

# Programming from the Ground Up in Controls Laboratories Using Graphical Programming<sup>‡</sup>

Todd D. Murphey

Jeannie Sullivan Falcon

Electrical and Computer Engineering National Instruments

University of Colorado

Austin, TX, USA

Boulder, CO, USA

## Abstract

This paper describes a recent, innovative educational approach taken at the University of Colorado for teaching a senior-level undergraduate controls course. By using software written in a graphical programming language (in this case LabVIEW<sup>TM</sup>) to control the hardware, the students were able to do all the programming themselves. The advantages of this approach over “canned” lab approaches are many. The students feel more responsible for the final product and they are able to apply control techniques learned in class in a more fundamentally creative way. We present evidence from the class taught in Autumn term, 2005 and 2006, that suggests that students will often use tools from beyond the scope of the class to solve problems in creative and occasionally unexpected ways.

**Keywords:** Education, Control Education, Laboratory Education

---

\*This paper was original presented at the IFAC Advances in Controls Education Conference and was consequently recommended for IJEE publication.

<sup>†</sup>This work was supported by National Instruments, the National Instruments Foundation, and the University of Colorado.

# 1 Introduction

This paper describes an innovative approach to teaching a controls laboratory aimed at upper division undergraduate students. We use a less structured approach to the laboratory than is typically used. In particular, the students are in charge of all programming, the labs are more open-ended than typical controls labs, and the students work in rotating groups (which forces them to incorporate “legacy” code from previous groups). This course structure gives the students a sense of the difficulties and ambiguities often associated with solving real problems. The basic thesis of this paper is that if one uses a graphical programming environment, a course can be designed to give students much more control by allowing them to do all the programming and giving them more open-ended assignments. The consequence of this is a deeper understanding of the application of controls to real problems, which we illustrate using in-class assessment and through the students’ ability to generalize what they have learned in class.

Traditional control courses often involve pre-written, or “canned,” software (often written in an intermediate-level programming language like C or C++). Moreover, partially because of the constraints imposed by the pre-written software, labs are highly regimented. This choice is made largely because it is not feasible for the students to write all the code—there simply is not enough time during a one-term course for students to cover as much material as we would like as well as do all the programming themselves in an intermediate-level language.

By using a graphical programming environment, however, it is possible to create a course where students do nearly all of the programming themselves. This prepares them well for jobs where they will be expected to know how to integrate software and hardware, and it additionally helps them learn controls concepts more deeply by allowing them more flexibility in exploring control designs. Moreover, this course structure has other advantages as well. In particular, it is possible to give students more open-ended labs (that do not include step-by-step instructions), which is typically more engaging for the students. This also leaves more room for creative application of concepts the students are learning.

The decision to let students write all of the code was inspired by noticing that graphical programming environments (such as LabVIEW, Matlab/Simulink<sup>TM</sup>, and others) are faster environments for programming. We used LabVIEW for all laboratory development and students used LabVIEW for analysis, simulation, and experimental implementation. (For a useful, if not comprehensive, introduction to LabVIEW, see [2].)

Over the course of the term, the students learned to write all of their own code, both for simulation and for hardware experiments. They quite literally started with blank code, or would incorporate student-created code from previous labs. They used LabVIEW in a manner similar to that reported in [1]. However, instead of the instructor and teaching assistant writing the code for the students, the students were given very short tutorials on graphical programming and wrote the code themselves.

Assessment of the course is addressed using two modalities—student questionnaires and generalization. These two assessment modalities are both appropriate (and common) for small classroom sizes, but are nevertheless not statistically significant. The use of generalization as a measure of assessment can be particularly useful when looking at small numbers of students, and has already been used by the first author in [8]. Moreover, as a consequence of the present work, we have obtained funding from the National Science Foundation for formal assessment of this laboratory teaching strategy using an Expert/Novice criterion that is appropriate for small class sizes [9]. This is the main focus of our future work.

