

- ⁶J. Bechhoefer, "Feedback for physicists, a tutorial essay on control," *Rev. Mod. Phys.* **77**, 783–836 (2005).
- ⁷K. Ogata, *Modern Control Engineering*, 4th ed. (Prentice Hall, New York, 2002), pp. 682–699.
- ⁸P. K. Dixon and Lei Wu, "Broadband digital lock-in amplifier techniques," *Rev. Sci. Instrum.* **60**(10), 3329–3336 (1989).
- ⁹P. K. Dixon, "Specific-heat spectroscopy and dielectric susceptibility measurements of salol at the glass transition," *Phys. Rev. B* **42**(13), 8179–8186 (1990).
- ¹⁰P. K. Dixon, "Third harmonic dielectric response of a water-in-oil emulsion," *Phys. Rev. B* **55**, 6285–6295 (1997).
- ¹¹P. Horowitz and W. Hill, *The Art of Electronics*, 2nd ed. (Cambridge U.P., Cambridge, 1989), pp. 297–303.
- ¹²D. Damjanovic, "Ferroelectric, dielectric and piezoelectric properties of ferroelectric thin films and ceramics," *Rep. Prog. Phys.* **61**, 1267–1324 (1998).
- ¹³K. Uchino, *Ferroelectric Devices* (Marcel Dekker, Oxford, 2000), p. 215.
- ¹⁴G. H. Haertling, "Ferroelectric ceramics: History and technology," *J. Am. Ceram. Soc.* **82**(4), 797–1615 (1999).
- ¹⁵D. Hennings, "Barium titanate based ceramic materials for dielectric use," *Int. J. High Technol. Ceram.* **3**(2), 91–111 (1987).
- ¹⁶See EPAPS Document No. E-AJPIAS-75-011706 for the supplementary text. This document can be reached through a direct link in the online article's HTML reference section or via the EPAPS homepage (<http://www.aip.org/pubservs/epaps.html>).

Temperature dependence of the capacitance of a ferroelectric material

John Bechhoefer^{a)} and Yi Deng

Department of Physics, Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada

Joel Zylberberg

Department of Physics and Department of Chemistry, Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada

Chao Lei and Zuo-Guang Ye

Department of Chemistry, Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada

(Received 12 September 2006; accepted 16 March 2007)

We present an alternate version of the undergraduate laboratory experiment developed by Dixon [*Am. J. Phys.* **75**, 1038–1046 (2007)] that is suitable for second-year students. We study the temperature variation of the capacitance of a ferroelectric ceramic derived from barium titanate, the $\text{Ba}(\text{Ti}_{0.9}\text{Sn}_{0.1})\text{O}_3$ solid solution. The ratio of tin to titanium is chosen to provide a convenient Curie temperature near 50°C . Using careful temperature control and real-time capacitance measurements, we track the time evolution of the capacitance in response to temperature changes at 5 Hz for runs that last up to a day. At temperatures well above the Curie temperature, T_C , the capacitance relaxation is well-described by a single exponential decay. Near T_C , the relaxation is linear in the logarithm of time over more than three decades. For $T > T_C$, the permittivity deviates from the Curie–Weiss law and follows another phenomenological form commonly used to describe relaxor perovskite-ceramic capacitors. © 2007 American Association of Physics Teachers.
[DOI: 10.1119/1.2723801]

I. INTRODUCTION

Over the last few decades, an extraordinary rate of technological growth has led to significant advances in computer hardware and software, as well as associated hardware such as data acquisition devices (DAQs). While such advances have transformed the technological world we live in, they have had rather less impact on physics instrumentation and physics experiments, particularly on those used for undergraduate instruction. While books and articles describing some of these advances in the context of undergraduate physics labs have been written,^{2,3} they only touch the surface of possibilities that have been created. In addition, the very rapid advance in technology has meant that even relatively recent books, such as that of Essick,² do not reflect significant possibilities now available for creating experiments that once would have been possible only at a more advanced level.

In this note, we describe an experiment that measures the real part of the dielectric permittivity of a ferroelectric material as a function of both time and temperature. It differs from traditional versions of similar experiments^{4,5} in forgo-

ing specialized instrumentation—temperature controllers, function generators, impedance analyzers, and the like—in favor of general-purpose DAQs and signal conditioning electronics with a computer implementing the instrumentation in software. Our experiment was inspired by the work presented by Dixon in the related note in this issue.¹ In our version, we substitute a custom-synthesized ferroelectric ceramic for the commercial capacitor used in Ref. 1. The custom material has a more convenient Curie temperature (transition between paraelectric and ferroelectric phases).⁶ Also, we have developed the lab for a second-year undergraduate course, with no electronics prerequisite, while the experiment in Ref. 1 was designed for third-year students who have taken a first electronics course. A more detailed discussion of the differences between the two approaches is provided in Sec. VI.

While the ferroelectric-capacitor experiment is interesting in its own right—opening windows to the worlds of phase transitions, critical phenomena, complex relaxational dynamics, and more—our overall goal is to give an exemplar of the type of experiment that can be created using modern design principles and present-day technology. We use a stripped-

down apparatus, consisting simply of a data acquisition device (DAQ), a preamp that was actually a part of the DAQ, and a power amplifier. Compared to previously published designs for the undergraduate lab, e.g., Refs. 4 and 5, our design is cheap and simple to make. Making many copies of the set-up to accommodate larger classes where all students do the same experiment in parallel then becomes feasible. By leveraging the power of present-day computers, our design outperforms in some ways commercial research equipment that costs a hundred times more. It is particularly satisfying that, although the physics touched on by this experiment is complex enough to remain topics of active research, we have implemented this lab successfully at the second-year undergraduate level. While many issues must be approached in a qualitative, intuitive way, we can nonetheless show beginning students some of the power and, we hope, the excitement of building and carrying out a sophisticated experiment with modest technical means.

