

Setting Up and Running a Control Research Laboratory

Useful advice for the aspiring experimentalist.



©MASTERSERIES & EYEWIRE

By Dennis S. Bernstein

If you've spent most of your research career proving theorems and running computer programs, then you might find it satisfying to operate control hardware and experience first hand the challenge of applying control concepts to real systems.

Getting started on control experiments, however, can be an intimidating experience, and, like all challenges in life, you'll have to face the fear of failure. Without a doubt, some of your experiments will not work as you'd hoped, and you'll wrestle with the vagaries and idiosyncrasies of hardware and software. As a researcher you know that success arises from failure. So take heart and

forge ahead. In the end, you'll gain great satisfaction as well as insight into technology.

Setting up and managing a control research laboratory can be a time-consuming and labor-intensive undertaking. I'd like to share with you some helpful advice based on my own experience, much of which I learned by sometimes painful trial and error. I hope that this advice will help you avoid some of the difficulties that you might otherwise encounter.

This article is not about how to design a control experiment to investigate a specific research objective such as robust, nonlinear, or adaptive control. Some of the refer-

ences are more relevant in that regard. Instead, this article addresses the nitty-gritty challenges of developing a control laboratory and building control experiments that work.

The following remarks are directed primarily at the objective of developing a control research laboratory. However, many of these remarks apply as well to developing a control education laboratory.

Should You Build or Buy?

If you want to operate control experiments, your first decision will be whether to buy off-the-shelf (OTS) experiments or build your own (BYO). There are some obvious advantages to buying OTS experiments: they're already engineered, they're ready to run, and they're completely supported by the vendor, and thus they're easier to maintain. For a control education laboratory, OTS experiments are definitely the way to go. Furthermore, if you're not a mechanical, electrical, and software guru, then this is a good way to get started. When your OTS experiment stops working for whatever reason, be it electrical, mechanical, or software problems, then you merely call the vendor and let them diagnose and solve the problem.

On the other hand, there are disadvantages to buying OTS hardware. First and foremost is the fact that most of the engineering has already been done for you, and this will prevent you from being challenged by many of the design and engineering tasks that control engineers must confront. In addition, many OTS control experiments have been "cleaned up" by good engineering, and thus many interesting real-world physics and control challenges have been minimized, for example, by reducing friction and using low-hysteresis actuators.

More importantly, BYO experiments have the distinct advantage of forcing you to become familiar with the physics of the process, and this familiarity can provide insights into controller design and control theory. BYO also forces you to understand the restrictions and reality of engineering hardware in terms of sensors (resolution, bandwidth, linearity, drift, repeatability) and actuators (stroke, low signal bandwidth, high signal bandwidth, authority, linearity, hysteresis).

A more subtle point is that OTS experiments are designed for a preconceived experimental environment, and they're often difficult to reengineer to allow innovative control experiments. When you design your own control experiments, you'll have the opportunity to formulate your own experimental objectives. At the very least, your personally designed experiment will be unique. For these reasons, I believe that BYO is more desirable for graduate students engaged in research. However, it is important to keep in mind that OTS experiments can be set up much more quickly and are often less expensive than BYO.

If you do choose to build your own experiments, you'll want to be sure that the experiment provides an effective environment for controller implementation. In my experi-

ence, the most crucial elements are: 1) adequate sensor dynamic range (resolution divided by signal range) with good accessibility for feedback, performance assessment, and diagnostics; 2) adequate actuator authority (stroke, force, bandwidth); 3) sufficiently interesting dynamics in terms of nonlinearities and dimensionality; and 4) adequate processor throughput to implement sophisticated algorithms and signal processing. The first three objectives generally require careful engineering since it's easy to build a high-dimensional plant but difficult to adequately instrument it. With sufficient funds, the last objective is generally easy to achieve by purchasing a high-end turn-key control processor.

How Much Lab Space Do You Need?

The first thing you'll need is lab space. If you simply want to operate an OTS tabletop experiment, then you won't need much more than a desk and a chair. On the other hand, if you're designing and building your own experiments, then your need for space may grow rapidly to accommodate all of the infrastructure you'll need. For example, you'll need a workbench for soldering, drilling, and assembly; tables for PCs, test equipment, and your experiment; storage cabinets for tools and parts; and file cabinets for manuals and software. All of this adds up to a lot of infrastructure, and I haven't even mentioned special mounting tables (such as an optical table), electrical power distribution, water cooling, air supply and distribution, and Ethernet ports. Some experiments have environmental requirements on noise, vibration, temperature, and ventilation, which can be expensive to implement.