This paper is organized as follows. Section 2 discusses the hardware setup, the lab organization, the tasks the students were expected to complete, and the general level of background students had coming into the course. Section 3 has evidence that this structuring of the class led to more internalization of controls techniques as well as more creative approaches to open-ended problems. It additionally describes student reactions to the course, including course evaluation results. Section 4 discusses the necessary trade-offs made when teaching a course using the described approach. Section 5 discusses the various advantages and disadvantages to the instructor using this method. We end with conclusions in Section 6.

## 2 Lab Organization

The Electrical and Computer Engineering Department at the University of Colorado offers two senior elective control courses: Control Systems Analysis (ECEN 4138) and Control Systems Laboratory (ECEN 4638). Both courses typically have 20-30 students. This last autumn was the first time in the last several years the laboratory was offered using hardware, made possible through funds provided by the university and National Instruments. Typically students take both courses concurrently, although they may take the lab if they have taken the lecture course previously.

The goal of the class was two-fold. First and foremost, the class aimed to reinforce the students' understanding of introductory controls concepts—Modeling, PID, Root Locus, Bode, Nyquist, and state-space techniques. A secondary goal of the class was to reinforce the fact that the models we write down do not *actually* reflect the true dynamics of a system; they only approximate them. We wanted to see students discover control designs that work in principle but fail in practice. The more students are in charge of discoveries of this sort, the more we expect they will retain key concepts past the class.

All the students worked in randomly selected, rotating groups of three people each. These groups changed five times over the course of the term. This structure forced the students to combine and use legacy code from each member's previous group. One of the byproducts of this was that by the end of term, all the students had very good, stable code for running experiments. Moreover, they typically understood the code and could replicate it as needed.

The course included the following topics.

1. Modeling
2. System Identification
3. PID tuning
4. Lead/Lag Controllers

5. Root Locus Analysis

6. Bode/Nyquist Analysis

Moreover, the students did a final project.

## 2.1 Student Background Prior To This Laboratory

Students were predominantly seniors with a standard electrical engineering background. All but one student was taking the lecture course (ECEN 4138) concurrently (the one student had taken it a year previously). Most had never seen any more controls than a PID controller.

Almost none of the students had any graphical programming experience, but several had seen Matlab/Simulink and LabVIEW in demonstrations in other classes. Hence, students view these languages more as interfaces than computer languages. Because of this, the students had somewhat of a predisposition against using the software, but this was overcome by pointing out that it would be nearly impossible to have them write all the code in C or C++.

## 2.2 Hardware

The experimental set-up used in the laboratory was comprised of a torsional disk system from Education Control Products, seen in Figure 1. These experiments are relatively robust, which was particularly important given the level of control we gave the students. The software used was LabVIEW 7.1 (as it has been used in other classrooms, see [1]), the Simulation Module 1.0 [6] and Control Design Toolkit 2.0 [5] that both run as part of LabVIEW. Our input/output capabilities were all provided by a National Instruments FPGA (NI FPGA 7831R) [4, 7]. The students modeled the system, performed system identification, applied the basic techniques of PID tuning, Lead/Lag, Root Locus, Bode, and Nyquist analysis. The course ended with a final project and competition.




Figure 1: The ECP Model 205a Torsional Plant (photograph taken from the manufacturer's website <http://www.ecpsystems.com>)

This system has a  $2 \text{ N} \cdot \text{m}$  DC motor at the bottom that drives the first disk. Two more disks are connected via a torsional spring that allow for torsional displacement between disks with an associated torsional spring constant. Thus, the dynamics of this system are sixth-order, making this a nontrivial control system. Moreover, weights (also seen in Fig. 1) are included that allow the user to change the inertia of each disk. This system is basically linear, making it a good choice for a linear controls laboratory. The biggest downside to this experiment is that it is neutrally stable, unlike other classic control systems such as an inverted pendulum. (For a description of using these plants with LabVIEW, see [3].)