II. EXPERIMENTAL SET-UP

In this section, we briefly describe the experimental set-up, calling attention to some of its most important and original features. Those readers interested in actually building and implementing the experiments described here should consult the supplementary material.²³ Our description of the temperature regulation is more detailed, as control theory is an important subject that usually gets short shrift in physics courses. Also, our implementation of temperature control deals with a number of subtle points that are not usually discussed.

Sample preparation. One of the major differences in our version of the experiment relative to that of Ref. 1 is the use of a custom-made ferroelectric material. Being able to prepare or synthesize the material used has several advantages: The dielectric material of commercial ceramic capacitors shows a very broad peak near room temperature. In our material, the peak is tuned, by substitution of Sn^{4+} for Ti^{4+} , to occur near 50°C , allowing us to explore both sides of the transition. The transition is also noticeably sharper than that of a commercial capacitor. The disadvantage of using a custom material is that it takes time and a modest amount of money to make special materials, but this burden is not placed on the students.

Sample assembly. The sample-assembly protocol was developed and the samples assembled by Jeff Rudd. The main challenges were to provide mechanical robustness so that the sample could survive sometimes-rough handling by students, give good thermal contact between the heater and both the thermistor and capacitor, and be able to withstand the repeated heat shocks created during heating and cooling cycles. Details are given in the supplementary material.²³

Power amplifier. The output from the DAQ's A/D converter is not nearly enough to drive the heater. While Ref. 1 makes the construction of a power op-amp a part of the students' task, our students have not yet taken an electronics course and we thus needed to provide an amplifier that could handle the heater's load. Because we wanted 25 set-ups, it was far cheaper to design and build our own circuit than to buy a commercial product. The power amplifiers were built in our electronic shop following a design based on a single, high-power operational amplifier.

III. TEMPERATURE CONTROL

Temperature control is an important part of the lab. At the beginning, we ask students to measure the capacitance of the sample using the capacitance meter of an ordinary digital multimeter (without shielding the sample from air currents). The sensitivity to temperature variations reaches $(1/C) \times (dC/dT) \approx 0.1/^\circ\text{C}$ at the steepest parts of the $C(T)$ curve, which is large enough that readings drift continuously on a three-digit digital multimeter. This point may be further emphasized by gently touching the capacitor or by breathing on it.

Having motivated the need to control temperature in order to get a reasonable measurement of the capacitance, we have the students set up a temperature-control loop. This actually occurs over two laboratory sessions of three hours each. We believe that the importance of control concepts, coupled with the pedagogical effectiveness of the setting here, justifies the time spent. Because the sample has a low mass (≈ 5 gm) and because the heater is powerful (3.7 W), the temperature changes quickly when the heater power is changed, up to 3°C/s . The rapid response makes it easy to observe the effects of changes in the control loop on the performance of the regulation.

A. Characterization of the thermal system

In the first session, the students characterize the thermal part of the sample system. They first calibrate the thermistor signal, establishing a relationship between its temperature and the voltage across it, when the thermistor is placed in series with another resistor. Having calibrated the input, they do the same with the output, establishing a relationship between the DAC voltage and the power across the heater resistor. It is important to work with the output power rather than the DAC voltage, as is usually done. This makes the relationship between the input and output variables of temperature and power a linear one. As we discuss below, it is much easier—and the results are much better—when one controls a linear system rather than a nonlinear one.

After linearizing the thermal system, students explore the effects of averaging readings, measuring time series of temperature readings, with each reading the average of N readings taken at 48 kHz, the maximum allowed by the DAQ. For small N , the standard deviation of successive readings decreases as \sqrt{N} : more averaging decreases the measurement noise. Eventually, for large enough N , the deviation begins to increase as one begins to average over real variations in temperature. At $N=2000$, which corresponds to an averaging time of just under 40 ms, the read noise is about 0.5 mK. The choice of N is a compromise between the benefits of signal averaging and the need to respond to temperature variations due both to the environment and the heater.

B. Implementation of the temperature control loop

In the second laboratory session, the students build and implement a temperature-control loop in software. Again, in the last few years, there has been a nearly complete adoption of digital control loops which replace the dedicated electrical circuits of analog design.⁸ Because students do not have access to LabVIEW or to the experimental apparatus outside of the lab session, we supply a PID module written in LabVIEW. The module, however, is easy to explain and we go

over it in class. (The main complication is the use of a shift register.) The students, however, program the connections to the sub-VI, as well as the temperature input, signal averaging, and heater output. The sub-VI implements standard discrete-time PID control

$$P_n = - \left[K_p e_n + K_i (Int_n) + \frac{K_d}{\Delta t} (e_n - e_{n-1}) \right], \quad (1)$$

where P_n is the power output at time $n\Delta t$; K_p , K_i , and K_d are the proportional, integral, and derivative gains, respectively; and e_n and e_{n-1} are the error signals at times $n\Delta t$ and $(n-1)\Delta t$. Here, Δt is the sampling time, 200 ms. Prior to evaluating Eq. (1), we calculate the integral term Int_n using the trapezoidal rule

$$Int_n = Int_{n-1} + \frac{\Delta t}{2} (e_n + e_{n-1}). \quad (2)$$

We also protect against integral windup⁹ by not updating the integral term when the calculated output power would exceed the maximum (3.7 W) or be less than the minimum (0 W). Finally, the integral term is set to zero if $K_i=0$, to prevent saturation when the integral gain is first turned on.

