellectual content *does* justify academic treatment. But I suspect that the real reason was pig-headedness. By the time it came to make a decision, I had become quite enthusiastic at the prospect - and I don't much take to being beaten by challenges. I did make sure to my own satisfaction that, despite my earlier reservations, I would be able to get together sufficient of a course that students' time would not be wasted; but then I took the plunge, and, with the department's agreement, started planning for a full lecture course.

HOW THE COURSE WAS BUILT.

The course syllabus was determined primarily by its title and by the students' background. The title - Robotics and Real-time Control - was clearly the wrong way round, as ideas of control must precede their application to robotics, but at least delineated the boundaries of the subject. The students' knowledge determined the starting point : they were highly knowledgeable and experienced in what might be called pure computing topics, but I could assume no familiarity at all with any sort of engineering. A significant part of the course would therefore necessarily be given to an introduction to ideas of control.

After that, the material must be predominantly concerned with the computing aspects of the field. Departments in the Faculty of Engineering were already presenting courses which covered control in mechanical, electrical, and chemical engineering, and it was inappropriate for a Computer Science course to wander too far into detailed discussion of such areas.

I tried, without success, to find a textbook. There are books on control in all imaginable branches of engineering; there are books at any level from introductory to incomprehensible on control theory; there are even books on what might be termed advanced aspects of computing for real-time control. But the advanced books are too advanced, and the elementary one are for people who know more about engineering than computing; there seems to be nothing for computists with little knowledge of engineering. (This point is underlined by a response from one publishing company which I approached in search of a suitable text : it was an invitation to write one !)

It seemed that I had invented a new subject. Well, perhaps that's something of an exaggeration, but that's how it felt. It was up to me to drag together an appropriate set of topics, and to try to impose some sort of order and structure upon them. During the long vacation of 1985-6 I was inclined to be even more abrupt than usual with people inquiring as to how I was enjoying my holidays.

Yet another factor which influenced my choice of material to cover was my perception of the aim of the course. It was clearly not sensible to try to produce expert control engineers in a course of 30 to 40 lectures which would be about one sixth of a participant's workload; it wasn't even desirable, and any such attempt would rightly be greeted with derision. It seemed much more practicable, and much more important, to address the problem of communication between computist and engineer, and to present the sort of information which would help members of the two professions to cooperate fruitfully when they meet. The aim was therefore to give a wide coverage emphasising principles, rather than to follow up particular topics to greater depth.

One principle from the reading course survived : that I could make the students do some of the work. It was no longer appropriate to see a major part of the material as a do-it-yourself effort; an important feature of a reading course is the individual guidance resulting from discussions of the material gleaned between lecturer and student, and such personal attention was out of the question for a large class. But it did seem possible to ask the students to prepare material on some single topic, and a good combination of activities seemed to be for me to present a basic course outlining the subject and introducing the important topics, then for the students to present seminars covering more advanced work in a variety of tasks. This would illustrate the many ways in which the ideas of real-time control could be applied in practice.

I sketched out a rough syllabus, incorporating what I saw as the important topics, and worked on it until I had what seemed to be a list of headings each of roughly equal importance : those were to be the lecture topics for the course. I took a sheet of paper for each topic, and wrote the name of the topic and the head, then sat back with some satisfaction : I had made a start on my lecture notes.

I will not bore you with a blow-by-blow account of how I spent my summer "holidays". I aimed for 20 lectures. Some of my topics turned out to be smaller than I had expected, or to be harder to find out about, so I finished up with a few short lectures. Many more topics went to the opposite extreme, and the main task was to select the appropriate material. I'd expected that, of course, and intended to reorganise the material later; as it turned out, I was overtaken by time, so that's a job which still remains to be done.

PRACTICAL WORK.

Lectures and seminars, be they never so brilliant, are not all : practical experience of a subject like real-time control is indispensable. Even back in the days of the reading course, I had wondered how to incorporate appropriate practical experience; with the development of the plan into a fully-fledged lecture course, the question became even more pressing.

The Computer Science department owned exactly one object which was, even by stretching the imagination, appropriate as a subject for experiment in real-time control : a "turtle", which could respond to instructions from a microcomputer to move forward or backwards, to turn, or to beep. It was also fitted with four contact sensors which the computer could interrogate. In some ways, such a simple device is well suited to illustrating the principles of computer control - but a single turtle didn't seem enough to support the practical work of a dozen Masters' students.

What else was available ? I remembered reading an article many years earlier about the equipment used at the University of Newcastle-upon-Tyne to teach the principles of control : an elaborate model electric railway. The article was sketchy, but the idea seemed good. Negotiation with my son gave us access to a model railway - not, true, on the same scale as the Newcastle system (they had 31 locomotives; we only had two, one of which was in marginal condition), but, again, sufficient to illustrate some basic principles.