### 2.3 Programming From the Ground Up

As previously mentioned, the students wrote all of the code for the course. Hence, for simulation and analysis, the students wrote their own simulation code from scratch and integrated analysis tools (such as root locus, bode, and Nyquist plots) into their simulation. With the exception of some basic hardware safety (such as turning the system off if the torsion between two disks became too high), the students also wrote all of the experimental

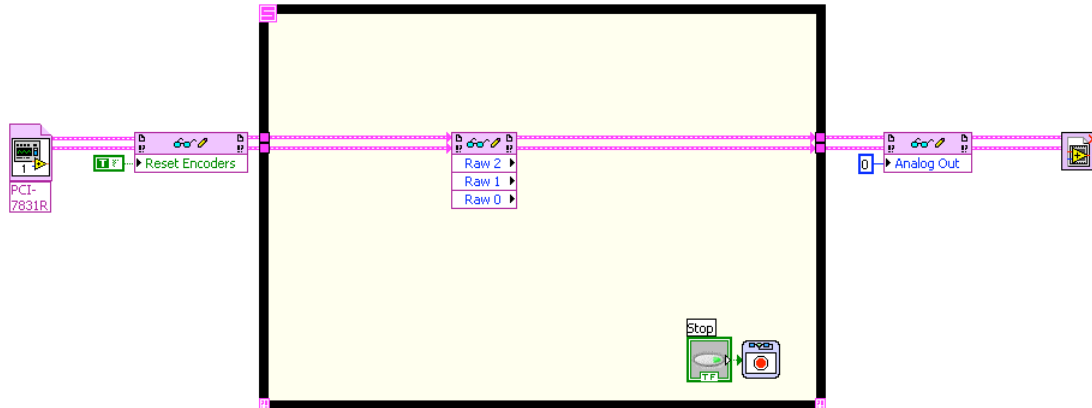


Figure 2: One of the templates the students used for communicating with the hardware. It provided the signals going to the DC motor and coming from the encoders, as well as a halt feature to gracefully reset the FPGA if needed. This was all the students were provided with—everything else needed for the laboratory they created themselves.

software. The communication and hardware safety was given to them in the form of a template, an example of which is seen in Fig.2.

Labs were generally open-ended. They focused on problem solving rather than following instructions. In particular, all labs were divided into high-level tasks (known in the education community as “authentic” tasks), most of which were not mathematically defined for the student, described next.

## 2.4 Examples of Assigned Tasks

In designing the tasks for this lab, we were trying to avoid the more traditional approach of giving students a long series of steps to follow. Instead, we tried to give them tasks that were more similar to verbal tasks they might be asked to accomplish in a job, but that nevertheless required the technical detail and insight generally learned in the lecture course.

Examples of assigned tasks included the following:

1. Model the ECP system as a single-input-single-output system first using the bottom disk as an output and then using the top disk as an output. Simulate this using both the symbolic transfer function and symbolic state-space model representations

in LabVIEW. (Note to reader: The sample code given to students regarding how to create state space models in LabVIEW consisted of what is contained in Fig.3—more than enough to solve the problem, but in no way a “canned” solution for them to use.)

2. Design a controller that is stable for all possible locations of two weights on disk 3. Is it possible to do this and guarantee a rise time of 0.5 second?
3. Run the experiment and your simulation simultaneously. What are the differences? If there are differences, where are they coming from and can you fix them?
4. Consider the system with the bottom disk as an output and with the top disk as an output. Which is “more” stable using a PID or Lead/Lag controller? Why? You may approach this problem in whatever way you find most convenient—just detail it in your writeup.
5. Plot the root locus for the parameter you think introduces the largest amount of uncertainty. You may wish to “lump” some uncertainties together. Does a parameter variation necessarily give a valid root locus?<sup>1</sup>
6. We know that using a PID controller with the top disk as the output requires a very low gain on the controller. Design a higher performance controller that allows you to push the overall system gain up higher while not destabilizing.

The labs generally became shorter and shorter as the term progressed and the students required less and less guidance in the labs. This, it seems, is a good indication that the students were becoming more competent. By the time the end of term drew near, and it was time to do the final project, it was possible to give the final project using the following four bullets:

1. make the bottom disk go from 0 to  $\pi$  and stabilize back to 0 as quickly as possible,

---

<sup>1</sup>The answer to this is “no” because the root locus may have higher order dependencies, depending on how the parameter enters the transfer function.