We note that this temperature-control algorithm is already more sophisticated than many commercial controllers. In particular, while some controllers offer the possibility of linearizing a thermistor signal, we do not know of any stand-alone controllers that linearize the power output. If the voltage rather than the power were output, the optimal PID coefficients would be different at different temperature setpoints. Since we want to sweep through a wide range of temperatures without having to readjust the PID coefficients, this would be an important limitation. Perhaps more important, most stand-alone controllers do not provide for the kind of low-pass filtering that we in effect do when averaging 2000 readings. The amount of noise limits the strength of the derivative term that can be applied, while the strength of the derivative term is crucial for improved performance. In the end, the difference in performance is significant. While most commercial controllers regulate to 0.1 °C, the controller here performs to ≈ 0.001 °C, as shown below. The increased temperature stability is not needed for measuring the basic $C(T)$ curve but is useful for exploring the slow internal dynamics of the ferroelectric transition, as discussed in Sec. V A. The performance of the temperature controller here is comparable to that in Ref. 1; however, the controller of Dixon uses a dedicated 6-1/2 digit DMM, while ours uses only the standard analog input of our DAQ.

The bulk of the second session is then devoted to an exploration of basic concepts in control theory. As one of the authors (JB) has argued elsewhere,⁹ control theory is an important subject that historically has not received its due in physics courses. We believe that introducing the concepts of a PID controller at the second-year level is valuable. Because we do not assume that the students know Fourier analysis, we keep the concepts informal and intuitive. But these are already valuable. Specifically, we start with pure proportional control and verify that, for low gains, the system comes to an equilibrium at a temperature that is $\Delta T = a/(b + K_p)$ below the set point, where a and b are empirically determined. As the gain increases, there is an instability at a critical gain K_p^* . The students track the performance of their loop at different times by looking at the absolute range of

temperature fluctuations over a brief period, e.g., 30 s. Measuring the performance of the control loop against some such standard is the essential idea behind optimal control—here introduced in a simple and natural setting.

Having measured the critical gain K_p^* , the students then reduce K_p and introduce derivative gain, K_d . The derivative term allows one to increase the proportional gain with a corresponding improvement of the control. Its value is limited by the noise of the system, which is reduced by making each “reading” the average over many temperature readings. Then one introduces the integral term, K_i , which removes the “proportional droop” (Ref. 9), allowing the system to settle to the desired setpoint. Finally, one “touches up” the values of all three parameters to optimize the control.

Figure 1 shows measurements of the thermistor temperature, after the PID parameters have been fine-tuned in accordance with the above discussion. Figure 1(a) shows a time-series of over 400 000 points collected at 5 Hz over an entire day, with each point the average of 2000 measurements at 48 kHz. The series shows no diurnal temperature variation and has a maximum excursion of 6 mK. Typical fluctuations are much less, and the rms deviation is 1.16 mK. At lower temperatures, the inability to cool the sample actively introduces a nonlinearity in the feedback response that degrades performance. At 30 °C, the typical performance is worse by a factor of 2. This still easily meets the needs of the experiment.

As good as the temperature control is, it could still be better. Figure 1(d) shows the autocorrelation function of the temperature deviations and points the way to a more sophisticated analysis. The autocorrelation shows the response of the system to perturbations and can also be explored by manually perturbing the temperature, e.g., by momentarily disconnecting the heater from the amplifier. One can clearly see the influence of the system’s dynamics. In Fig. 1(d), the response is slightly more underdamped than it should be. If overshoot is allowed, the fastest response is one that overshoots once and then relaxes monotonically. Here, there is a small second overshoot. In the autocorrelation function, sensor noise from the input shows up as noise on the DAQ output that is input to the system. Such notions are not pursued in our second-year version of the lab but could be in a more advanced version. By correlating the temperature output with the power input, one can measure the frequency response of the dynamical transfer function and then design a matching control loop in frequency space.⁹

A second limitation of the temperature-control system is that the thermal stability of the thermistor is not the same as the thermal stability of the capacitor one is interested in. Because the thermistor and capacitor are located symmetrically about the heater and are in good contact with it, first-order temperature variations can be eliminated by choosing equal masses for the two components. With good insulation from the environment, second-order variations can be small. In our implementation, however, the masses are not perfectly matched, and we use only a small piece of transparent tubing to shield the system from air currents. While using transparent insulation makes for good pedagogy—students can see the components as they work on them—it leads to larger thermal gradients. From the intermediate-scale variations of capacitance, as seen in Fig. 2(a) below, we estimate that the capacitor has temperature variations of 30 mK over the time scale of tens of minutes, which we attribute to variations in the thermal load caused by air currents and the overall tem-

