WORKING NOTE AC85

Alan Creak
5 January 1994

SPACE FOR ROBOTICS AND REAL-TIME CONTROL

In this paper, I present a case for certain space requirements for the Robotics and Real-time control course, 07.473. There is also a quantity of background material which I added because I though it was interesting. Those who just want the punch line should go straight to the last section.

The course *07.473 : Robotics and real-time control* has been listed in the Calendar since the Computer Science department came into existence. In retrospect, we owe a lot to Richard and Rob (or whoever it was) for devising, long before there were people around to offer the courses, a set of topics which have served us pretty well.

R&R (no relation) languished until 1985, when Steven Lomas noticed it, and thought that it would fit into his plans for his master's study. On making enquiries, he learnt that no such course had ever been presented, so started a lobbying campaign. Though others in the department - Richard and Rick, at least - had some experience in the area, the buck stopped at me, probably because I was the most encouraging. (Or perhaps as a variation of the Peter principle ?)

I had previously taken part over several years in a project (largely in collaboration with George Blanchard of the Mechanical Engineering department, who has been a pillar of strength throughout the existence of the course) to develop a programming language and supporting system for controlling sequential processes such as assembly lines, and had thereby acquired a reasonable understanding of the field, but I was far from confident that I had an adequate basis for a stage 4 treatment. I therefore demurred at the suggestion that I should present a formal course in real-time computing, but agreed to supervise Steven in a reading course, which I hoped would both serve his purposes and give me a good enough view of the topic to decide whether or not I wanted to carry on with it.

Fortunately or unfortunately, Steven told his friends about it. Over the next few weeks a trickle of stage 3 students visited me, inquiring about this new course. Having a constitutional difficulty in saying no to people with perfectly reasonable requests, I ended up by painting myself into a corner from which the only escape seemed to be to present the course I had tried to avoid.

I had a very busy summer vacation.

R&R came to life, of a sort, in 1986. I have never claimed special expertise in the area, and have therefore always run it with a strong emphasis on teach-yourself. Most of the time it has worked fairly well. The enrolment has been around 10 each year; the 1993 class is the biggest ever. Every year there have been a few people who have been very enthusiastic, and have learnt a lot; there have usually also been a few who didn't really want to do it, but needed another four credits at stage 4, and haven't been enthusiastic at all. Most have done a creditable amount of work, learnt something (perhaps despite themselves), and passed.

PAST.

The "teach-yourself" emphasis requires experimental equipment, and somewhere to use it - which is what this paper is about. Of course, not everything is machinery; we can legitimately work on topics which only require computers - simulators (see below) and languages are prime examples - but the machinery must be there too.

For equipment, we began with a turtle, which was the extent of the controllable machinery available in the department, and a BBC microcomputer, which I'd acquired some years earlier from a research grant in connection with my real-time language work. Someone in the group wanted to build some hardware, so I brought in my son's electric train to give us something else to control. For the rest, we survived by grace of the Mechanical Engineers, as embodied in George Blanchard. He arranged for us to have access to their control laboratory, in which there were several desirable items of machinery with computers attached.

Against all odds, the course worked, but it became clear that we would need more equipment of our own to support it properly. (It also became clear that my son's train was not up to scratch, so I persuaded a mildly embarrassed head of department to buy me a new one. It has been very successful, and well worth the smallish sum we paid for it.) We continued to use the engineers' equipment for a few years, but access was always a problem : their students, reasonably enough, used the machinery from time to time during the day, and it proved politically difficult to arrange out-of-hours access.

Over the years since then, we have therefore acquired some more equipment. We have two more BBC machines and Acorn A500 machines, two small robots, a conveyor belt with some accessories which can be used with the robot to build a small workcell, two small analogue computers with which we can simulate a variety of systems with different dynamic behaviour, a pair of shaft-position encoders, a number of kits for "Stiquito" "robots" (small insect-like machines which depend on a memory-alloy wire for their operation and are controlled simply by switching on and off a 9v supply to each leg), and other small items which we've bought from time to time.

I have aimed always to get equipment which would illustrate principles rather than demonstrate the current state of industrial practice, and I think that's appropriate. The turtle is simple enough to understand easily, but is an adequate mobile robot with simple sensors and stepper motors. The robots are the simplest and cheapest which I could find with five axes and a gripper and feedback of all actuator coordinates; they have been ideal for the job. The train is both an exercise in simple motor control and in sequential processes, and has also worked well. Students learn a lot about interfacing and the low-level realities of control systems by building fairly primitive electronic gadgets to connect various sensors and controllers to computers. The analogue computers are still fairly new toys, but I intend that they shall give better insights into the principles of control theory and continuous control, and sidelights on analogue-digital conversion.

It had been my hope that over the course of time we would build up for ourselves more ambitious equipment and software as students built on the work of earlier students to construct more elaborate systems. This hasn't worked; in practice, students have almost always started from scratch again. I think that there are several reasons for this lack of enterprise, including poor documentation in the earlier assignment reports, inability to find earlier equipment and software because of my poor organisation of the laboratory, and the lack of any sort of atandardisation in the developments which means that one needs a thorough understanding of a probably idiosycratic design before one can safely do anything with it.

