

Writing Guide for Stanford Physics 107

Thaddeus D. Ladd and Dennis Ugolini

January 11, 2006

Abstract

Intended for students of Stanford's Physics 107 course, this manual introduces the issues of style and structure expected of a technical journal article such as *Physical Review Letters*. We emphasize clear exposition of a central thesis with logically flowing supporting material in a structure catering to the reader's expectations. These concepts are revealed through examples of both student and professional writing.

1 Introduction

Technical writing is an art, and not an easy one. Poor examples of it are easily found, and not just among student writing. Published literature contains an unfortunately large amount of unclear, muddled prose.

However, good writing is a critical part of good physics. Published papers are the working currency for a modern practicing physicist. Journal articles are often the only lasting result of research labor. They are an important product on which physicists are judged for funding grants and job candidacy. Most importantly, journal articles are the most broad-reaching and lasting means by which scientists reveal their ideas and discoveries to the community. In physics and in most of academia, the quality of our research is inevitably colored by the quality of our writing.

So what makes technical writing good, and what makes it bad? This is clearly a subjective question, but we may first approach it from extremes.

Consider first the structure of a paper. On one extreme, rambling from one topic to the next without apparent logical structure, such as a "stream of consciousness" novel, is clearly a poor choice for technical prose. But the other extreme is an over-sectionalized, rigid document without any natural flow, such as a tax form or a legal contract. While the critical information might be present in either form of paper, neither will be well comprehended.

Consider now the style of a paper. Some bad writing is too colloquial. Consider this sentence, which may have appeared in the conclusion of the first published report of optically trapped atoms: "We trapped lots more atoms than in other experiments, and made them much colder too." Surely this sentence strikes us as "informal," but its real problem is that it is vague. Some bad writing is not colloquial enough; the same sentence might be written: "Trapping of a particular atomic species with manifold improvement in atom density and substantial reduction in atom temperature was achieved, where the improvement is reported with respect to previous experiments involving either the technique of magnetic quadrupole trap formation or the use of radio-frequency alternating currents to form a stable

potential for ionized atoms.” This sentence is not vague, and it is certainly not informal, but it is nearly unreadable. The issue separating good from bad writing style is not simply the degree to which it is formal or colloquial.

How do we find the “good” middle-ground for technical writing? In our examples above, the first, colloquial sentence was too vague; it did not deliver *enough* information to the reader. The latter sentence delivered too much information, with language that was perhaps *too* formal, resulting in a sentence that most readers have trouble understanding. “Good” scientific writing must deliver sufficient and correct information, but must rely on more than just formal structure and accurate language: it must *actively attempt to communicate* that information to the reader.

How does one effectively communicate complex information? Again, it is a difficult art learned by experience. However, a key concept to remember from the start is that one should write from the point of view of the *reader*. Remember that the reader does not simply read; the reader *interprets* what he or she is reading based on context.

In this manual, we will briefly outline some basic rhetorical principles in both structure and style to help students write clear, flowing, logical, *communicative* prose that caters to the expectations of the reader. Often, these principles will be best demonstrated with several examples of technical writing, some from typical student lab reports, and some from a Nobel-prize winning Letter from a former Stanford professor [S. Chu et al., *Physical Review Letters* **57**, 314 (1986), hereafter referred to as the “attached Letter”]. The sentences about atom trapping above, for example, were mutations of the second sentence of paragraph 13 of the attached Letter. In what follows, we follow a “top-down” approach to discussing structure and style, as this is the approach in which actual writing is best accomplished. Correspondingly, we begin with the overall structure of a typical journal article.

2 The Structure of a Scientific Article

On the largest level, the basic structure of “introduction, body, conclusion” is essential for nearly all types of technical writing. In the presentation of experimental physics research, a finer structure is expected.

The usual format in *Physical Review*, plus suggested length limits for a Letter, is as follows:

1. *Abstract*. ($\frac{1}{2}$ pg.) A concise statement of the experiment’s goals, *unique* methodology, and results.
2. *Introduction*. (1 pg.) Motivation, context, and theory necessary for the understanding of the experiment.
3. *Experimental Methods* (2 pg.) Description of physical setup and procedure, sufficiently detailed to allow the experiment to be repeated by a reasonably knowledgeable expert.

4. *Results and Discussion*. (2 pg.) The data obtained and quantitative analysis of that data. The physical interpretation of the data should be discussed, and the degree to which it agreed with the theoretical hypotheses of the introduction.
5. *Conclusion*. ($\frac{1}{2}$ pg.) A restatement of the important results, as well as applications or possible future improvements.

Before proceeding, we must address issues of length and typesetting. There are no constraints to the *minimum* length in technical writing, and indeed, the more concisely one can convey the critical information, the better. There are certainly constraints on the *maximum* length. For example, a *Physical Review Letter* is strictly limited to 4 single-spaced, two-column pages, and nearly all Letters are nearly this long. The length guidelines presented in the list above assume double-spaced, 12 pt. font on 8.5×11 " paper with approximately 1 inch margins; this format is preferred for "preprints" for ease of editing. The length limits are not strict, but if they are exceeded then a more concise rewrite should be considered.

The reason a good paper follows the above structure is because the reader expects this structure. Every sentence in a paper is read with a context in mind; the same sentence can take on different shades of meaning depending on which part of the presentation the reader believes he or she is reading. If the results are presented before the experimental methods, for example, then the reader will become confused as to what the results mean, what was known before and after the experiment, and the relevance of the ensuing discussion of methods. However, this structure is also not a strict form. Occasionally part of the results *are* presented before the methods, in order to clarify aspects of the methods section to follow.

In what follows, we describe the expected content of each section not in the order of the final presentation but rather in the rough order in which they should be written, or at least conceived.

The Thesis. This is not a section of the paper, but rather the overarching reason for it, and therefore should always be the first part planned by the author. Deciding how to order material in a paper should always follow the decision of what material to include in that paper in the first place! Any publishable scientific experiment is extremely complex, and requires the experimentalist to develop many ideas and technologies. Unfortunately, very little of this work actually makes it into the published article. (This work should still be written down in lab notebooks, dissertations, and other forms of technical documentation.) The writer must choose a *thesis*, a central argument of the paper to which all material in the paper is somehow related.

Carefully choosing the thesis is the most important and often the hardest part of writing a paper; we offer only brief advice here. Student laboratory reports will usually describe an experiment which has been done before, measuring a physical quantity that is already well known. For this

reason, students make the assumption that the goal of the experiment is to get as close as possible to the “right” answer. This situation is very different from “real” physics, however, in which the “right” answer is seldom known, and a thesis limited to describing how close the results came to this answer will rarely result in an excellent report. Most measurements are sufficiently rich that alert and thorough students may discover previously unknown effects leading to error or imprecision; frequently, the best theses focus on one such effect. There is *always* new physics to be found, even in old experiments.

As one example, in a nuclear lab performed by Stanford students for many years, students learned that the efficiency for detection via the photoelectric effect decreases with increasing energy. Yet when studying emission from ^{57}Co with a proportional counter, a peak at 14.4 keV invariably had more counts than the 6.5 keV peak. For years students wrote off this contradictory result as a statistical glitch or a result of absorption in air, without understanding or testing their assumptions. Finally in 1997 three students proved that the plastic encapsulating the ^{57}Co source was preferentially absorbing the 6.5 keV radiation, using a variety of experimental tests. Their report featured as its thesis this hypothesis and a series of convincing results, indeed finding new aspects of an old measurement. Students may not always be able to successfully discover the cause of experimental errors or anomalies, but careful awareness, study, and written communication of them will have more ultimate pedagogical value than repetition of old results.

The Title. The title should begin the paper by expressing the thesis as completely and succinctly as possible. Both the goal of the experiment and any technique or apparatus unique to the paper’s approach should be mentioned. Any excess words should be trimmed away, however, and the title should be expressed as a fragment rather than a full sentence. If the title runs more than two lines, it is likely that too much information has been included or the thesis needs to be more narrowly defined.

Consider these example titles:

The Proportional Counter
The Proportional Counter and its Uses
Counter Resolution and Dependence of Gain on Bias Voltage

The first reveals nothing except one of the tools used in the experiment. The second is equally unclear, and the vague word “uses” implies that the thesis is still too broad for an effective paper. The third title is better, but includes two different concepts that would be hard to tie together; one of the topics should be dropped. The title

Energy Resolution of a Proportional Counter

however, immediately shows that the author has explored the energy resolution of the counter, while

Variation of the Energy Resolution of a Proportional Counter with Bias Voltage

is even better, highlighting exactly what quantities were measured and hinting at how this information could be useful (i.e., helping the reader choose a voltage that optimizes energy resolution).