In most universities space is contested and in short supply. In fact, in some universities space usage is reviewed annually and must be continually justified. In other universities, space is assigned and rarely relinquished despite fluctuating need. If your experiment is relevant to ongoing work in another laboratory, you may be able to borrow space and save yourself the need to develop a lot of infrastructure.

My conceptual ideal for a lab is the gleaming, brightly lit, spacious laboratory of the villainous enterprises in the James Bond movies, where the sinister weapons work flawlessly the first time. In reality, a working laboratory is often filled with clutter and chaos. More on this later.

How Much Money Do You Need?

Obviously, this depends on what you want to do, but it's trickier than you might think. Over the years I've realized the importance of balancing the hardware cost with the manpower cost. I know some faculty who have had success with used equipment at a fraction of the cost of new equipment. However, I usually avoid used equipment unless it's fully documented and factory support is available. In fact, I've found that it's often difficult enough to get

new equipment to work properly, where the main problems are usually due to electrical and software glitches. The cost of a month's time of a graduate student trying to get a piece of new or used equipment to work is sometimes greater than the added cost of higher quality equipment, not to mention the delay to the project and the student seeking a degree.

Keep in mind that custom components are often extremely expensive due to nonrecurring design, setup, and documentation. This cost is an excellent reminder that much of the economic benefit we get from engineering is due to volume.

There are also unexpected costs. For example, when you're trying to choose components, it's sometimes necessary to purchase items from several different vendors and simply test and compare them. Manufacturers' specifications are essential to get you started, but they rarely address the specific circumstances of your application. The necessary testing may be time consuming and may require the purchase of mechanical and electrical test equipment. Another expense that's hard to avoid is the continual update of software. Control experiments typically involve software from multiple vendors, and updating one program generally necessitates updating others, in a kind of recurring wave.

Finally, there are the unexpected costs of repairing equipment and components that fail. Failure can occur for mechanical, electrical, or thermal reasons. Dropping delicate equipment and reversing power leads are sure ways to damage equipment. Drawing too much current can easily destroy circuits and melt motor components. I always ask vendors about the ruggedness of their product to mild and severe abuse. What never ceases to amaze me is the ability of equipment to fail even when it's not in use. High-tech equipment often has a shelf life at least partly due to the cyclical departure of graduate students who are often the only ones who know how to operate it.

What's the Best Way to Design a Control Experiment?

When not attending meetings, a lot of what engineers do in practice is analysis, especially error and failure analysis. However, what engineers thirst for is the creative act of *design*, which is surely one of the most satisfying of all human activities.

Books have been written on design, and it's definitely an art. On the other hand, a scientific approach to design is beneficial and can save enormous amounts of time, expense, and effort. I'll give a few unscientific observations about design based on my own experience and the advice of my colleagues in academia and industry.

When I design anything, I work backwards from the constraints, which include performance requirements, money, time, weight, volume, power, and environmental effects. If

something must fly on a spacecraft, then weight is usually a driver; otherwise, weight may be a parameter that I can squander. If something must operate in heat, cold, under water, or on a vibrating surface, then I have to worry about those effects; otherwise, I don't bother with them. If I don't have hard performance specifications and money is limited, I merely optimize the specifications subject to the funds I have available.

Once I decide which components I need, I must decide how much engineering to do myself and how much to contract to others, thereby allowing me to focus on issues more directly related to control objectives. However, it is nontrivial to find companies to build exactly what you want. I've found that the willingness of a company to do custom work is inversely proportional to the size of the company. In dealing with a small company you can build relationships with engineers who will often take the time to help you after the product has been delivered. In addition, small companies often appreciate the publicity that universities can provide. On the other hand, some companies are wary of the support that universities may require. Don't be surprised when custom engineering takes longer than everyone expects.

No matter how hard you try, you can't design something perfectly the first time. The 747 is a great airplane, but the manufacturer had a lot of practice building other planes. Ideally, you would have several chances to design and build something. The first design would allow you to check out the major functional issues. The second design would allow you to correct the initial flaws and get the system working acceptably. The third design would allow you to refine the fabrication and packaging for implementation. Just before you complete the third design, however, your understanding of the whole system will be so much greater than when you started (not to mention the fact that the available components may have improved) that you'll want to begin again from scratch to implement vast improvements. Of course, if you don't resist this urge you'll never finish what you set out to do.