Maintenance has been something of a problem. The robot and the turtle, in particular, need better attention than they get, but I am well aware of the pressures on our vast technical staff. It is fair to say that the machinery keeps working because of the good offices of enthusiastic students, who choose to mend various bits and pieces as their assignments. Educationally, this is great - the students learn lots about the details of the systems they are mending - but it's probably illegal.

The illegality comes in because the government, determined to stifle enterprise at all costs, decrees that electricity is a nasty dangerous thing, so almost everyone must be kept away from most of it - for their own good, of course, which they are clearly incapable of working out for themselves. Regretfully leaving that line of argument for another day, I have therefore tried to give my students little bits of electricity by asking for low-voltage power supplies. These have finally materialised in small quantity.

A way to avoid overcrowding on the machinery is to have students work on simulators. This is all right up to a point, but doesn't work as well. They're too perfect, and don't break down enough. Some students like it, because it saves them having to learn about machinery; it could be argued that this isn't the aim of the course.

So far as space is concerned, there isn't a lot of significant history to recount. We have rubbed along so far in holes and corners. Before we had any of our "own" space, we overflowed into the engineeering laboratories. Since then, we've used the old "hardware laboratory" and most recently a corner of the "fishbowl".

Over the years, we have accumulated a vast library of assignment reports. Some are poor, most are not bad, a few are excellent. The students use them as background information and as a source of ideas for new assignments. All are records of some sort of work carried out by a student, and the great majority show evidence of careful thought invested in some aspect of control computing. All in all, they're impressive, and they are my evidence that the system works.

PRESENT (which is to say, 1993).

The basic pattern of the experimental work has served us well, and I intend to continue it, developing as and when it seems appropriate.

This year there are more students in the course than ever before. So far, we're managing, but by putting a number of students onto work with simulators. A major difficulty is maintenance : when this note was written, the multimeter and the oscilloscope were dead (which didn't make the analogue computers easy to use), the turtle was thought to be dead, and the older robot was fairly sick. Since then, some things have improved, but largely by the efforts of an enthusiastic student.

FUTURE.

The major need is for support. I remember that some time ago (perhaps when Peter-s came ?) we were asked to say what sort of technical help we needed; I said then that someone who could help with straightforward maintenance and simple repairs to the 473 equipment would be good. Peter is very good indeed - when he's available, which is hardly ever.

I am convinced that the approach of demonstrating principles is right, so I want to acquire a lot of bits which can be connected to computers, and if appropriate to each other, in lots of ways. I don't really want lots of specific devices which will talk to nothing but themselves.

I want people to be able to build bits of electronics if they so choose. There's little reason nowadays to go above 12v to do this (a few motors may need a little more), so that should be all right with appropriate power supplies.

I therefore want to aim for a laboratory which provides computers of various types (particularly more or less conventional microcomputers for "intelligent" control and programme development, and single-board computers or equivalent for local control functions), interesting pieces of machinery on the pattern of our present equipment (I almost wrote "on the lines of", but that's ambiguous), sensors, motors, etc., and - very important - interfacing gear.

To make this run, it needs power. One would expect that many self-contained machines will work from mains power, but for the rest low voltages should suffice. 12v, 5v, and increasingly 3v supplies are useful for the electronics; some machines may need up to about 20v. All supplies should be capable of giving several amperes. Obviously, all should be protected against overload.

There are at least three ways to do this, which are not necessarily mutually exclusive : provide separate power supply units (perhaps the most flexible); provide low-voltage power rails as laboratory supplies; or require all "loose" electronics to be constructed on standard cards which communicate through and draw power from some standard bus. I am not yet sure which is the best way to go, and there is certainly more to say about each of these possibilities, but I don't think that the answer affects the requirement for raw space.

To help in constructing circuitry, there should be basic tools, breadboarding facilities, soldering (or wirewrap ?) equipment, etc. I don't expect that lots of people will want to do this, but some gain a lot from it.

SPACE. (AT LAST.)

Raw space.

The basic requirement obviously depends on how many students are likely to want to do the course. History suggest 10; this year there are about 15. I shall assume for the sake of argument that we cater for 12. While not all students want to be there at once, most of those working on experiments need to leave the gear set up over a week or two at least, so there should be bench space for 12 experiments. The train is inevitably clumsy, but apart from that working spaces of size comparable to those in our other laboratories should be ample.

In addition, there should be storage space for equipment and documentation. Apart from the train, I see no reason why items current or future equipment should be much bulkier than the Mentor robot, and most is likely to be considerably less so. We have a chest of drawers which will accommodate much of the circuit-board-shaped equipment, if I ever get it in there; something of that style augmented by cupboard or shelf space for any larger items which are not used sufficiently frequently to leave out all the time would perhaps be satisfactory.

We have a cupboard which we use for documentation. At present, it's bursting, but that's largely because of the project reports; it's taken us about six years to fill it up, so something about twice the size should keep us going for a while.

Cooked space.

The space should be decorated in certain ways to make it useful. Each workspace needs mains electricity and some source of low voltage power, which could either be provided by separate power supply units or by reticulation from a central supply. Reticulated supplies may need to be augmented by local units where variable voltages are required.

One or two workspaces should be equipped with gear for light electronic, and very light mechanical, construction tasks. From experience, I don't think any more will be necessary.