The Abstract. The abstract should also begin with a statement of the thesis, more detailed and less abrupt than the title. Any important results or novel experimental techniques should also be included. The reader should gather the paper's important points from just the abstract, since in reality this may be all that most colleagues will read. Nonetheless, the shorter the abstract, the better.

If the author has not narrowed his or her focus to a clear thesis, the abstract will reveal it. Here is an example from a student report:

Abstract

Using the 6.5-keV and 14.4-keV spectral lines of a ^{57}Co source, I found that the energy resolution of our proportional counter varies from $6.8 \pm 0.9\%$ to $13 \pm 4\%$. The resolution improved as the energy of incident radiation is increased and also as the voltage applied to the proportional counter is raised. The other half of the experiment verified an inverse square theory that I will propose in this paper. The experimental data indicated that the radiation detection rate is approximately equal to a constant, which represents the background radiation rate, plus a term proportional to the inverse source-detector distance squared.

There are two theses in this abstract: the study of the energy resolution of the counter, and the verification of the inverse square law. The two are dissimilar enough that the student should concentrate on only one of them. Since the student has done at least two useful sets of measurements for ΔE (vary E , vary bias voltage), and since the inverse square law is nothing new, the student may want to focus the whole paper on the energy resolution studies, and cut out of the abstract everything after "voltage applied...is raised."

The Conclusion. The conclusion should echo but not directly copy the abstract. It should restate the results cited in the abstract, but also summarize the limitations of the experiment, such as the total uncertainty of the results and the dominant source of this uncertainty. The author may then wrap up the paper by discussing future work planned on the subject, possible applications of the work, or suggestions for other techniques. All of this information helps the reader quickly absorb the important quantities, determine if they are useful in his or her research, and decide if the precision of the result can be improved by other techniques.

The most important consideration in writing the conclusion is consistency with the thesis. Here are the abstract and conclusion from a student report on sodium iodide scintillators:

Abstract

We explain the origin of the hitherto unexplained 662.71 keV peak in the scintillation spectrum of ^{60}Co . Experimental measurements of the peak conform perfectly with the theoretical expectation for the energy absorbed when one of the annihilation photons from the 1.173 MeV pair production event escapes detection. We present similar explanations for other features in the ^{60}Co spectrum, and verify them within experimental error.

Conclusion

Notable features in the ^{60}Co spectrum can be predicted based on interactions of radiation with matter inside the scintillation detector. Certain features which are expected are shrouded due to Bremsstrahlung radiation and require a source with better beta absorbers. The spectrum of ^{60}Co has several effects superposed on each other, therefore, it would require a detector with superior energy resolution than 10% in order to fully verify the more subtle features expected.

This example was written by a harried student attempting to write a report straight through in one sitting. As new ideas occurred to her, her focus changed midstream. The abstract puts the most emphasis on explaining an anomalous peak at 662 keV as a single escape peak from pair production. As she wrote, however, she realized that her evidence was not as convincing as she first believed (in fact, the peak was from a ^{137}Cs source sitting nearby). By the time she wrote the conclusion, this topic had been dropped entirely. Furthermore, the vagueness of both paragraphs (no other particular feature is discussed) makes the rest of the paper even more difficult to read. This example should emphasize the importance of planning: one should *begin* by choosing a clear thesis and making an outline for the paper with a logical structure to support that thesis. Having done this, one might choose to write the abstract and conclusions first.

For a more positive example, compare the abstract and final paragraph of the attached Letter, further discussed in Sec. 4.

The Introduction. The introduction must, like the abstract and conclusions, eventually declare the thesis clearly. Sometimes it will do so using the same language as the abstract; this is acceptable. But the introduction should do more than this: it should give the “why” of the experiment, its purpose and theoretical motivation. In most articles, a summary (and subtle judgement) of other related research is a critical component of the introduction. Key equations and concepts necessary to understand the rest of the paper are highlighted here. The section should be summed up with a restatement of the thesis, and possibly a sentence or two on the plan of action as a smooth transition to the experimental methods.

The introduction must not dominate the paper, however, as it usually presents work not originated by the author. Concision and clarity can be made consistent here by first minimizing the concepts to those strictly needed for what follows in the paper, and citing the other papers and texts from which the presented concepts and equations originate. An introduction which rederives well-known results or delves too far back into the history of a line of research will take up too much of a paper's two most valued resources: space, and the reader's attention.

How much knowledge of the reader should the author assume? In professional writing, this question can be difficult to answer. For reports in Physics 107, students should consider their audience as their peers, so if the student knew a given concept before preparing for the lab, then they should assume that the reader knows that concept when reading the report.

For an excellent example of an introduction, see the first two paragraphs of the attached Letter, further discussed in Sec. 4.

In contrast, consider this example from a student laboratory report about measuring the linearity of a proportional counter:

Introduction

The proportional counter is one of the most commonly used devices for recording the energy of X-ray radiation. The anode wire in the center of the counter is kept at a high bias voltage, and in the case of the RSG-61 proportional counter, the maximum voltage is 2 kV. However, if the voltage does not fall within a certain range, the proportional counter will not behave proportionally. Ion-electron pairs will not be accelerated enough to create the collisions that produce more ion-electron pairs, and the pulses coming from the counter will not be proportional to the energy of incoming particles.

Because a proportional counter that is not in its proportional region is next to useless for gathering meaningful data, it is important to understand what levels of bias voltage will yield proportional results. ^{57}Co is a radioactive isotope with well-known spectra, producing X-rays with energies of 6.4 keV and 14.4 keV. If the counter is behaving proportionally, the ratio of energies recorded by the counter at its peaks should be $14.4/6.4$, or 2.25. By altering the bias voltage, and recording the output, we are able to see whether the counter is in a proportional region.

The first paragraph has little purpose other than to define proportionality. Unfortunately, the language suffers from constant repetition of the word "proportional," not approaching a physical definition until the last sentence, in which "proportional region" is only indirectly defined. Details such as "...in the case of the RSG-61 proportional counter, the maximum voltage is 2 kV" are not appropriate for this introductory paragraph. In particular, more background material could be added to this introduction; in particular the reader may be wondering what physical processes *limit*

the proportional range. The reader might also wish to know whether the current results are limited to the RSG-61 counter or whether they are more general; if they are specific to one kind of counter, perhaps the scientific relevance of that specific counter could be more explicitly discussed.

The second paragraph provides an appropriate summary of why the student makes the measurement, why ^{57}Co is chosen for it, and how the measurement will be made. Further details are not needed for this part of the discussion; indeed this paragraph transitions nicely to the ensuing methods section. However, precise language is always important; in particular the words “see whether” are vague and informal, whereas a more precise phrase such as “test the lower limit of...” would better clarify the scope of the experiment for the reader.

Experimental Methods. The experimental methods section contains descriptions of both the experimental setup and the types of measurements made. Some explanation of why certain measurements are made (what specific effect the experimenter hopes to study) is also appropriate. A good rule of thumb for the methods section is to include just enough detail to allow an *expert* reader to repeat the experiment. Excessive description of common equipment will waste space; in a published paper, it will often be only those ingredients that are new or novel that will be described with any detail. Certainly, only those experimental ingredients that are relevant to the thesis should be described at all.

A frequent problem with student writing is the use of time-order words. Writing “first we did this, then we attached that” makes the section sound like a recipe, with no concern for why the procedure was done this way. Also, chronological order is often not the best way to understand a process. A sentence such as, “First we connected the preamp to the amplifier, then linked the amp to the MCA, then hooked up the proportional counter,” has no logical flow. If the writer instead follows a signal through the system, as in the sentence, “An x-ray triggers a current in the counter, which is enhanced by the preamp and amplifier and read by the MCA,” the reader sees more clearly the purpose of each piece of equipment in the procedure.