That was still hardly enough. There was no immediate question of the department's buying any equipment specially for the course; the 1986 run was an experiment, and the course's continued existence was to be reviewed in the light of the first year's experience and the general circumstances of the department.

The gap was filled by the generosity of the Mechanical Engineering department, who allowed us to use their equipment - in particular, a robot arm, a charge-coupled diode array camera, and a microcomputer connected to a general interface module. Without their help, the course would still have run; but it would not have been nearly so effective. Even with the engineers' contribution, there was a tense time when it seemed that even more equipment would be needed to serve a potential class of 17, and I had begun inquiries which I hoped would give me access to an analogue computer; but the threat receded, and all was well.

What did we do with the equipment ? My plan was to set three assignments, so that each group of two or three students would use at least three of the different pieces of machinery at our disposal. I didn't wish to circumscribe the work to be done too closely, as I was dealing with postgraduate students who were eager to learn, and quite capable of using their own initiative to find useful and interesting avenues to explore. (The other reason was that I didn't know enough about the machinery myself to set specific assignments !)

I therefore devised three rather general tasks. The first was to find out how a piece of machinery was, or could be, controlled, and to write it down for future reference; this exercise should convey some of the general ideas of controlling machinery, and incidentally teach something about the difficulty of preparing adequate documentation.

The second task was to use a different piece of equipment, and the previous group's documentation, to perform some fairly straightforward operation; the third task I didn't define too carefully at the beginning of the course, for it was possible that experience gained as the course proceeded would suggest a different form for the assignment, but my contingency plan was to tackle some more challenging task using yet a third machine.

HOW IT WENT.

15 people enrolled for the course, two of whom dropped out during the first term. (One found he had taken on too many courses; the other, who was a part-time student, changed his job and was unable to continue.)

My 20 lectures, which I had intended to finish during the first term, in fact stretched well into the second; they just finished in time for the students' seminars to begin. The seminars themselves were, by and large, very good; they were certainly a valuable part of the course.

The practical work, unfortunately, did not go so well. We got off to a slow start (largely my fault, through slow organisation), and only completed two of the assignments. Some of the difficulties were the sort of teething troubles one might expect during the first year of a new course : bits and pieces had to be gathered, documentation had to be hunted out, facilities had to be found, and so on. Another difficulty concered access to the engineers' equipment. Our MSc students are accustomed to unrestricted access to the facilities they need, and tend towards a pattern of work in which the daytime hours are spent on lectures, tutoring, and the like, while practical work happens in the evenings, when it is possible to spend long uninterrupted periods at work. The engineering laboratories were only accessible during the normal working day, so most of the experimental work was pushed into the vacations. Nevertheless, I believe that the students learnt a great deal from the assignments - though it wasn't all about real-time control.

ASSESSMENT.

Assessment is not my favourite topic. One can argue about its necessity, but I believe that, at the very least, it occupies time both of teacher and student which could be better used. Nevertheless, as the University is organised at present, it has to happen, and in particular it had to happen in this course.

I decided early on that the assessment should be based entirely on work done in the course, and not on a final examination. There were two reasons for that decision: first, I was far from confident that I would be able to set an examination of appropriate level in a subject which was almost as new to me as it was to the students; second, and more importantly, I was concerned that the importance of the practical work should be emphasised. The marks given were therefore based on the students' performance in the assignments, and on the seminars which they presented.

In both parts of this assessment, I included a contribution from the students themselves : I asked them to assess each seminar delivered, and also to assess the documentation they inherited from other groups during the assignment work. The object of this was not to relieve myself of the work (I marked them too, and the final marks reflect my own assessment in greater proportion than the students'), but to underline the importance of the intended audience in any presentation.

WHAT WE LEARNT.

If the students learnt a quarter as much as I did, they got good value from the course. For one thing, I learnt something about real-time control and a lot about robotics. But I also learnt a lot about the course, rather than its subject matter, and that's what this section is about.

The course material was about right, though some emphases need shifting. I laboured some descriptive areas rather more than they warranted, and left too little time to cover the robotics satisfactorily, but I think the topics included were representative of the field, and gave a reasonably complete coverage of the essential points.

I've already mentioned that the seminars were good. The standard was, of course, not uniform, but they made a very significant contribution to the course, and I'll certainly want to include them in future courses.

Having said that, it may seem ungracious to criticise, but a few comments are in order. I had hoped, and intended, that each student would put in some work to follow up current developments in the assigned topic, and tell us about the way computers and computing were used in that field. While some followed that pattern more or less, some didn't. Some seminars went little further than the lecture course (which could, I suppose, mean that the students hadn't understood the lectures); some gave an elementary description of a slightly related topic; rather few concentrated on the computing component of their topics, most emphasising the mechanical and electronic wonders. In the end, I'd hoped to have accumulated a wide range of references to recent work in many fields; in fact, my harvest has been much leaner than I expected. None of this is necessarily the students' fault - but now I know what to ask for, I can ask for it more specifically next year.