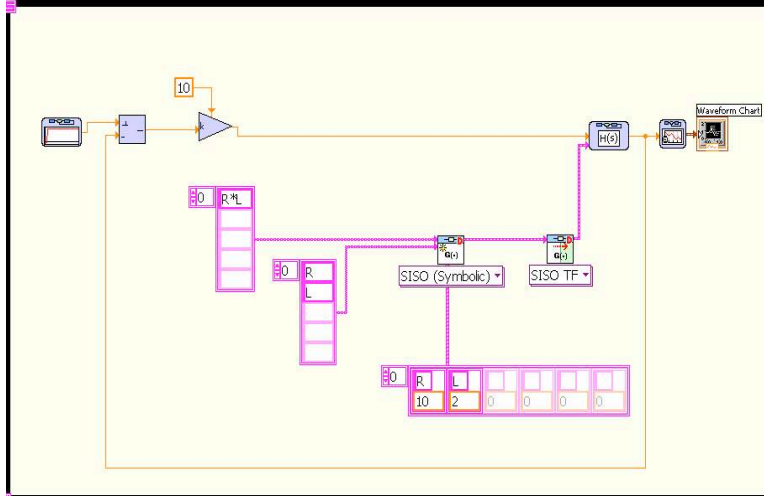


Figure 3: Sample code given to the students to help them in creating state space models in LabVIEW.

2. accomplish the same task for the top disk,
3. do so without breaking the ECP unit,
4. and do so knowing that the instructor will move the weights up to three centimeters away from their nominal position on the ECP machine.<sup>2</sup>

### 3 Classroom Assessment

There were several different elements of this endeavor that we consider substantial successes. These include the increasing quality of computer code the students were writing, the improved understanding of block diagram representations (that was a clear consequence of having to create block diagrams repeatedly while programming in the graphical environment), and the creativity shown in the final projects. By the end of the term, the students were treating control problems as *actual* problems to be solved, and were using an appropriate balance of physical intuition and analysis skills from the control analysis lecture course.

Often, because of the lack of structure in the teaching of this laboratory, students would come up with controller designs that worked in principle but not in practice. This, too, is a

<sup>2</sup>Hence, introducing uncertainty into the system.

valuable aspect of the open-ended nature of the way the laboratory was designed.

### 3.1 Final Project—Generalization of Concepts

This section describes anecdotal evidence, based on this one class, that the “programming from the ground up” teaching philosophy does indeed have positive consequences. As noted above, the final project was very simple. Students were to make the bottom disk go from 0 to  $\pi$  and stabilize back to 0 (within 2 degrees or 0.035 radians) as quickly as possible and then were to do a separate design with the same goal for the top disk. The students then competed on the final day of class for who had the best performance for the bottom disk and who had the best performance for the top disk. In addition, the students knew that on the day of the competition the instructor would move the weights to introduce some plant uncertainty (although not of a particularly malicious type).

Using the bottom disk as the output, of course, did not present much difficulty. All the students knew that they should go with some version of a PID or Lead/Lag controller. However, using the top disk as the output generated surprisingly creative approaches. This is because the top disk is the “noncollocated” problem (where the actuation is not located at the same place as the sensing), which is substantially more challenging than the collocated problem. Moreover, we believe that some of the approaches students used would have been impossible to expect in a class that gave the students “canned” software for control design.

A few notable designs (using the top disk as the output) by students:

1. The noncollocated problem has a transfer function with six poles and no zeros. One student tried a controller that had five poles and no zeros. This controller worked in simulation (with terrific performance and stability) but did not work on the experiment. He discovered that the controller was requiring  $10^4 N \cdot m$  while the DC motor we had could only provide  $2 N \cdot m$ . Fortunately, he introduced a saturation into his simulation after this.