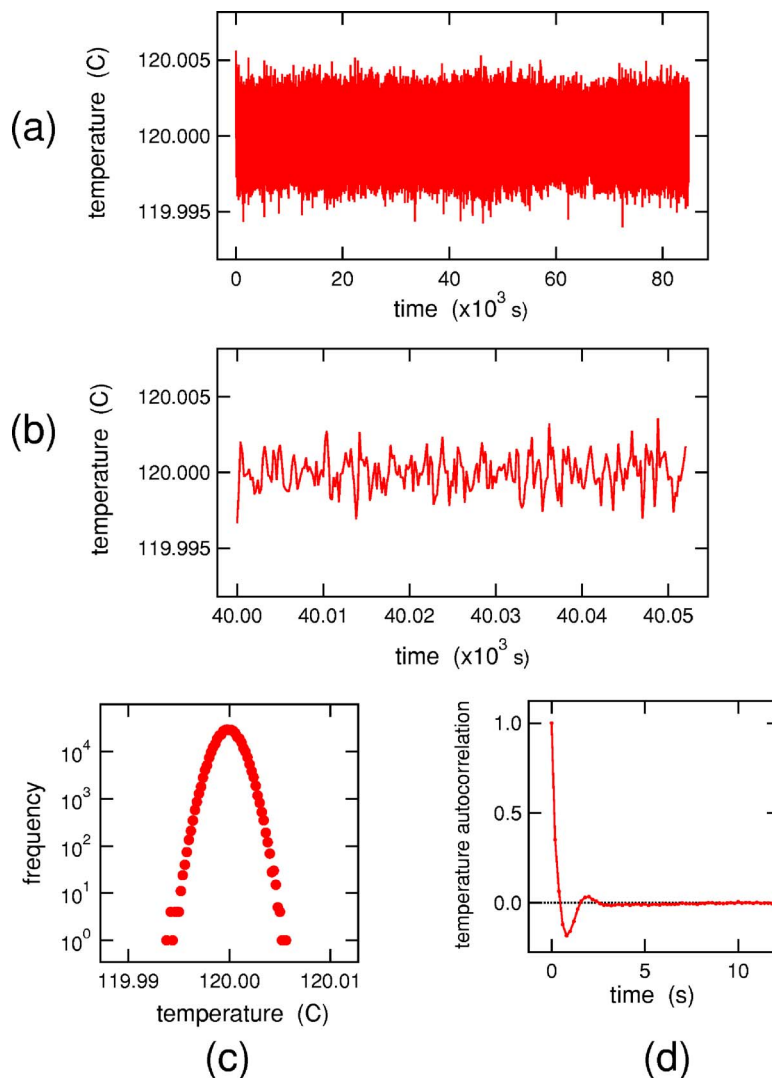


Fig. 1. Performance of the temperature control system. (a) Raw time series: one day's worth of data, acquired at 5 Hz. (b) Detail of (a). (c) Histogram of data in (a) showing Gaussian distribution of fluctuations. (d) Autocorrelation function of temperature deviations, showing rapid stabilization when perturbed.

perature variations of the laboratory room. For the experiments reported here, such variations are much smaller than the 100 °C temperature range that is probed in large steps of 1 °C. See the supplementary material for more discussion.²³

IV. DATA ACQUISITION AND ANALYSIS

One important feature of this experiment is the structure of the data acquisition and analysis programs. We separate the two functions by using one program, written in LabVIEW,¹⁰ for data acquisition and another, written in Igor Pro,¹¹ for analysis.

The data acquisition is written using a simple state-machine architecture, a “real-time, interrupt-driven” style of programming that most students have not encountered before. The essence of such a program is that everything is synchronized to a master clock, here running at 5 Hz. Each time the program cycles through the main loop, it takes a number of actions based on the state of the program, including those influenced by external data. Details are given in the supplementary material.²³

V. RESULTS

While the first two laboratory periods are devoted to characterizing the thermal properties of the sample and to learning how to control temperature, the last two are devoted to the capacitance measurements. In the third session, the students put together the circuitry for measuring capacitance at a fixed temperature. Time permitting, they explore the slow relaxation effects present in such measurements. In the fourth session, the students write a program to sweep the temperature, measuring capacitance, in order to construct full $C(T)$ curves.

A. Capacitance at constant temperature

In the third lab period, students learn to measure the capacitance at a fixed temperature. We begin with an example of a typical capacitance discharge curve. Figure 3 shows data from a RC decay, along with a fit to a single exponential. As the residuals and the semi-logarithmic plot in the inset show, the fit is nearly, but not quite, well-described by a single time scale. In this experiment, we fit all the data by a single ex-

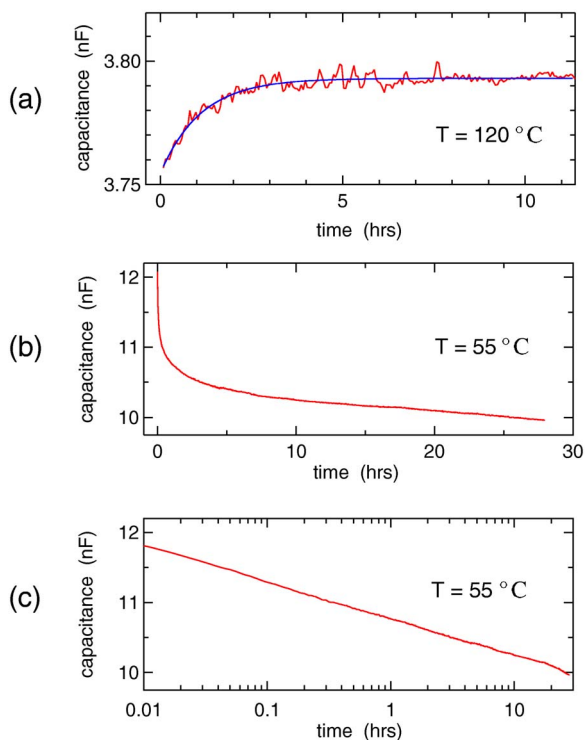


Fig. 2. Relaxation of capacitance. (a) Data at 120 °C. (b) Data at 55 °C. (c) The decay is approximately linear in the logarithm of time over more than three decades.

ponential, neglecting the small nonexponential deviations explored in Ref. 1. In Fig. 3, the value of the decay constant inferred from the fit to an exponential over the whole range of 0 to 4 ms in the primary figure differs by less than 1% from that inferred from the fit shown in the inset, which covers only the linear portion of that plot.