If you have only one chance to design and build something, then you must face the fact that you simply can't think of every possible difficulty that can arise. Therefore, it is helpful to take some precautions. For example, if a part must mate with a component provided by a manufacturer, then it may be helpful to have the component on hand to verify its specifications and geometry. Sometimes you must choose multiple components at the same time and with incomplete information. When this happens you may need to guess and hedge your bets or else you'll end up in design gridlock.

My research group has had problems with assembly and operation due to overlooked details as small as the size of the head of a screw. We've also had to fabricate tools of odd size and shape to assemble mechanisms.

You can't be too careful making drawings for a machinist. An error can be costly if a part needs to be remade. A hole can be made larger but not smaller, so triple check. I've had drawings interpreted upside down, so be wary of parts with symmetry. An experienced machinist can suggest ways to improve an experiment's design (easier and less costly to manufacture, easier to assemble, etc.) So develop a relationship with them, tell them how the parts will be used, and how they will be put together. All of these comments also apply to working with an electronics technician.

When choosing materials, steel and aluminum are the most convenient. For the same part, steel is about three times as heavy as aluminum and about three times as stiff. Whenever possible, I like to make parts out of polycarbonate plastic, which looks high tech but cannot be machined as precisely as metals.

One of the most overlooked aspects of control experiment design is the cables. It's worth the extra effort to design these with special care since they will be the most suspect part when your experiment doesn't work. When ordering custom equipment, be sure to specify cable lengths, the type of connectors, and the connector gender (male or female). Give unique and descriptive names to all cables, connectors, and pins. Beware of variations in nomenclature. For example, most manufacturers make optical encoders with standardized input and output channels, but different manufacturers use different signal terminology. Excessive cable length can usually be bundled (although it may be unsightly or a source of noise) but a cable that is too short is harder to fix. Physically label all cables and just about everything else you can think of. A label maker is a great laboratory investment.

Stress on cables and connectors is a constant source of problems. Any cable attached to a moving component is a potential problem since the cable and connector can get stressed and fatigued. BNC connectors, which have a locking mechanism, are good for single wires, and Dsub's are great for multiwire cables. Flimsy connectors like Molex are to be avoided whenever possible. Use shrink wrap to cover exposed metal at wire connections.

Design mechanical mounts such as brackets to accommodate cables. Once I was so determined to avoid unsightly, dangling wires that I designed a support bracket so that the wires could be routed inside the bracket. Keep in mind that cables can be passed through small holes but many connectors cannot.

Today, wireless links can be used for signals (but not power). Wireless links eliminate cables and are a boon whenever you can afford to put the necessary electronics at both ends of the link.

There will always be components that are hard to mount or keep stationary. Therefore, some of your best friends in putting the final touches on an experiment are

Velcro, cable ties, and foil (not duct) tape (but beware that the surface of foil tape is conductive).

How Much Should You Worry About PCs and Software Maintenance?

PCs are cheap nowadays, but it's important to keep in mind that the software may be more expensive than the computer. Many instruments such as oscilloscopes exist in stand-alone or virtual (PC-based) versions. The PC version sounds great, but you must maintain the software, and you may have to dedicate a PC to the instrument. However, laptop-based instruments are portable and thus convenient for data acquisition and testing outside the lab.

As I already mentioned, it's a constant battle to maintain compatible, updated software. Don't overlook the importance of carefully organizing and storing disks and manuals. PCs can crash, and then you may need to reload the entire machine. Furthermore, as hard disks fill up, you may need to install new drives and reload the operating system and software. These tasks are not a major problem if you're dealing only with your personal machine, but it's a continual problem if you have to maintain several lab PCs. Remember that computers have a tendency to consume money with each software update. Finally, it helps to assign responsibility for each machine to a particular student. Maintaining each machine in terms of virus protection, memory management, and other tasks requires constant care and attention.

What Electrical Problems Should You Look Out For?

Unless you're an electrical guru, you'll find yourself, like me, continually amazed by the subtleties of circuitry. It's helpful to keep in mind that your system will generally involve power circuits, with high volts and amps, intermeshed with signal circuits, with low volts and amps. The trick is to wire everything in such a way that these circuits don't interfere with each other.

Good grounding techniques are especially important to avoid sending excess current through sensitive components and to prevent noise generated by power circuits from affecting signal circuits. It's also a good idea to protect delicate and expensive components by putting fuses, capacitors, and diodes in their circuits. This practice is especially helpful if components such as amplifiers are drawing lots of current that might get discharged into the wrong circuits.