2. Some groups decided that despite the fact that *we had not covered state-space control in the laboratory*, they would use a state-space design for the final project. They learned LabVIEW's tools for doing so and used both an LQR and pole placement technique for their designs.
3. Another group also chose to use a state-space design. However, they noticed that in fact the ECP unit has all three outputs available. This, they reasoned (correctly), implied that they could get a much better estimate of the system using all three outputs rather than getting an estimate only using one output and using a sixth order estimator. Their performance ended up being nearly twice as fast as any other group. This is an excellent example of the groups choosing a perfectly valid control strategy that creatively applies what they have learned in class and in lab, but nevertheless goes beyond the scope of what they have explicitly learned. Moreover, this strategy would have impossible to implement, much less conceive of, in a more structured lab where all the code was provided a priori.
4. The second best performance came from a group that realized that the top disk was passively stable with respect to the bottom disk. Hence, they chose to use a PID controller on the bottom disk and allowed the natural dynamics of the system to stabilize the top disk. Hence, they were using a form of passivity control (certainly without ever even having heard the phrase). There is no question that this solution was somewhat outside the *intended* approach, but this does not mean that it does not reflect real insight and learning on the part of the students.

## 3.2 Student Reactions

Students were asked to fill out a traditional survey, where they ranked items related to how much they perceived they had learned, as well as other aspects of the class. Some of these questions and student responses, are in Table 1. Basically students almost uniformly

	Question	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
1	I gained valuable hardware experience in this class	0	0	1	10	7
2	The labs were appropriately open-ended	0	2	5	8	3
3	I would recommend this class to someone looking for a laboratory class next year	0	0	0	7	11

Table 1: Table of student responses. Students were asked how much they perceived they had learned and asked how much they agreed with some statements about the effectiveness of the course. Notably, even though some of the students did not *like* the open-ended nature of the laboratories, they nevertheless uniformly recommended the course to other students.

believed they had learned more in this class than they had in an “average” course, and moreover uniformly recommended the course for others.

The students reacted positively to nearly all portions of the course. Our university administers mandatory course evaluations at the end of every course, with the most important metrics being “Course Rating” and “Instructor Rating.” In the Autumn of 2005, the students gave the course an average overall rating of 3.3 out of 4.0 and an instructor rating of 3.6 out of 4.0. After the second time that the course was taught in the Autumn of 2006, the students gave the course an average overall rating of 5.2 out of 6.0 and an instructor rating of 5.9 out of 6.0. (Our university changed its administration of the rating system between the two years.)

These course ratings can also be examined in comparison to the averages for the College of Engineering (COE) in the University of Colorado at Boulder. The course ratings were in the 79th percentile for the department and college, while the instructor ratings were in the 92nd percentile for both the department and college. Thus, the course was uniformly above average in its ratings both years. Although this does not indicate how much the students actually learned, it does indicate that the programming from the ground up approach did

not alienate the students.

Students specifically mentioned the fact that programming was a valuable experience (partially because many of them were asked in job interviews whether they were familiar with graphical programming techniques). Moreover, the majority of students felt that they had learned control techniques that they would feel comfortable applying in realistic settings. (One student (out of 22) complained that the labs were not “industrial” enough.) Roughly two-thirds of the students mentioned that the open-ended labs were more interesting and fun than previous lab courses they had taken.

The only part of the course that the students did not like was the rotating, assigned groups. This is because of the fact that they had to incorporate legacy code into their projects at the beginning of each new rotation. Although they viewed this negatively, the authors feel that the uniform improvement in the quality of code over the course of the term is an indication that this was a useful part of the course. Therefore, despite the fact that students disliked the practice, we will continue to have rotating, assigned groups in future versions of the course.

Lastly, as part of our future work in this course, we have obtained funding from the National Science Foundation for formal assessment of this course and the pedagogical style taken here.

## **4 What is lost in this approach**

As with all choices in teaching, some things were lost because of having the students write their own code in the graphical environment. One of the most significant of these losses is that the students never had to concern themselves with digital control—LabVIEW took care of all the translation of continuous time design into digital control on the FPGA.