The ability to record long stretches of data allows one to explore the dynamics of how the capacitance relaxes to its equilibrium value. Naively, one expects that the value of the capacitance measured via RC decays, such as the one shown in Fig. 3, would follow any temperature changes. Following

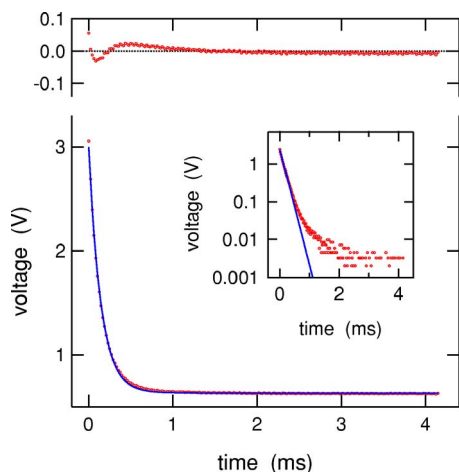


Fig. 3. Discharge of capacitor at 120 °C. The primary figure shows the data points, along with a single exponential fitted to the data. The fit residuals are shown at the top of the graph. The inset shows a semi-logarithmic plot of the decay, with the long-time offset subtracted.

a step in the temperature set-point, the thermistor temperature stabilizes to 5 mK within 30 s. The striking observation, though, is that the measured time constants of the RC decay change over much longer times, implying that processes intrinsic to the capacitor itself have long time scales. Figure 2 shows time-series of measured capacitance values taken at two temperatures, 120 °C and 55 °C, near the paraelectric-ferroelectric transition temperature. Both time series begin one minute after a jump in the temperature set-point. At the higher temperature, the relaxation is a small effect, about 1%, and is well-fit by a single exponential with a time constant of 1.1 h. Near the transition temperature, the relaxation is much stronger and slower. In Fig. 2(b), the capacitance relaxes nearly 20% over the course of a day. Figure 2(c) shows that the relaxation is approximately logarithmic in time, or linear in $\log t$, with an “aging rate” of 4.3%/decade. A stretched exponential form also fits well. After a day, the relaxation shows little sign of saturation. Such large aging effects in ferroelectric materials such as the barium-titanate-tin-based solid solution explored here were discovered in the late 1940s and extensively explored in the 1950s.¹² While there is some agreement that the long time scales reflect a wide distribution of relaxation times associated, in particular, with the ferroelastic effects produced by the cubic-tetragonal structural transition associated with the paraelectric-ferroelectric transition, theoretical explanations and models continue to appear to this day.¹³ The logarithmic decrease has been followed for *years*.¹⁴ This raises an interesting point—nice to discuss with students—about what “measuring the capacitance” means when the value you obtain will vary with the time you take for the measurement.

As an amusing aside, the variations in capacitance with time are significant enough that manufacturers must take them into account. The technical literature of one company notes that manufacturers will reheat their capacitors above the Curie temperature, thus resetting their aging properties, just before shipping. The reset value is designed to be larger than the target value so that, after ten days, the capacitance will have drifted down to its specified value.¹⁵ Of course, the value will continue to drift down, and after a long period of time may be significantly lower—users beware!

Logarithmic relaxation—“aging”—is found in many other complex systems. Examples include the density of a pile of sand subject to periodic tapping,¹⁶ the Hall magnetization of a high-temperature superconductor,¹⁷ and the volume of glasses under pressure.¹⁸ Explaining the generality of logarithmic decay in glassy systems remains a theoretical challenge.¹⁹ Because of student time constraints, we have not asked them to explore relaxation and aging in any detail. By recording a few minutes of data, they see that the capacitance dynamics are much slower than those of the temperature or power. It would be better to do this at 120 °C, as shown here, but the glue used in the current version decomposes and becomes brittle at those temperatures. Exploring the very long times, i.e., hours to days, associated with aging is awkward since our students do not have access to the experiment outside the one afternoon a week scheduled for the course. If students have small amounts of time spread out over several days, they could begin to explore some of the issues connected with slow relaxations.

B. Temperature sweeps

In the fourth and final lab period, students complete the LabVIEW program to sweep the temperature and then ex-

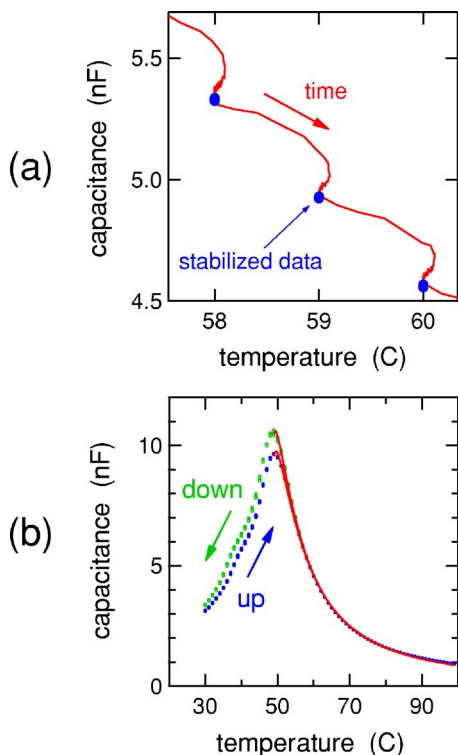


Fig. 4. Typical data from temperature sweeps. (a) Solid line: logged data, including transients; markers show stabilized, “filtered” data used for analysis. (b) Capacitance vs temperature, measured at $f=5$ Hz and processed as illustrated in (a). Solid lines in (b) are fits to Eq. (3).

plore the capacitance measured while increasing and decreasing the temperature. Since the slow relaxation effects described above always lead to a decrease in capacitance, the two curves will not bracket the equilibrium measurement. Nevertheless, they illustrate how the measured capacitance depends on the experimental protocol.