Further improvements require more subtlety. Diagrams, for example, are almost always useful to clarify a complicated setup, but two issues must be kept in mind. First, if parts of the diagram will be discussed in the text, always refer the reader’s attention to the figure first. Otherwise the reader is confused and must scan the text again to connect it to the diagram. Second, if the setup is easily described in words, a diagram becomes unnecessary. For the sample sentence in the previous paragraph, an accompanying diagram would only consist of four boxes labeled “counter,” “pre-amp,” “amp,” and “MCA,” connected in series by straight lines. This diagram neither adds to nor clarifies the text, and should not be included in the paper.

Taking all of the above suggestions into account, we can examine a real methods section written by a Stanford student:

Experimental Methods

We used a Reuter-Stokes RSG-61 x-ray proportional counter tube to measure the x-rays produced in the decay of the radioactive cobalt. The signal produced by an event in the detector—a burst of current proportional to the energy of the radiation causing the event—was passed through a series of amplifiers, which converted the current pulse to a voltage pulse, and amplified it to a measureable size. This pulse was then directed to a multi-channel analyzer (MCA), which measured the number of signals at each of a range of voltages. Using an MCA, we measured the shifts in the number of x-rays of each energy recorded by the detector. Figure 1 shows our experimental setup.

Radioactive decay of ^{57}Co produces gamma rays at 136.5 keV, 122.1 keV, 14.4 keV, and 6.4 keV; our detector, however, was only able to detect the latter two of these. The higher energy gamma rays are presumed not to have had any effect on our data and were ignored.

First we placed the cobalt source on the window of the proportional counter. Then we hooked up the pre-amp, amplifier, and MCA. We set the RSG-61 bias at 1700V and ran for five minutes. Then we increased the voltage to 1750V and ran for another five minutes. We repeated this process up to 2000V. Each time we saved the spectrum to the MCA's memory for later study.

The above example suffers from a lot of unnecessary information. In the third paragraph, the order in which the equipment was hooked up is a less informative echo of the first paragraph, and the voltages used can be expressed without specifying which was tested first. The workings of a proportional counter and an MCA, while new to the student, should be known to peers, so an explanation of how they work and why the higher energy x-rays are ignored is not needed. Finally, as we mentioned before, Figure 1 in this instance will probably add nothing to the paper. Thus this sample section could be trimmed to the following:

Experimental Methods

A Reuter-Stokes RSG-61 proportional counter tube was used to measure the 6.4 and 14.4 keV x-ray decays of ^{57}Co . The detector signal was processed by a pre-amplifier and shaping amplifier and read out by a multi-channel analyzer (MCA). The source was placed on the center of the detector's beryllium window. Spectra were measured in five minute intervals for bias voltages from 1700 to 2000 volts in steps of 50 volts.

Results and Discussion. The results section follows naturally from the methods. It is natural for a reader to ask, after the description of the experiment, "so, what happened?" The discussion section answers the naturally ensuing question, "and what does it mean?" In these sections a crucial consideration is the maintenance of a logical flow, directed by the central thesis.

Unordered lists of data, even if explained with verbose discussion, will be insufficient if the reader does not follow the logic of its presentation.

In the presentation of numeric data, tables are very important. Measured numbers should almost never appear in the text of the paper; a string of numbers is difficult to digest in one read and not as revealing as when they are placed side by side. Consider the following student example:

Peaks were apparent around 1240 and 2700 channels, corresponding respectively to the 6.4 keV and 14.4 keV emission lines. Their ratio was found to coincide (within the uncertainty of the measurements) with the anticipated value of 2.25 (14.4/6.4). Although it was encouraging to find the peaks so well related, the resolution was disappointing after the results of our calibration; the 6.4 keV line had a ΔE of 17%, the 14.4 keV peak 10%.

Here it is difficult to tell which numbers are data and which are conclusions. In this case, the calculated ratio of 2.25 is the important number. The other values, even though there are only 6 of them, should be set off in a table, and the text should simply cite the table and state the inferred result:

The counter output for the two prominent ^{57}Co X-ray peaks are listed in Table I. The count ratio was found to coincide (within the uncertainty of the measurements) with the anticipated value of 2.25. Although it was encouraging to find the peaks so well related, the resolution (ΔE) found after calibration was disappointing...

Channels	Corresponding Energy	Resolution (ΔE)
1200	6.4 keV	17%
2700	14.4 keV	10%

Table I: Prominent ^{57}Co X-ray peaks.

The use of a table not only simplifies the accompanying text—it also provides the opportunity to visually offset the key results from the related discussion. Tables should appear with well-labeled headings, all units labeled at least once, and a short, descriptive caption or title.

Figures For the presentation of data, a table is better than text and a graph is better than a table. This important point is best revealed by a simple example. Consider this exciting passage:

As we heated the temperature of the resistor from 300 K to 305 K, the RMS voltage increased from $12.86 \pm 0.05 \mu\text{V}$ to $12.9 \pm 0.04 \mu\text{V}$. As expected, the RMS voltage continued to increase as we raised the temperature, reaching $13.15 \pm 0.07 \mu\text{V}$ at 310 K, continuing to $13.18 \pm 0.04 \mu\text{V}$ at 315 K, and ending up at $13.30 \pm \pm 0.06 \mu\text{V}$ at the highest temperature we went, 320 K.




Figure 1: Johnson Noise vs. temperature

While the principles of story-telling can help in the construction of clear prose, it is rarely effective in the presentation of data. In this case, a table is more concise but still not illuminating:

Temperature (K)	RMS Voltage	st. dev.
300	12.86	0.05
305	12.90	0.04
310	13.15	0.07
315	13.18	0.04
320	13.30	0.06

Table II: RMS Voltage measured at five temperatures

Most people have trouble seeing even the simple pattern in a list of numbers like this. However, even a poorly drawn graph, as shown in Fig. 1, shows the point of the data. There are, however, many problems with the graph in Fig. 1, and all of them may be seen in student writing and (unfortunately) in too many professional presentations. The problems include

1. The axes are not labelled; both the physical measurement and the units should be clearly labelled.
2. The caption does not explain the graph; although it provides a title that gives us some idea of what is being shown, it should briefly let us know what this data is and what measurement yielded it (without being too long or redundant to the text.)
3. The lines through the points are inappropriate and undescribed. Is this a fit? Is this a spline? A fit would be most appropriate; curved lines to “guide the eye” should be used sparingly.

4. Error bars, though known, are not featured on the plot.
5. Superfluous material irrelevant to the information presented (“Series 1”) is included.

In addition to these, it is wise to begin making figures that work well when printed: Fig. 1 is unnecessarily filled in, rather than leaving a white background, and it includes a superfluous box around it. An improved version of this graph is shown in Fig. 2.




Figure 2: The average RMS voltage measured across the 10 MΩ resistor is shown against the temperature of the resistor. For this data, the bandwidth is 1 kHz. The error bars show the standard deviation from five independent measurements. The dashed line is a linear fit to the data, showing a slope of 22.0 ± 0.1 nV/K.

For more examples in the use of figures, consult the attached Letter. In the methods section, for example, the authors present as Fig. 1 graphs of the expected potentials experienced by atoms due to applied laser fields. These pictures are surely more valuable than a written description; they are also easier for the reader to understand than a mathematical formula. Note that all axes are clearly labelled with titles and units.

Do not hesitate to clarify the presentation with figures. Besides graphing physical data and mathematical functions pertinent to the discussion, diagrams may be used in the introduction and discussion sections to explain concepts, while schematics are often the clearest way to describe an apparatus in the methods section.

Often, the explanation of figures provides the clearest and easiest logical structure for the text. This occurs in the results section of the attached Letter. Starting with paragraph 7, the results section has the general form of explaining what each figure shows, and then discussing what it means.

This structure is very convenient for both the writer and the reader. When structuring a section this way, always refer the reader to figure first.

Nearly all papers would benefit from more figures, but unfortunately in a Letter they take up much space and must therefore be judiciously chosen.

The structure of the results and discussions section can vary greatly depending on the nature of the paper's thesis. Indeed, the internal structure of *any* section may be written quite differently depending on the nature of the work. Common to all papers, however, is the *style* in which these sections should be written, a subject to which we now turn in more detail.

3 Style

Our discussion of the general structure of an article has emphasized logical flow; such flow does not come automatically from the structure, however, but rather from the construction of paragraphs, sentences, and even individual word choice. In this section we offer helpful suggestions for maintaining a clear, flowing style.