Moreover, the students did not learn anything about the embedded systems (in particular, the FPGA) they were using. This is because LabVIEW automatically compiled all the code

they were writing to the FPGA. Whether or not this is actually a negative or a positive aspect of the class is up for debate, particularly since the course is not aimed at embedded systems. Such a course would, in all likelihood, not allow students to actually implement much more than a PID controller by the end of a one term course.

Another downside is the fact that there was only one experiment. The motivation for this was that the students had to derive all the equations of motion themselves, and this was reasonably time-consuming. Moreover, obtaining and maintaining additional hardware units creates added expense. However, there were several times during the course when the laboratory was substantially ahead of the lecture course. Therefore, what is needed are substantially simpler (e.g., lower-order) systems that are reasonably easy to model but that have different properties from the ECP unit. In particular, we would like to use a plant that is unstable, such as an inverted pendulum.

## 5 Pros/Cons of Teaching This Way

This method of teaching is easier for the instructor prior to a course beginning and harder once the course has started. The instructor *must* be competent at programming in the chosen graphical programming language. This is because the students are actually learning not just how to program in the graphical environment, but they are simultaneously improving their more generic programming skills. Another advantage to having students do the majority of the programming is that software upgrades are much easier to integrate into the classroom because there is little to no legacy code from previous versions of the course.

Another aspect of a course like this is that students get to see the instructor solving problems in class. Although in principle this can undermine the instructor's authority, the first author's experience teaching this class is that the students generally appreciated the spontaneous nature of the interaction, even when it occasionally produced enigmatic results that were not immediately resolvable.

## 6 Conclusion

This paper describes an approach to teaching a hardware-based controls laboratory that allows students to be in charge of all the programming required for the course. This was made possible through the choice of a graphical programming environment, specifically LabVIEW 7.1, though any graphical programming language would accomplish the same goal. We found that students were more involved in the course and offered more intrinsically creative solutions to problems than they would in a course where the software is pre-written for them. Additionally, the laboratory used open-ended labs, which was facilitated by the fact that the students wrote the software themselves. We are going to revise the course again to involve more experiments, to cover state-space methods, and to make the assignments more open-ended. Moreover, we are going to engage in more formal assessment of the course, partially supported by the U.S. National Science Foundation. This formal assessment will use the correlation of concept maps between the students and control systems experts.

### Acknowledgements

The author gratefully acknowledges the support for this work provided by the NSF under CAREER award CMS-0546430.

## References

- [1] Imre Benyo, Gyorgy Lipovszki, and Jenő Kovacs. Advanced control simulation tools in LabVIEW environment. In *Proc. 6th IFAC Symposium on Advances in Control Education*, 2003.
- [2] Robert Bishop. *Learning with LabVIEW 7 Express*. Pearson Prentice Hall, 2003.
- [3] Robert Bishop and Nick Lin. Control lab setup: Using LabVIEW to control ECP plants. Technical report, Connexions, 2005.

- [4] Jeanne Sullivan Falcon. Graphical programming for field programmable gate arrays. In *Proc. American Control Conference*, 2006. Tutorial Session.
- [5] Finn Haugen. *Introduction to LabVIEW Control Design Toolkit*. TechTeach, [www.techteach.no/publications/labview\\_control\\_design/](http://www.techteach.no/publications/labview_control_design/), 2005.
- [6] Finn Haugen. *Introduction to LabVIEW Simulation Module*. TechTeach, [www.techteach.no/publications/labview\\_simulation/](http://www.techteach.no/publications/labview_simulation/), 2005.
- [7] John Limroth, Jeanne Sullivan Falcon, Dafna Leonard, and Jenifer Loy. Labview real-time for networked/embedded control. In *Handbook of Networked and Embedded Control Systems*, pages 447–470. Birkhauser, 2005.
- [8] T. D. Murphey. Teaching rigid body mechanics using student-created virtual environments. *IEEE Transactions on Education*, 2007. Accepted for Publication.
- [9] R. J. Shavelson, M. A. Ruiz-Primo, and E. W. Wiley. Windows into the mind. *Higher Education*, 49:413–430, 2005.