Figure 4 gives typical results from an upward sweep and a subsequent downward temperature sweep in the range 30–100 °C. Because the LabVIEW program is set up as a data logging program, all of the data, including transients, are recorded to disk. The solid curve in Fig. 4(a) shows a portion of the raw data for the increasing temperature sweep. The loops are caused by the overshoot of the temperature control when the set point is increased. The isolated markers are “filtered” data: these are points that are selected offline after the data collection is complete, in accordance with various criteria. In Fig. 4(a), the markers are averages of 100 points taken over 20 s, selected only if the temperature over the 100-point window has a standard deviation less than 0.01 °C and the standard deviation of the capacitance over the same window is less than 0.03 nF. These conditions are implemented in a simple “data filter” routine written in the Igor programming language. Our protocol is to log the raw data at 5 Hz, i.e., once each cycle of the main loop, and then to process it later by averaging and selecting the stable points. Its great advantage is that one can explore different filtering criteria without having to retake data. Alternatively, one would have to decide beforehand how long to wait before taking a data point. Here, one can revisit any criterion, to decide afterwards the best filtering algorithm. The only disadvantage is that more data must be stored, but this is not a problem on current computers. In the end, our protocol is

just an updated version of the traditional method of recording data on a strip chart and analyzing it afterwards. But here, one can reach quantitative conclusions more easily and can optimize the data filtering without redoing the experiment.

Figure 4(b) shows the main result: capacitance vs temperature, for increasing and decreasing temperatures. The raw data are acquired at $f=5$ Hz and filtered, as described above. In the data shown, the program waits for 100 s at each point, then increments the temperature by 1 °C. In practice, students usually set the wait time at 30 s, in order to keep the runs short (~ 35 min). Nonequilibrium effects are, as a result, slightly larger than those shown in Fig. 4(b). The raw data are then filtered, as described above. One immediately sees the effects of slow relaxations. Interestingly, the largest differences between the heating and cooling curves are observed in the vicinity of the capacitance peak. Since the peak is associated with a weak first-order phase transition,²⁰ one expects particularly slow dynamics near the transition temperature. Local parabolic fits to the peaks gives a transition temperature of 49.0 ± 0.2 °C, with the temperature difference between the peaks of the two curves coming from the thermal hysteresis associated with this first-order transition. Here the heterogeneous microdomains, each with their own transition temperature, cause the hysteresis to appear as anomalously slow relaxation effects. The important point is that neither curve is an equilibrium measurement.

The data in Fig. 4(b) suggest a divergence due to a second-order phase transition. Because the grains and domains are of finite size and because the curves are not equilibrium measurements, a simple theory such as the Curie–Weiss description of a paraelectric-ferroelectric transition does not accurately describe the data from “relaxor” materials such as that used in this experiment. One can, however, use another phenomenological law for $T > T_C$:²¹

$$C(T) = \frac{C_{max}}{1 + K(T - T_C)^\gamma}, \quad (3)$$

where T_C is the transition temperature, C_{max} is the capacitance at the transition, K is an amplitude, and $\gamma \approx 1-2$. As one can see, this phenomenological form fits our data well. In Fig. 4(b), both the up and down data give $\gamma = 1.40 \pm 0.01$. For temperatures greater than 200 °C, well beyond the maximum working temperature of this experiment, the capacitance *does* seem to follow the Curie–Weiss law.²²

Finally, note the hint of a “shoulder” in the $C(T)$ data of Fig. 4(b) below T_C , at $T \approx 43$ °C. This nearly hidden peak is associated with another structural transition in the barium titanate solid solution, from a tetragonal to an orthorhombic lattice.^{7,21} The sequence of phases is inherited from that of pure BaTiO₃. This transition is seen more clearly in the variation of the dielectric loss tangent as a function of temperature, as shown in the supplementary materials.²³

VI. DISCUSSION

While we have gone into a careful and fairly extensive discussion of the physics of our ferroelectric material, we emphasize that, in our version of this experiment for second-year physics majors, we do not ask the students to consider these questions in anywhere near this detail. Rather, the focus is on building the experiment, understanding its various elements, and obtaining basic results—in particular, the $C(T)$ curve of Fig. 4(b). We dwell on other issues such as relax-

ation times for two reasons: first, even if the students are not asked to go into great depth, it is useful for the instructors and TAs to understand the issues. If nothing else, it will help to answer the inevitable questions and to “debug” problematic results. Second, we have tried to give points of departure for students who have more time and would like to go into greater depth. For example, one could study as a function of temperature the relaxation of either voltage, as in Fig. 3, or capacitance, as in Fig. 2. One could also look at the interplay between the two: is there a link between the nonexponential features seen in each decay? Looking at the form of each kind of relaxation as a function of temperature would be a way to approach this question.