Students should be reminded that writers are often the poorest judges of their own writing. Meaning is lost to the reader for two reasons: first, and most obvious, the reader does not have the prior knowledge possessed by the author. Second and more subtle, however, are the contextual interpretations added by the reader. Each sentence presents an expectation for the next, and this expectation can color the meaning of that next sentence. The same system of expectation is true within a single sentence. This effect can be detrimental or advantageous, depending on the author's rhetoric. Often, stylistic problems can only be detected by a reader separate from the author, emphasizing the importance of finding someone to read your writing before final editing.

3.1 Simplifying Subjects and Verbs

The first rhetorical principle that cannot be sufficiently emphasized is simplicity. The concepts to be explained in physics are sufficiently complicated that complex, flowery language is more likely to hinder communication than help it. The author of a physics paper should seek the most direct, simple way of stating his or her meaning.

Joseph Williams (see Sec. 5) has argued that the most effective writing follows the basic principles of narrative prose. While technical writing should not sound like a story written for children, it does tend to be most effective when composed as "characters" that perform "actions." In technical reports, those characters are rarely people; they may be instruments, samples, or even ideas or theories. However, whatever they are they should be identified and, if it is reasonable to do so (bearing in mind Sec. 3.2), made the subject of the sentence. The "actions" these characters perform should be the verb of the sentence. Keeping characters and actions clear in writing leads to very readable prose. A related concept is to ask what the important

stress is in each sentence, and make sure it is presented in an appropriate part of the sentence for the reader to perceive it as such.

Consider this abstract we already saw in the section on structure on page 6. Try to follow the long second sentence:

Experimental measurements of the peak conform perfectly with the theoretical expectation for the energy absorbed when one of the annihilation photons from the 1.173 MeV pair production event escapes detection.

Cutting the unnecessary words “experimental” (already implied by “measurements”), and “perfectly” is not enough to make this readable. The real problem here is the presence of five prepositional phrases, all acting as very long-winded nouns: “Experimental measurements **of the peak**,” “the theoretical expectation **for the energy absorbed**,” “one **of the annihilation photons from the 1.173 MeV pair production event**.” These overblown phrases are called *nominalizations*, and they obscure the reader’s ability to identify characters or actions, and the important stress of the sentence.

What is the subject of our sentence? It seems to be “experimental measurements,” but are “theoretical expectation” and “one of the annihilation photons” both objects of “measurements?” No, the latter is the object of “theoretical expectation;” the theory predicts something for the photon. And what does it predict? The “measurement.” So if we let “theoretical expectation,” or more simply just “theory,” be our subject “character” and open our sentence, the other two nouns can follow as objects.

Our “action” verb will be whatever our subject, “theory,” is doing, which in this case is “predicting” a result. The main object, the thing being predicted, was originally written as “experimental measurements,” but is there a better way to phrase this? What is being predicted? The answer is the energy measurement given at the end of the first sentence, 662.71 keV. So we can simplify our first object to just “this energy,” referring back to the previous sentence.

Now let’s place our problem sentence, with the subject, verb, and main object more clearly defined, back in the abstract:

Abstract

We explain the origin of the hitherto unexplained 662.71 keV peak in the scintillation spectrum of ^{60}Co . Theory predicts this energy when one of the annihilation photons from a 1.173 pair production event escapes detection. We present similar explanations for other features in the ^{60}Co spectrum, and verify them within experimental error.

The abstract can now be digested in one pass.

3.2 Old, then New

The second important rhetorical principle in constructing sentences and paragraphs is that that new concepts should *follow* previously presented

concepts. Although this principle sounds obvious, it is frequently abandoned even in professional writing.

What is the problem with this introductory paragraph?

Radiation from ^{60}Co interacts with matter via both photoelectric absorption and Compton scattering. All of the energy of an incident gamma is absorbed by an electron, causing a scintillation, during photoelectric absorption. The full energy peak, or "photopeak," of the spectrum corresponds to the energy absorbed by these electrons. A broad range of energies, however, may be transferred from the gamma to the electron, depending on scattering angle, during Compton scattering...

When we read, we try to link each new piece of information with what we saw previously. In this case, after digesting the topic sentence our brain is looking for descriptions of the two main objects, "photoelectric absorption" and "Compton scattering." But in the later sentences, these familiar names are not mentioned until the end. When we reach the final few words of each sentence, we realize that what we just read, already partially forgotten, was important in relation to the topic, and we are forced to go back over the sentence again to forge this link properly. Astute writers always begin sentences with known information before introducing ideas, signaling the reader immediately as to how these new facts are connected to the old. See how simply placing the topical terms first in each sentence smoothes the flow of information:

Radiation from ^{60}Co interacts with matter via both photoelectric absorption and Compton scattering. During photoelectric absorption all of the energy of an incident gamma is absorbed by an electron, causing a scintillation. These electrons correspond to the full energy peak, or "photopeak," of the spectrum. During Compton scattering, however, a broad range of energies may be transferred from the gamma to the electron depending on scattering angle.

Note that one reason the passive voice exists, and the reason it is so popular in scientific writing, is that it allows old information to be shifted to the front more easily. Active sentences like "A broad range of energies **characterizes** Compton scattering" can be shifted to passive sentences such as "Compton scattering **is characterized** by a broad range of energies." Many students are told to avoid the passive voice in grammar and high school; surely excessive use of it will muddle any writing, but complete avoidance of it will prohibit clearly flowing technical writing.

3.3 Less is More

Many writers fear that their work sounds too simple or naive when their sentences are too short. This defensiveness backfires when too many long,

complex sentences render a text unreadable. When we struggle to finish reading a sentence, it is most often because we get lost in wordiness halfway through. There is nothing inherently wrong with long sentences; if the characters and actions are matched to objects and verbs, as discussed in Sec. 3.1, even long sentences can be readable. However, most long sentences in student writing are decorated by unnecessary material which only lengthen the gaps between concepts.

We use more words than we need in several ways, and good writing requires scanning all of our sentences and removing the fluff. We can demonstrate many of these bad habits with the following paragraph from the methods section of a scintillation lab:

As noted earlier, our experiment was basically composed of three physical elements: the radiation source, the absorption material, and, of course, the detector. From the detector, the given signal was amplified and captured using a standard amplifier and multi-channel analyzer, where each and every energy was assigned to one of 4096 discrete bins to be subsequently analyzed.

Meaningless words. There are many words we use automatically to soften the abruptness of information flow in our sentences, but which have no real meaning. Words such as “virtually,” “basically,” “practically,” “various,” and so on, give some shape to spoken language but are out of place in scientific writing. In the example above, we could cut words from “was **basically** composed,” “and, **of course**, the detector,” “the **given** signal,” and “to be **subsequently** analyzed” without loss of content.

Repeated words. English includes several phrases consisting of two words with identical meanings connected by an “and.” A few are “First and foremost,” “one and only,” and “various and sundry.” In almost all cases only one of the words is needed; in our above example, “each and every” should be reduced to just “each.”

Richard Feynman tells a story of a writing exam on “The Importance of Science in Aviation” in his high school days:

I was about to write a simple theme about this dumb question when I remembered that my literary friends were always “throwing the bull”—building up their sentences to sound complex and sophisticated. I decided to try it, just for the hell of it... So I wrote stuff like, “Aeronautical science is important in the analysis of the eddies, vortices, and whirlpools formed in the atmosphere behind the aircraft...” —I knew that eddies, vortices, and whirlpools are the same thing, but mentioning them three different ways *sounds* better! ... The teacher who corrected my examination must have been impressed by eddies, vortices, and whirlpools, because I got a 91 on the exam.

—Richard Feynman, *What Do You Care What Other People Think?*, 1988.

Feynman went on to explain that his high score may have been due to his English teacher's lack of ability to fully understand the topic. "Throwing the bull" to either unnecessarily decorate or even obfuscate a topic is highly discouraged in technical writing.

Unnecessary modifiers. Many words imply others, yet we include both words in our sentences and overburden the reader. Often these extra descriptive words take the form of adjectives or adverbs: "**completely** finish," "**final** outcome," "**future** plans." The best example in our passage above is "**physical** elements", where "elements" is clearly sufficient. A slightly trickier situation is when unnecessary modifiers appear as nouns or prepositional phrases, as in "patient **manner**" or "rectangular **in shape**." Omitting these words usually requires altering the modified word in some way. "Absorption **material**," for example, can be changed to simply "absorber."