Signals can be read in reference to ground or differentially. Differential measurements are more immune to noise, and they're useful if your system doesn't have a common ground. Optical links are helpful for isolating signal circuits from power circuits to reduce noise. Unfortunately, you won't know how much noise corrupts your signals until you test your system after it's built, although careful engineering can improve the outcome.

Be aware that all circuits have nonideal input and output impedances, which will affect performance when subsystems are interconnected.

Finally, all circuits will cause signals, whether viewed as inputs or outputs, to have a bias. Ideally, a robust controller will not be affected by a dc voltage offset, but this isn't always the case. A dc voltage offset can drift due to grounding or thermal effects. Hence, it may be useful to determine the thermal drift of your equipment. This knowledge may be important in nonlinear systems when the operating point shifts, which can make feedback linearization difficult.

Getting started on control experiments can be intimidating, and, like all challenges in life, you'll have to face the fear of failure.

Is It Easy to Find Good Filters Cheap?

As control engineers, we appreciate the importance of filters to remove noise and anti-alias signals. Digital filters can be implemented in software, but only an analog filter can truly provide anti-aliasing. Amazingly, there are remarkably few sources for good analog filters at a reasonable cost. The lowest I've seen is about \$600 per channel. You can build your own from filter chips costing less than \$20 each, but devising low noise circuits is a nontrivial engineering task. Battery power is helpful to avoid power conversion noise.

When you buy filters, first decide whether you need low or high pass; some instrumentation filters can be operated in either mode. Many low-pass filters do not actually go all the way to dc (in jargon, they are ac coupled). Some filters are purely analog, while others are hybrid; that is, they use a combination of analog and digital circuits. A few well-known filter vendors are Ithaco, Alligator, and Krohn-Hite.

Also note that many signal amplifiers (that is, low-current amplifiers for sensors) have built-in filters. Since you'll probably need a signal amplifier anyway, this may be a good way to get the filtering you need. However, these filters tend to have a fixed bandwidth, so you can't change the rolloff frequency for control studies involving different bandwidths.

Don't forget to include your analog and digital filters in your plant model. These transfer functions can adversely affect closed-loop behavior due to phase shift. If you happen to implement either an analog or digital filter with a nonminimum phase zero, then you'll be limiting your closed-loop performance just as if your plant had such a zero (which it effectively will). Of course, when you identify a plant with filters for signal conditioning, your identi-

fied model will include the filter dynamics. This is one reason that it's desirable to use the same hardware for both identification and control.

Is It Easy to Find a Good Power Amplifier?

For almost all kinds of actuators, your amplifier is your most critical component. There are two main types: PWM (pulse width modulated) and linear (continuous signals). Your amplifier may be designed to supply the voltage you command, or it may be regulated (transconductance) to provide a commanded current. Most amplifiers require careful study to operate well. Manufacturers' data sheets and schematics are helpful but are often cryptic, especially to the novice. Like filters, some amplifiers such as audio amplifiers will go to low frequency, but not to dc. If you wish to provide a constant voltage or current offset for your actuator, make sure your amplifier is dc coupled. Check for bandwidth and phase shift as well. Also, the performance of many amplifiers is sensitive to the inductance of the load. In summary, don't treat your amplifier like a black box.

The settings in many amplifiers require the insertion of resistors and jumpers, while some models can be configured using software. In fact, some of the newest amplifiers (such as the Copley Accelus model) can self-tune proportional-integral-derivative loops, reducing the need for control engineers!

How Much Should You Worry About Batteries?

Sometimes you don't have the luxury of plugging your experiment into the power grid. This happens, for example, if your system is a vehicle that can't be tethered, in which case you may need to run it on batteries. You'll be surprised how many things you need to worry about when your system depends on batteries. First, there are many different kinds of batteries to choose from (lead acid, nickel-cadmium (nicad), nickel-metal-hydride, lithium) so you need to understand the tradeoffs. Critical specifications include cost, weight, and volume, as well as voltage profile during operation and peak amp capability. Some batteries, such as lead acid batteries, can be damaged if they're discharged too quickly by drawing too many amps, while other batteries, such as nicad batteries do not mind fast discharging. Equally important are maintenance issues. Some batteries can be damaged if they're recharged too long or too quickly, or if they're discharged too low or too fast. It's helpful to think of batteries as living entities that need to be constantly cared for. A sophisticated battery charger is a good investment.