It is instructive to dwell on some of the differences between the experiment presented here and the version of Ref. 1, as one sees how differing goals and constraints can lead to differing outcomes. Both experiments explore the temperature dependence of the capacitance of a ferroelectric material in an undergraduate setting. The important differences are as follows:

- Custom capacitor material: This gives a more convenient T_C , allowing one to explore both the paraelectric and ferroelectric sides of the transition, without needing to cool the sample.
- Simplified temperature control: even without using a multimeter, we can explore delicate material relaxation effects requiring high temperature stability to observe.
- Data-logging structure of the code at 5 Hz: This allows one to explore relaxation effects—for example, to distinguish between the “trivial” thermal relaxation of the system and the additional relaxation processes due to the internal dynamics of the ferroelectric material. The structure is also useful pedagogically, in that it makes explicit the data-filtering step, which is always present, if only in the protocol for taking data.

On the other hand, in aiming the lab at a lower level of student, we give up some things, too. In particular, in Ref. 1, students focus more on the data acquisition programming and thus learn the programming techniques more deeply. In our case, we have chosen to focus on teaching the structure of an experiment. This is our students’ first introduction to an experiment that has the major elements of a “real” experiment: regulation of certain variables and controlled sweep of one particular variable, coupled with the measurement of the system’s response.

For this reason, our students use high-level Express VIs in LabVIEW wherever possible. We also supply sub-VIs that implement the PID and the curve-fitting to an exponential of the RC decay data. In Ref. 1, students learn lower-level LabVIEW routines and program all parts of the experiment. Some of that programming is done at home, as all students are expected to have their own version of LabVIEW. In our case, we do not require students to purchase LabVIEW. Still, at the end of the four weeks, the students will have built up—and understood—reasonably elaborate LabVIEW and Igor programs that are much more complex than what they could have achieved easily on their own.

VII. CONCLUSIONS

We have described an experiment on the measurement of capacitance as a function of temperature in a ferroelectric

material undergoing a transition from paraelectric to ferroelectric phases. By building the experiment from standard lab equipment with only an inexpensive amplifier as a custom part, we were able to afford to build many stations (12 each at two different campuses), rather than a single experiment through which students must rotate.

Although the experiments were inexpensive, they produced high-quality results, with temperature control at the mdeg level and capacitance measured to the 1 pF level. Moreover, the tight integration of all elements in software allowed data to be taken at 5 Hz for indefinite periods of time—days, if desired. This creates possibilities for exploring slow relaxation effects, along with the more easily accessible peak in the capacitance at the phase transition temperature.

The experiment is aimed at second-year physics majors taking a course on data acquisition and analysis. While one goal is to explore the physics of the system, in the context of our course, an even more important goal is to teach the structure of a typical modern physics experiment at an early stage in a student’s education. We want to emphasize, in particular, that there is no “best” way of structuring such an experiment. Each experiment in a lab course should fit in with the rest of the course, and each course should fit in with the other courses in a student’s program. Our version of this experiment was thus optimized against our particular set of constraints. In this article, including the supplementary material,²³ we hope to have provided enough detail that others can adapt this experiment to their own situation.

ACKNOWLEDGMENTS

This work was supported by the teaching budget of Simon Fraser University and by NSERC (Canada). We thank Paul Dixon for informing us of his work and for allowing us to present our work alongside his. We thank Jeff Rudd for his tireless efforts in developing a protocol for assembling the samples. We thank Pawel Kowalski for designing, building, and testing the power amplifiers.

^{a)}Author to whom correspondence should be addressed. Electronic mail: johnb@sfu.ca

¹P. K. Dixon, “The bug: a temperature-controlled experiment on a proto-board,” *Am. J. Phys.* **75**, 1038–1046 (2007), preceding paper.

²J. Essick, *Advanced LabView Labs* (Prentice Hall, Upper Saddle River, NJ, 1999). While a good book, parts of this work are already somewhat dated. As a specific example, the introduction in LabVIEW 7.0 (in 2003) of Express VIs such as the “DAQ Assistant” significantly reduce the programming effort needed to interface data-acquisition devices. Without such a simplification, we could not have implemented such a complex experiment at the second-year level, as the effort involved in teaching programming would have required too much time from an already crowded course and would have skewed the balance we were seeking between techniques and physics.

³P. J. Moriarty, B. L. Gallagher, C. J. Mellor, and R. R. Baines, “Graphical computing in the undergraduate laboratory: Teaching and interfacing with LabVIEW,” *Am. J. Phys.* **71**, 1062–1074 (2003).

⁴M. Trainer, “Ferroelectrics and the Curie Weiss law,” *Eur. J. Phys.* **21**, 459–464 (2000).

⁵M. Trainer, “Ferroelectricity: Measurement of the dielectric susceptibility of strontium titanate at low temperatures,” *Am. J. Phys.* **69**, 966–969 (2001).

⁶M. P. Marder, *Condensed Matter Physics* (Wiley-Interscience, New York, 2000), pp. 609–611.

⁷N. Yasuda, H. Ohwa, and S. Asano, “Dielectric properties and phase transitions of Ba(Ti_{1-x}Sn_x)O₃ solid solution,” *Jpn. J. Appl. Phys., Part 1* **35**, 5099–5103 (1996).