Throat-clearing. Many tired phrases exist to help ease the transition of starting a new sentence. "In addition," "and therefore," "in a manner of speaking," and "I would like to mention" are a few common examples. This throat-clearing in preparation for writing is unnecessary in a scientific paper; just get to the important information. Thus, "As noted earlier" can be lifted from the above paragraph.

Combining all of these suggestions give us a leaner, clearer opening paragraph:

Our experiment was composed of three elements: The radiation source, the absorber, and the detector. From the detector, the signal was amplified and captured using a standard amplifier and multi-channel analyzer, where each energy was assigned to one of 4096 discrete bins to be analyzed.

These sentence-level examples of repetition are fairly easy to avoid or edit away. A more difficult problem occurs when the writer unnecessarily repeats concepts in order to fix improperly structured flow. Recognition, diagnosis, and correction of this kind of repetition can be challenging, but is a critical skill for preparing technical documents.

We demonstrate with an example from a student introduction concerning atomic energy levels:

Introduction

Through measuring the wavelength of emitted light, it is possible to determine both the energy level of the orbiting electron and the specific substance emitting the light. When an electron circling an atom's nucleus changes from a higher energy orbit to a lower energy orbit, the atom emits electromagnetic radiation of a specific wavelength or color. This emitted wavelength depends only on the change in energy level. As a reminder, under the Bohr model, electrons can remain only at certain quantized energy levels when they are orbiting a nucleus. Elements and their isotopes can be differentiated since each has unique energy levels at which the electrons are stable. In this experiment we found four wavelength ranges emitting light, α , β , γ , and δ , each range consisting of a hydrogen and a deuterium peak (a heavy isotope of hydrogen with two neutrons instead of one). These peaks correspond to the energy transitions as an electron moves from a higher energy level, $n = \text{an integer}$, to a lower energy level, $n = 1$. α , β , γ , and δ ranges of wavelength match with transitions starting from $n = 2, 3, 4$, and 5 . This quantized range of emitted wavelengths of light is called a Balmer series.

Balmer and other founders of quantum mechanics have discovered mathematical relations to describe the wavelengths of hydrogen spectral lines. Using these equations, it is possible to calculate the ratio of the mass of an electron to a proton based on the reading for the wavelength difference between hydrogen and its isotope deuterium.

$$\frac{\Delta\lambda}{\lambda_{\infty}} = \frac{1m_e}{2m_p}$$

From the wavelength difference (or isotope shift) and a calculated value of the wavelength of light emitted from a particle of infinite mass, this ratio can be determined. It makes sense that masses are related to the change in wavelength, since this isotope shift is simply the result of an additional atomic weight. For a more in depth coverage of these principles, any introductory text on modern physics may be of use.

Just understanding what the author is trying to say probably took more than one read. We will see that much clarity may result from simplification. This introduction distills into the list given below:

1. Spectrum tells energy level and atom type
2. Energy level determines wavelength
3. Wavelength depends on energy level
4. Bohr model says electrons are in energy levels
5. Energy level depends on atom type
6. We saw four lines
7. Lines depend on energy level
8. Called Balmer series
9. Balmer predicted lines
10. $\Delta\lambda/\lambda_{\infty} = m_e/2m_p$
11. From $\Delta\lambda$, can find m_e/m_p

12. This makes sense

Now we can more clearly see some of the problems. There is much repetition; lines 3 and 7 echo line 2, and lines 1 and 5 are also similar. There are concepts out of order, such as the definition of an energy level (line 4), which should come before energy levels are discussed. Finally, there is some unnecessary information about Balmer which does not directly affect the measurement. By removing the redundant and unnecessary lines and rearranging the rest, we can make a much stronger argument while still preserving the original author's voice:

1. (Prev. 4) Bohr model says electrons are in energy levels
2. (Prev. 5) Energy level depends on atom type
3. (Prev. 2) Energy level determines wavelength
4. (Prev. 10) $\Delta\lambda/\lambda_\infty = m_e/2m_p$
5. (Prev. 11) From $\Delta\lambda$, can find m_e/m_p

To this restructuring, the original text has also been edited to follow the principle of "old, then new." A possible introduction now reads:

Introduction

Under the Bohr model, electrons can remain only at certain quantized energy levels when they are orbiting a nucleus. These energy levels are unique to each isotope of each element. When an electron circling an atom's nucleus changes from a higher energy level to a lower energy level, the atom emits electromagnetic radiation of a specific wavelength or color. This wavelength depends only on the change in energy level, allowing different isotopes to be distinguished.

This Bohr model allows the calculation [citation] of the ratio of the mass of an electron to a proton based on the observed wavelength difference, $\Delta\lambda$, between hydrogen and its isotope deuterium. In terms of a calculated value of the wavelength of light emitted from a particle of infinite mass, λ_∞ , this ratio can be determined from

$$\frac{\Delta\lambda}{\lambda_\infty} = \frac{m_e}{2m_p}.$$

The author's discussion of the taxonomy of Balmer spectral lines is more appropriate for the methods section of the text, rather than this introductory section. Further discussion of the equation is not really necessary, since it may be found in an appropriate citation.

3.4 Other Bad Habits

Over-selling. A typical problem with introductions and conclusions is the use of sweeping gestures and grand language to describe the importance of the paper. Many students feel that every scientific paper must include a result of earth-shattering importance to be worthy of publication. But even the most important papers are written in so matter-of-fact a style

that without the benefit of hindsight, the reader would never know their historical importance. The writer should present his or her findings and possible applications, and leave the reader to decide the value of the results.

Over-hedging. Physicists try to avoid absolute statements, especially in their conclusions. Feigning 100% confidence that your results are incontrovertible can antagonize colleagues, leaving them even more skeptical of your report. But students can tend too far in the other direction, resulting in sloppy sentences: “We believe that it is possible, under certain circumstances, to carefully align the source in such a way that the observed effect may partially disappear...” One or two of these hedging words is almost always enough to express caution without needlessly complicating a sentence.

Colloquialism. Words like “see,” “look at,” and “check” are part of our spoken language, but are not useful if they alone describe an observation. Imagine a methods section that stated, “The signal was plugged into the amplifier.” Where? In the bipolar or unipolar output? If this is a critical step if the experiment is to be successfully repeated, more precise information is necessary than “plugged in.” If not, the sentence need not be present at all.

Some academics believe that the use of the first person (“I” or “we”) is too colloquial. Over the centuries the direct use of the first person has slowly diminished in scholarly writing. In modern physics literature, however, the first person plural (“we”) still appears frequently, and it should therefore be used as the subject when the authors are the appropriate “character” of the sentence. A more difficult issue is the case when a paper has only a single author. Should the author use “I”? Few modern single-authored papers do, instead employing either “we” (presumably referring to the author and the reader) or avoiding the issue altogether with passive voice. Our advice is to not let this awkward element of scholarly style impinge on clarity, but rather to use the most appropriate subject for the message of the sentence.

Repetition for Clarity or Emphasis A common problem in student writing, especially in the results section, is when the author attempts to emphasize or clarify a point by repeating the same idea using different words. Attempting to explain a concept multiple ways is useful for writing a draft, but of these ways only the clearest should remain in the final draft.

3.5 Improving Your Writing

Ultimately, clear writing cannot be reduced to a finite list of rules. The gateway to being a good writer is to first be a good reader. A student can greatly improve his or her own writing ability by taking note of when material they

are reading is particularly clear or unclear, and rereading critically to find out why.

4 A Professional Example

We have discussed the attached Letter some above; we now take a closer look. Written by former Stanford professor Steven Chu and coworkers, this paper is the first published report of the optical trapping work that led to Chu's Nobel prize. Comments in *italics* are from Chu himself about this paper.

Title. *Descriptive, simple.* The quality of this title is perhaps revealed by counterexample: the previous *Letter* describing the generation of "optical molasses" which made the trapping work possible had the complicated title "Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure." Chu writes that this title *is not a good title because it is too convoluted. A better title would be "The demonstration of 'Optical Molasses.'"* It would have attracted more readers and would have been better for historical purposes as well.