Despite these difficulties, batteries have an important advantage over line power, namely, they're free of the noise that arises from ac-to-dc conversion.

How Well Should You Document Your Lab Results?

Ideally, every student should maintain a lab notebook to document their progress. Lab notebooks are of extreme importance traditionally, but tend to be overlooked in the computer age. Instead we have stacks of printouts all over the lab. It helps to have a supply of printer ink, paper, a puncher, and three-ring notebooks to organize documents. These supplies are essential as students join and leave a project, since otherwise you'll have little continuity. A digital camera is handy for documenting experimental setups for reports and papers.

Which Vendors Are Good to Deal with?

I've already pointed out that smaller companies tend to be more receptive to university business, which is important in procuring custom components. For OTS parts I've had good luck with: Dell for PCs; MCM for connectors, wire, and electrical parts; DigiKey for specialized electronic components; Grainger for industrial-type equipment; Copley Controls for amplifiers; McMaster-Carr for hardware; and Precision Industrial Components for specialized mechanical parts. It's essential to find suppliers that provide the quality, service, and price that works for you.

Always check availability and lead time. Electronic chips can quickly go out of stock. If you unknowingly design your system around a component that's been discontinued, then you'll waste a lot of effort.

Compare prices when you have the luxury. Beware that some companies have minimum order amounts, but even when that is not the case, shipping can cost as much as what you order when you need only a few parts. For each company you normally deal with, it helps to keep a running list of "things to get" that are not time critical. So, when you find you need something in a hurry, you can clear off your list and reduce the cost of shipping relative to the cost of the components.

How Much Should You Worry About Lab Safety?

Last, but most important, is lab safety. A laboratory is a place of constantly shifting layout, with unusual configurations that are not addressed by safety regulations. That's why you must take extraordinary precautions for safety. Anything that even remotely appears dangerous should be given serious thought. The dangers of high currents, high-speed moving parts, high-power lasers, and high temperatures should be given special attention.

When you run your experimental control systems, you'll realize that much of the engineering that goes into OTS experiments is intended to protect the equipment against damage and the user against harm when an unstable control system is implemented. With BYO experiments, you'll have that complete responsibility. It's doubtful that you can predict every possible failure mode, and, hopefully, when failure does occur, the cost of repair is the only issue.

So, When Should You Get Started?

When I started building a control lab, I was completely intimidated. Luckily, a colleague helped get me started, and it wasn't as hard as I thought. The important thing to keep in mind is that, as a control engineer, you can learn a lot from even the simplest hardware, which will be noisy, nonlinear, uncertain, and you name it. Get yourself a good control processor with reliable software, build yourself a simple plant, collect some data to identify the system, and close some loops. You'll never look back.

Acknowledgments

I would like to thank Carl Knospe and Rob Clark for numerous helpful suggestions.

References

Some perspectives on control experiments are given in

C.R. Knospe and E.H. Maslen, "Meaningful control experiments," in *Proc. Amer. Contr. Conf., Invited Session on Control Experiments: What Do We Learn from Them?*, Albuquerque, NM, June 1997, pp. 2703-2707.

D.S. Bernstein, "Control experiments and what I learned from them: A personal journey," *IEEE Contr. Syst. Mag.*, vol. 18, pp. 81-88, Apr. 1998.

D.S. Bernstein, "A plant taxonomy for designing control experiments," *IEEE Contr. Syst. Mag.*, vol. 21, pp. 7-14, June 2001.

Some ideas for designing control experiments can be elicited from the discussion in

D.S. Bernstein, "What makes some control problems hard?," *IEEE Contr. Syst. Mag.*, vol. 23, pp. 8-19, Aug. 2002.

There are many good books on laboratory techniques and instrumentation such as

J.H. Moore, C.C. Davis, and M.A. Coplan, *Building Scientific Apparatus*. London: Addison-Wesley, 1983.

J.P. Holman, *Experimental Methods for Engineers*, 7th ed. Boston, MA: McGraw-Hill, 2001.

J.A. Blackburn, *Modern Instrumentation for Scientists and Engineers*. New York: Springer, 2001.

For issues relating to sensors, see

D.S. Bernstein, "Sensor performance specifications," *IEEE Contr. Syst. Mag.*, vol. 21, pp. 9-18, Aug. 2001.

and the references therein.

Dennis S. Bernstein can be contacted at the Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109-2140, U.S.A., dsbaero@umich.edu. 