- ⁸D. A. van Baak, "Temperature servomechanisms using thermoelectric modules," *Am. J. Phys.* **60**, 803–815 (1992).
- ⁹J. Bechhoefer, "Feedback for physicists: A tutorial essay on control," *Rev. Mod. Phys.* **77**, 783–836 (2005).
- ¹⁰National Instruments, LabVIEW, v. 8.0. www.ni.com
- ¹¹WaveMetrics, Igor Pro, v. 5.04b. www.wavemetrics.com
- ¹²W. A. Schultze and K. Ogino, "Review of literature on aging of dielectrics," *Ferroelectrics* **87**, 361–377 (1988).
- ¹³D. C. Lupascu, Y. A. Genenko, and N. Balke, "Aging in ferroelectrics," *J. Am. Ceram. Soc.* **89**, 224–229 (2006).
- ¹⁴K. W. Plessner, "Aging of the dielectric properties of barium-titanate ceramics," *Proc. Phys. Soc. London, Sect. B* **69**, 1261–1268 (1956).
- ¹⁵R. Fiore, "Dielectric aging phenomena," excerpt from the *Circuit Designer's Notebook*, American Technical Ceramics Corp., www.atceramics.com (see Technical Notes).
- ¹⁶J. B. Knight, C. G. Fandrich, C. N. Lau, H. M. Jaeger, S. R. Nagel, "Density relaxation in a vibrated granular material," *Phys. Rev. E* **51**, 3957–3963 (1995).
- ¹⁷L. Fruchter *et al.*, "Low temperature magnetic relaxation in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$: Evidence for quantum tunneling of vortices," *Phys. Rev. B* **43**, 8709–8712 (1991).
- ¹⁸O. B. Tsiok, V. V. Brazhkin, A. G. Lyapun, and L. G. Khvostantsev, "Logarithmic kinetics of the amorphous-amorphous transformations in SiO_2 and GeO_2 glasses under high pressure," *Phys. Rev. Lett.* **80**, 999–1002 (1998).
- ¹⁹W. Götze and M. Sperl, "Logarithmic relaxation in glass-forming systems," *Phys. Rev. E* **66**, 011405 (2002).
- ²⁰M. E. Lines and A. M. Glass, *Principles and Applications of Ferroelectrics and Related Materials* (Clarendon Press, Oxford, 1977), pp. 71–80.
- ²¹L. Mitoseriu, V. Tura, M. Curteanu, D. Popovici, and A. P. Anghel, "Aging effects in pure and doped barium titanate ceramics," *IEEE Trans. Dielectr. Electr. Insul.* **8**, 500–506 (2001).
- ²²Chao Lei, unpublished data.
- ²³See EPAPS Document No. E-AJPIAS-75-012706 for a more detailed description of the sample preparation and characterization and for more details about the experimental apparatus. This document can be reached through a direct link in the online article's HTML reference section or via the EPAPS homepage (<http://www.aip.org/pubserve/epaps.html>).

Young's double-slit interference experiment with electrons

S. Frabboni^{a)}

*Department of Physics, University of Modena and Reggio Emilia and CNR-INFM-S3,
Via G. Campi 213/a, 41100 Modena, Italy*

G. C. Gazzadi

CNR-INFM-S3, Via G. Campi 213/a, 41100 Modena, Italy

G. Pozzi

Department of Physics, University of Bologna, v. B. Pichat 6/2, 40127 Bologna, Italy

(Received 30 November 2006; accepted 14 June 2007)

In this short Note we report a method for producing samples containing two nano-sized slits suitable for demonstrating to undergraduate and graduate students the double-slit electron interference experiment in a conventional transmission electron microscope. © 2007 American Association of Physics Teachers.

[DOI: 10.1119/1.2757621]

The double-slit electron interference experiment, whose pedagogical value in illuminating the wave-particle duality in modern physics cannot be underestimated,¹ was for a long time considered purely of the gedanken variety.

Its first realization is due to Jönsson,² who, with ingenuity, was able to produce slits in the micrometer range and to observe them using a dedicated electron diffraction apparatus that consisted of cylindrical and rotationally symmetric electrostatic lenses that were used to have the slits illuminated coherently by the electron beam and to suitably magnify the diffraction pattern.

An alternative didactic version of the two-beam interference experiment uses the electron optical analog of the Fresnel biprism to produce two virtual coherent sources, instead of the real ones represented by the slits. Moreover, the interference phenomena are observed in the Fresnel (i.e., near-field) instead of the Fraunhofer (i.e., far-field) region.^{3–5}

Today, advances in technology make it possible to perform Young's experiment using commercial instrumentation: transmission electron microscope (TEM) can perform the role of the diffraction apparatus and a focused ion beam (FIB) machining device allows easy fabrication of the slits.⁶

In this work we report some results obtained recently with the aim of demonstrating to our students this fundamental experiment, using the additional flexibility of modern electron microscopes for which the photographic plate is replaced by a charge-coupled-device (CCD) camera, so that the images can be displayed clearly and directly in real time on a monitor screen or even on a lecture-hall screen.

The slits were fabricated by FIB milling on a commercial silicon-nitride membrane window commonly used for TEM sample preparation.⁷ The sample consists of a 3-mm-diameter, 200- μm -thick silicon frame, with a 100 \times 100 μm square window at the center, covered with a 500-nm-thick silicon-nitride membrane. The membrane thickness was chosen to minimize electron transmission from regions other than the opened slits.

FIB milling was performed with a dual-beam apparatus (FEI Strata DB235M), which combines a 30-keV Ga^+ FIB with a thermal field emission scanning electron microscope (SEM), having resolutions of 6 and 2 nm, respectively. The system allows nanoscale machining by ion milling and a simultaneous control of the work in progress by high-resolution SEM imaging.