Abstract. *The abstract can be slightly technical, but remember it is the part of the paper that wants to lure the reader into reading more, so if it is too specialized, it will tend to scare away all but the few experts in the field.*

Note the extremely direct language of the abstract. (e.g. "We report the first observation. . ." instead of the passive voice.) Note also the ordering of concepts. For example, the second sentence could have been written "A dipole-force optical trap created by a single, strongly focused, Gaussian laser beam tuned several hundred gigahertz below the D_1 resonance transition captured sodium atoms cooled below 10^{-3} K in 'optical molasses'." This is the same material as in the original sentence but it is extremely difficult to read because the "action" of the sentence, "captured", is buried among long clauses. Here, using the passive voice clarifies the action of the sentence; it also emphasizes the topic "atoms," which is the topic of the next sentence as well.

Introduction. Before proceeding, we note that there are no headings in a *Letter*; the introduction is not labelled as such. This absence is in contrast to regular articles in *Physical Review*; it is a matter of style with origins in space limitations. Sufficiently good writing should flow through section divisions regardless of the presence or absence of heading labels; further, the nature of the section should be obvious to the reader without headings. Headings should only be present to act as navigational guides to locate material in a longer article.

Even without headings, it is obvious that the first paragraph of our example *Letter* begins an introduction. After repeating directly the thesis of

the paper, it provides the scientific context of the work. This disclosure of previous research often serves three goals: first, it clarifies the problem or question that the research is attempting to solve; second, it references the important work of colleagues that led to the present result; and third it “dismisses” previous work, showing why the present results are different or even more significant.

The second paragraph continues to be introductory; it is a *description of why previous workers had trouble trapping atoms and what new developments have contributed to the success*. Notice the space limitations do not allow us to describe the details of how optical molasses works. Beyond its content, this is a well structured paragraph; the first sentence provides a clear topic for the ensuing discussion. After reading the three reasons that trapping is difficult, the reader expects to hear resolution to each reason in order, and the authors deliver on this expectation, leading the reader smoothly into the subject of optical molasses as an evasion of the third difficulty.

Experimental Methods. The third, fourth, fifth, and sixth paragraphs provide a *detailed description of the actual trap used, and why we chose this trap versus others that have been proposed*. This section describes the workings of the trap, the basic physics ideas, and the all-important damping (laser cooling) of the atoms in the trap. Note that previously derived equations are only referenced and parts of the apparatus detailed in a previous paper are only briefly discussed. The sixth paragraph describes the working procedure, and is the most quantitative of the paper.

Results and Discussion. Paragraphs 7-11 provide a *description of the diagnostics of the trap*. This part of the paper answers the questions (1) *how did we know it was working*, (2) *how does the trap performance compare to what we expected*, (3) *how hot were the atoms in the trap*, and (4) *how many atoms were trapped*.

Recall that the thesis of this paper is the demonstration of a physical effect, so the primary results expected by the reader should be evidence that the physical effect occurred. A short topic sentence such as “Several tests were made to confirm trapping,” at the head of paragraph 8, combined with the reader’s expectation for such a topic, efficiently clarifies the logic of the presentation.

Paragraph 12 is *very short, but the first mention of the “optical tweezers” aspect of the trap*. This is mentioned last since it is not primarily a “result” related to the central thesis, but rather an indication of important future developments.

Conclusion. Many conclusion sections begin with “In conclusion,” “In summary,” “To conclude,” or similar, and paragraph 13 is no exception. Although the first sentence sounds repetitious to the abstract and the introduction, this is permissible as it reinforces the central thesis. This paragraph *summarizes what has been done, how things can be improved, and mentions*

where this work might lead in terms of applications. Note the tone in which the authors comment on the importance of the results; the degree to which this result exceeds previous ones is listed quantitatively, while the speculative comments for the future are stated with neither excessive certainty nor excessive doubt.

5 Further Reading

This guide has been intended only to help students begin thinking about effective technical writing. However, the perfection of such skills is a lifelong pursuit. We recommend several references for further advice:

- George D. Gopen and Judith A. Swan, “The Science of Scientific Writing,” *American Scientist* 78, 550–558 (1990). This article (currently available online) is an excellent place to start to read modern thoughts about scientific writing. Much of the current guide has been based on this article, whose emphasis is explained by its subtitle: “If the reader is to grasp what the writer means, the writer must understand what the reader needs.” This article is specifically about scientific writing; it is based primarily on the work of Joseph Williams, which is our second recommendation.
- Joseph M. Williams, *Style: Toward Clarity and Grace*, University of Chicago Press, 1990. For issues of style in any kind of writing, this book is extremely valuable. The writing examples—first showing problems, and then the solutions—allow one to understand with accuracy exactly what makes muddled, verbose writing so difficult to read. The first chapter on the causes of bad writing, including a “Very Short History of Bad Writing,” is well worth a read.
- Scott L. Montgomery, *The Chicago Guide to Communicating Science*, University of Chicago Press, 2003. More tailored to science than Williams, this text offers useful advice for structuring scientific writing. Chapter 2 is particularly interesting, showing how scientific writing has changed through history and providing an insightful analysis of Watson and Crick’s paper on the structure of DNA.
- The AIP Style Manual and the APS Style Manual. The American Institute of Physics and the American Physical Society both provide style manuals for authors who submit articles to their journals. Besides basic expectations for the style and structure of prose, these manuals also detail expectations for correctly typesetting equations and dealing with matters such as physical units, technical abbreviations, and the like.

Many more references are of course available, and do not forget the most valuable: your professors and teaching assistants.

6 Conclusion

We conclude by first acknowledging that clear, concise writing is very difficult, but not impossible, to generate. As an attempt to generalize and simplify the principles discussed in this guide, we suggest to the student author to write with several questions constantly in mind. First, when ordering material through the paper or even in a single sentence, ask “What is the reader expecting to hear at this part of the text?” The reader is expecting material that supports the thesis, with certain kinds of material in appropriate sections. Second, when choosing the order of words and concepts, ask “What does the reader already know, and what new information am I delivering?” Third, when composing a sentence, ask “*Who* is doing *what* in this sentence?” and whether the sentence could be phrased more clearly to reveal the answer. Finally, when editing the paper, ask yourself for each sentence, “Is there a *simpler* way to present this material?” If your writing correctly conveys the complex information appropriate for a publishable physics article but *seems* simple and straightforward, then you may call it a success.

7 Acknowledgements

Not discussed in Sec. 2 but not lacking in importance is the acknowledgements section. Any work worth writing about is done with the assistance of many people, and acknowledging those who assisted the most rarely hurts and often helps. The present authors wish to thank Claude Reichard, Rick Pam, Mark Kasevich, and Alan Schwettman for their assistance in preparing this guide.

Experimental Observation of Optically Trapped Atoms

Steven Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable

AT&T Bell Laboratories, Holmdel, New Jersey 07733

(Received 14 April 1986)

We report the first observation of optically trapped atoms. Sodium atoms cooled below 10^{-3} K in "optical molasses" are captured by a dipole-force optical trap created by a single, strongly focused, Gaussian laser beam tuned several hundred gigahertz below the D_1 resonance transition. We estimate that about 500 atoms are confined in a volume of about $10^3 \mu\text{m}^3$ at a density of 10^{11} – 10^{12}cm^{-3} . Trap lifetimes are limited by background pressure to several seconds. The observed trapping behavior is in good quantitative agreement with theoretical expectations.