The practical work, like the seminars, did not produce quite the benefit that I'd hoped. Part of this was perhaps inevitable, and can be put down to everyone's breaking new ground. There were a few more specific difficulties. I have already mentioned the awkwardness of using the engineers' machinery, and this was compounded by the machinery's breaking down from time to time. The engineers' generosity this year has been most welcome, not to say indispensable, but any such arrangement is bound to lead to conflicts of interest sooner or later, and the greater convenience and freedom of having our own equipment will make for a much better course. My son's train proved, on closer inspection, quite inadequate; but that problem at least has now been remedied by our acquiring a new one. Documentation for the various machines we used was hard to find, and not always adequate even when found.

There was another reason for the rather unsatisfactory contribution of the practical part of the course. This is perhaps my fault : it springs from a, maybe naively optimistic, belief that postgraduate students are intelligent people, motivated by the fascination of their subject; and that they can therefore be left to explore the subject in ways which seem profitable to them. I am unwilling to forgo this belief, because it still works - like thermodynamics, in an isolated universe. It didn't work in this course because the students were all under great pressure to complete assignments set in other Masters' courses, and the exploratory urge got pushed aside. I don't know what to do about this. I do not wish to set specific assignments to postgraduate students : they should not be treated like schoolchildren. But if I don't, they won't get the benefit of doing the practical work in this course.

A partial solution may lie in a change which I hope to make for next year's course. One of the inconveniences of having several quite different pieces of equipment is that students don't really find out about all of them in detail, and certainly not the practical detail which I believe should be an important feature of the course. I shall therefore aim to arrange a weekly meeting of all those in the course in which each group will give an account of work done, problems encountered, and problems solved since the previous meeting, so that each can benefit from the others' experience.

The course was hard to assess because of the grouped work in the assignment. Although I had foreseen this difficulty, and asked for an account from each group of who did what, and for a log of work done, no group handed in both these items, and some handed in neither. Next year I will have to be much more careful. The students' assessments of each others' seminars were interesting. They were (perhaps unsurprisingly) uniformly very generous to each other, but there was no common perception of what constituted a good seminar. In the end, there was very little variation between the marks given over the whole class.

A rather disappointing feature of the seminars was the rather low attendance. Once again, perhaps, interest took second place to pressure of work - which was unfortunate, because much good material was presented in the seminars. I'm not sure what to do about that.

WHAT DOES IT ALL MEAN ?

These are four conclusions which I drew from running the course on Robotics and Real-time Control in l986.

- The course should be continued. The subject matter is interesting, academically demanding, and not treated elsewhere in the department.

- We should try to get equipment of our own. Not only would that make for easier access, but we would feel less inhibited at experimenting with it. The equipment does not need to be very expensive or elaborate; but it does need to be varied and adaptable, to illustrate different aspects of control systems and robotics. To underline that conclusion, it's interesting to observe that, despite the many difficulties, the two pieces of equipment which lived in the Computer Science department - the turtle and the train - proved much more successful than the engineers' machines. (Indeed, no student using the engineering school machinery actually completed the second assignment in the way intended.)

- We also need some facilities which students may use for carrying out small hardware jobs. The class were able to cope this year by leaning heavily on the resources of the engineers and the Physics department; that's hardly satisfactory as a long-term arrangement ! One could wonder at the paradox of the Mechanical Engineering department's being far better equipped for computer electronics than the Computer Science department.

- We need more technical staff. Our students are working with equipment which is more or less comparable with that used by the stage 3 engineering students, and during this course we have had occasion to refer frequently to project reports produced by such students in the Mechanical Engineering department. The lists of acknowledgments in these reports make interesting reading : they almost always thank the technical staff for assistance, and frequently for constructing the equipment used. Our students tackle work which is at least equally demanding - and they have to do it by themselves.

ACKNOWLEDGMENTS.

- The students who enrolled for the course, for their courage, fortitude, and resilience in the face of the unexpected : Richard Seaman, Stephen Hawkins, Chris Fromont, David Masalehdani, John-Maarten Dales, Tom Young, Martin Kealey, Paul Qualtrough, Steven Lomas, Lex Miller, Robert Chew, Sean Fennell, Shane Clerk, Paul Reddy, Kerry Thompson.

- Members of the Mechanical Engineering Department, for their generosity, help, and patience with many silly questions : Associate Professor Günter Arndt, Mr George Blanchard, and Dr R. Rudziejewski.

DISACKNOWLEDGMENTS.

- To a number of New Zealand firms concerned with control systems. In the interests of acquiring up-to-date information and examples for the course I wrote to several such companies, explaining my position and asking for any literature which they could give. Not one replied, even to refuse. (But honourable mention to Philips and to Perkin-Elmer, both of whom responded helpfully in Auckland to initial requests I'd sent overseas.)