PACS numbers: 32.80.Pj

- 1 We report the optical trapping of neutral atoms by the forces of resonance-radiation pressure in a single-beam optical trap. At the time of the first demonstration of stable optical trapping and manipulation of small dielectric particles¹ it was predicted that similar effects were possible with atoms. Since then there have been extensive studies of the basic forces of laser light on neutral particles and atoms.^{2–8} The trapping and manipulation of neutral macroscopic dielectric particles from $\sim 30 \mu\text{m}$ to $\sim 25 \text{nm}$ has been demonstrated.^{5,9–11} With atoms, resonance-radiation forces have been used for atomic deflection,^{2,12–14} the guiding and focusing of atomic beams,¹⁵ the slowing and stopping of atomic beams,^{16,17} and the three-dimensional cooling of atoms to a temperature of 2.4×10^{-4} K.^{18,19} While there have been many proposals of optical traps for atoms in the last sixteen years,^{1,2,5,6,8} none has been demonstrated previously.
- 2 Optical atom traps are difficult to achieve for several reasons. (i) Their potential wells are shallow, typically 10^{-1} – 10^{-2} K. (ii) Their volumes, with the exception of alternating-beam scattering-force traps,^{8,20} are quite small. (iii) Once confined within a trap, atoms are heated by the random fluctuations of the light forces and will "boil" out of the trap in a fraction of a second.⁴ The single-beam trap demonstrated here has a well depth $\sim 5 \text{mK}$ and a volume $\sim 10^{-7} \text{cm}^3$. Use of the recently demonstrated "optical molasses"^{18,19} was crucial to the present experiment. Optical molasses (OM) is formed by the intersection of three pairs of counterpropagating, mutually orthogonal laser beams tuned half a linewidth below a resonance transition. The OM forms a highly viscous medium of photons capable of cooling the sodium atoms used in this experiment to 2.4×10^{-4} K. The dense collection of atoms confined within OM provides an excellent source with which to load an optical trap. The random-walk motion of the atoms and their long storage time allows atoms to continuously diffuse to the trap surface and be captured. OM was also used to cool atoms within the trap.
- 3 The optical trap demonstrated here is the single-beam gradient-force trap proposed in 1978.³ It consists of a single strongly focused Gaussian laser beam tuned about 10^4 natural linewidths below resonance. Conceptually it is the simplest of the proposed traps and offers many advantages. It has no standing waves and consequently minimal dipole-force heating.⁴ It is capable of giving deep optical potential wells, localized to a few optical wavelengths, and is ideal for achieving high atomic densities and for optical manipulation of atoms.
- 4 The properties of the single-beam trap result from the dominance of the dipole force over the scattering force in a focused Gaussian beam. The scattering force² is due to the spontaneous scattering of photons. Below saturation, it is proportional to the light intensity and points in the direction of the incident light. Quantitatively, $F_{\text{scat}} = (h/\lambda)(1/2\tau_N)[p/(1+p)]$ where τ_N is the natural lifetime and p is the saturation parameter. Also, $p = (I/I_s)[\gamma^2/4(\Delta\nu^2 + \gamma^2/4)]$ where I_s is the saturation intensity (20 mW/cm² for the D_2 line), $I = I(r,z)$ is the Gaussian beam intensity, $\Delta\nu$ is the detuning from resonance, and $\gamma = 1/2\pi\tau_N$ is the natural linewidth. The dipole (or gradient) force,³ which arises from the atomic dipole moment induced by the laser field, is proportional to and points in the direction of the intensity gradient. It can be derived from a conservative potential^{3,4} U where $U = (h\Delta\nu/2)\ln(1+p)$. For tunings below resonance ($\Delta\nu < 0$), the dipole force pushes atoms toward the high-intensity region of the light. This assures the radial stability of the trap. In the axial direction one has the additional complication of the scattering force. However, axial stability can be achieved since the axial gradient force can dominate over the destabilizing influence of the scattering force for tunings well below resonance and for strongly focused beams. The above considerations on the scattering- and dipole-force components on atoms are fully consistent with conservation of momentum in the optical-scattering process.^{1,4,11} The dipole potential U can also be understood in terms of the optical Stark shift of an atom. A laser tuned below resonance lowers the ground-state energy of an atom by an amount U and raises the energy of the excited state by the same

amount. Since an atom spends most of its time in the ground state its total energy is minimized in the region of the maximum laser intensity. Figure 1(a) shows the potential well along the z axis for three laser detunings. Figure 1(b) shows the axial well depth versus detuning below the sodium D_2 transition. It is the presence of the axial scattering force which accounts for both the asymmetry of the axial potential well and the fact that the depth of the axial potential is always shallower than the corresponding transverse well depth. The basic principles of single-beam gradient traps fed by viscously confined particles were demonstrated¹¹ in experiments on the trapping of Rayleigh particles as small as 25 nm in water.

5 Damping of the atomic motion in the trap is accomplished by periodically turning the trap off and applying OM by use of square-wave modulation.²¹ This avoids any interference with the damping process due to the optical Stark shift of the atomic resonance caused by the trap beam,⁴ or interference with the trapping process due to the large fraction of excited atoms caused by the OM beams. If the chopping cycle is short compared with one-half of the oscillation period of an atom in the potential well ($\sim 10 \mu\text{s}$),




FIG. 1. (a) The normalized, time-averaged axial potential $\bar{U}' = 2\bar{U}/h\gamma$ as a function of the axial coordinate $z' = z\lambda/\pi w_0^2$ for $w_0 = 10 \mu\text{m}$, $P = 220 \text{ mW}$, and for three values of trap-laser detuning below the D_2 resonance line. The calculation uses $I_s(D_1) = 2I_s(D_2) = 40 \text{ mW/cm}^2$. (b) The normalized, time-averaged, potential well depth versus $\Delta\nu$ for $w_0 = 10 \mu\text{m}$ and $P = 220 \text{ mW}$.

then an atom in the trap does not move much during the chopping cycle and its motion is roughly that of a damped harmonic oscillator with an average well depth of one-half the cw well depth. In our trap, the heating rate due to the trapping forces is so small that the temperature of the atoms in the trap remains close to the cw OM temperature.

6 Our experimental apparatus is similar to the apparatus used earlier to demonstrate OM^{18,19} with the addition of a focused trapping beam nearly parallel to one of the OM beams. Pulses of sodium atoms were created every 0.1–10 s by evaporating sodium metal with a 10-ns, pulsed, yttrium-aluminum-garnet laser beam. Atoms with a velocity of about $2 \times 10^4 \text{ cm/s}$ were slowed over a distance of 8 cm to about $2 \times 10^3 \text{ cm/s}$ by the radiation pressure of a frequency-chirped laser beam counterpropagating with the atomic beam. An electro-optic modulator at 856.2 MHz is used to create sidebands 1712.4 MHz apart to prevent optical pumping.¹⁸ The slowed atoms drift into the OM region where three-dimensional cooling and viscous confinement occurs. The OM-beam spot size is about 1 cm and the peak intensity is about 20 mW/cm^2 in the relevant sidebands. The optimum tuning of the OM beams ($\sim \gamma/2$ below resonance) is determined empirically by tuning for maximum storage time. We routinely achieve OM lifetimes of 0.5 s and densities of 10^6 atoms/cm^3 . Temperatures as low as $0.24 \times 10^{-3} \text{ K}$ were previously measured.¹⁸ The trap laser beam, obtained from a second dye laser, enters the vacuum chamber nearly parallel to one of the OM beams. The linearly polarized trapping beam has a power of about 220 mW and is focused to a diffraction-limited spot radius $w_0 = 10 \mu\text{m}$.

7 The optically trapped atoms were detected visually, by video camera, and photographically. Figure 2(a) is a 0.5-s exposure that shows the fluorescence from atoms in the initial slowing beam, the subsequently formed OM cloud, and atoms collecting in the trap. Figure 2(b) taken later shows the bright spot from trapped atoms which remain in the trap after most of the surrounding OM atoms have diffused away. The brightness of the fluorescence from the trapped atoms indicates a density much higher than the surrounding cloud of cold atoms. Virtually all of the fluorescence is caused by the OM beams since our trap is tuned far from the atomic resonance.

8 Several tests were made to confirm trapping: (i) The bright oval-shaped spot coming from trapped atoms occurs only for tunings within the range expected for axial trapping. For example, strong trapping was observed between -570 to at least -1300 GHz below the D_2 resonance with the deepest traps occurring at -650 ± 25 GHz. This agrees well with calculations based on Fig. 1(b). (ii) We observed visually that the lifetime of trapped atoms was longer than for




FIG. 2. (a) Photo showing the collimating nozzle, atomic beam, and atoms confined in OM. The distance from the nozzle to the OM region is 5 cm. (b) Photo taken after the atomic source and the slowing laser beam have been turned off, showing trapped atoms.

confined OM atoms. To quantify these observations we analyzed video tapes of the trap decay using a video waveform monitor. Figure 3(a) shows the signal of one horizontal video line through the trap center. The curve S_1 in Fig. 3(b) shows that atoms confined within only OM decay with a lifetime of about 0.5 s. Curve S_2 is for trapped atoms plus atoms in OM. The fact that S_2 is nonzero for more than a second after S_1 has fully decayed indicates that trapped atoms are longer lived. We surmise that background pressure is limiting the trap decay lifetimes to ~ 1 s. Since background gas atoms are at ~ 300 K and trapped atoms will be ejected by collisions transferring $< 10^{-2}$ K, we expect a large cross section for ejection of $\sim 10^{-13}$ cm². (iii) The shape of the trap fluorescence varied with tuning and trap-laser intensity as expected from the calculated axial potential profiles shown in Fig. 1(a). As an example, for strong trapping at -650 GHz we see a bright oval-shaped spot ~ 0.5 mm long, but not fully resolved. As the power is lowered, the oval lengthens and finally a long weak fluorescent streak appears only on the $+z$ side of the trap due to atoms confined transversely by the stronger radial potential well. (iv) Measurements were made on the ef-

fect of our varying the chopping period for the trapping and OM beams. At 220-mW and -650 -GHz tuning, good trapping was obtained for chopping periods between 0.4 and 10.0 μ s. Computer simulations of the radial motion by use of average optical forces shows that the atomic motion becomes unstable for periods longer than 8 μ s. We believe that trapping fails at fast chopping rates as a result of the generation of frequency sidebands, which is detrimental to optical cooling.

9 A measurement of the trap heating time was made by turning off the OM while keeping the trap on for varying times (0, 5, 10, 15 ms). At times of 5 ms the brightness of the trapped atoms decreases by about a factor of 2. This agrees with our calculated trap heating time⁴ of 4 ms.

10 We can determine the temperature of trapped atoms from the observation that at -750 GHz strong traps were obtained for 220 and 110 mW while no trapping was seen for 90 mW. This is done by calculating the intrinsic leakage rate of the trap as a function of power and temperature. The leakage rate is given by the escape-attempt frequency (twice the oscillation frequency of an atom in the trap) times the escape probability, $\exp(-\Delta U/kT)$. To agree with observation, we require intrinsic trap lifetimes of 1 s or greater for 110 mW and much less than 1 s for 90 mW. The result is, $0.6 \leq kT/h\gamma \leq 0.8$; thus the trapped-atom temperature is remarkably close to the quantum limit for atoms in OM,¹⁸ $kT = \frac{1}{2}h\gamma = 240$ μ K. The intrinsic trap lifetimes corresponding to $kT = 0.7h\gamma$ are 0.12 s for 90 mW, 1.0 s for 110 mW, and 3.7×10^4 s for 220 mW. The inferred temperature for the trapped atoms allows us to calculate the fraction of the trapping volume in which the atoms reside. For a detuning of -650 GHz, a trap beam power of 220 mW, and an atomic temperature of $0.7h\gamma$, the atoms are confined within a cylinder of length 210 μ m and diameter of 2.5 μ m, giving a volume of 1×10^{-9} cm³.

11 We deduced the number of atoms in the trap and




FIG. 3. (a) Video signal on a single horizontal video line through the trap showing scattered light, camera noise, atoms in OM, and a sharp spike due to the fluorescence of the trapped atoms. Data were quantified by use of boxcar integrators to measure the time dependence of the levels S_2 and S_1 . (b) Time dependence of S_2 (trap plus OM fluorescence) and S_1 (OM fluorescence).

their density in several independent ways. From an absolute calibration of the detection system (zoom lens, video camera, and waveform monitor combination), we conclude that our best trap signals are due to about 500 atoms. Based on the confinement volume of $1 \times 10^{-9} \text{ cm}^3$, the corresponding trapped-atom density is $5 \times 10^{11} \text{ cm}^{-3}$. An accurate estimate of the trapped-atom density is also obtained by noting that the signal from the trapped atoms is roughly equal to the signal from the atoms in OM. Using the f-stop number of the video-camera lens and assuming a OM radius of 0.5 cm, we estimate that the average trap density is 1.5×10^6 times the OM density. A OM density of $5 \times 10^5 \text{ cm}^{-3}$ yields a trap density of $8 \times 10^{11} \text{ cm}^{-3}$, reasonably consistent with the previous estimate.

12 Finally we note that the collection of trapped atoms could be moved easily at speeds on the order of 1 cm/s by manually scanning the location of the trap focal spot.

13 In summary, optical trapping of neutral atoms has been observed in good quantitative agreement with theoretical expectations. We have achieved orders-of-magnitude higher densities of trapped atoms and lower temperatures than reported previously using magnetic neutral-atom traps²² and electrodynamic ion traps.^{23,24} It may be possible to increase the density by $\geq 10^4$ and to decrease the temperature to $\leq 10^{-6} \text{ K}$. This would involve tighter beam focusing, additional cooling of OM with narrower linewidth transitions, or possibly beam-expansion schemes.²⁵ Confinement and spatial manipulation of cold trapped atoms to dimensions less than the optical wavelength are also possible. These capabilities should prove useful in studies of Bose and Fermi gas statistics and other interactions at high atomic density, atom-atom collisions, atom-surface interactions, atom diffraction and tunneling, molecular formation using individual atoms, modifications of atomic spontaneous-emission lifetimes and other collective effects, spectroscopy of the trap energy levels, and, possibly, high-resolution spectroscopy and atomic time standards.

The authors thank J. P. Gordon for valuable discussions.

¹A. Ashkin, Phys. Rev. Lett. **24**, 156 (1970).

²A. Ashkin, Phys. Rev. Lett. **25**, 1321 (1970).

³A. Ashkin, Phys. Rev. Lett. **40**, 729 (1978).

⁴J. P. Gordon and A. Ashkin, Phys. Rev. A **21**, 1606 (1980).

⁵A. Ashkin, Science **210**, 1081 (1980).

⁶V. Letokhov and V. Minogin, Physics Rep. **73**, 3 (1981).

⁷*Laser-Cooled and Trapped Atoms*, edited by W. D. Phillips, Progress in Quantum Electronics Vol. 8 (Pergamon, New York, 1984).

⁸A. Ashkin, Opt. Lett. **9**, 454 (1984).

⁹G. Roosen, Can. J. Phys. **57**, 1260 (1979).

¹⁰A. Ashkin and J. M. Dziedzic, Phys. Rev. Lett. **54**, 1245 (1985).

¹¹A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm, and Steven Chu, Opt. Lett. **11**, 288 (1986).

¹²R. Schieder, H. Walther, and L. Wöste, Opt. Commun. **5**, 337 (1972).

¹³P. Jacquinet, S. Liberman, J. L. Picque, and J. Pinard, Opt. Commun. **8**, 163 (1973).

¹⁴J. E. Bjorkholm, R. R. Freeman, and D. B. Pearson, Phys. Rev. A **23**, 491 (1981).

¹⁵J. E. Bjorkholm, R. R. Freeman, A. Ashkin, and D. B. Pearson, Phys. Rev. Lett. **41**, 1361 (1978).

¹⁶J. Prodan, A. Migdall, W. D. Phillips, I. So, and H. Dalibard, Phys. Rev. Lett. **54**, 992 (1985).

¹⁷E. Ertmer, R. Blatt, J. L. Hall, and M. Zhu, Phys. Rev. Lett. **54**, 996 (1985).

¹⁸Steven Chu, L. Hollberg, J. E. Bjorkholm, Alex Cable and A. Ashkin, Phys. Rev. Lett. **55**, 48 (1985).

¹⁹Steven Chu, J. E. Bjorkholm, A. Ashkin, L. Hollberg, and Alex Cable, "Methods of Laser Spectroscopy" (Plenum, New York, to be published).

²⁰We have recently learned of a newly proposed scattering-force optical trap that can have volumes and well depths comparable with alternating-beam scattering-force traps: D. E. Pritchard, private communication, and D. E. Pritchard *et al.*, preceeding Letter [Phys. Rev. Lett. **57**, 310 (1986)].

²¹J. Dalibard, S. Reynaud, and C. Cohen-Tannoudji, Opt. Commun. **47**, 395 (1983).

²²A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, Phys. Rev. Lett. **54**, 2596 (1985).

²³D. J. Wineland, W. M. Itano, J. C. Bergquist, J. J. Bollinger, and H. Hemmelt, Prog. Quantum Electron. **8**, 139 (1984).

²⁴W. Nagourney, G. Janik, and H. Dehmelt, Proc. Natl. Acad. Sci. U.S.A. **80**, 643 (1983).

²⁵S. Chu, J. E. Bjorkholm, A. Ashkin, J. P. Gordon, and L. W. Hollberg, Opt. Lett. **11**, 73 (1986).




FIG. 2. (a) Photo showing the collimating nozzle, atomic beam, and atoms confined in OM. The distance from the nozzle to the OM region is 5 cm. (b) Photo taken after the atomic source and the slowing laser beam have been turned off, showing trapped atoms.