Preface

Writing for Computer Science is an introduction to doing and describing research. For the most part the book is a discussion of good writing style and effective research strategies. Some of the material is accepted wisdom, some is controversial, and some is my opinions. Although the book is brief, it is designed to be comprehensive: some readers may be interested in exploring topics further, but for most readers this book should be sufficient.

The first edition of this book was almost entirely about writing. This edition, partly in response to reader feedback and partly in response to issues that arose in my own experiences as an advisor, researcher, and referee, is also about research methods. Indeed, the two topics—writing about and doing research—are not clearly separated. It is a small step from asking *how do I write?* to asking *what is it that I write about?*

As previously, the guidance on writing focuses on research, but much of the material is applicable to general technical and professional communication. Likewise, the guidance on the practice of research has broader lessons. A practitioner trying a new algorithm or explaining to colleagues why one solution is preferable to another should be confident that the arguments are built on robust foundations. And, while this edition has a stronger emphasis on research than did the first, nothing has been deleted; there is additional material on research, but the guidance on writing has not been taken away.

Since the first edition appeared, there have been many changes in the culture and environment of research. The web has become universal, whereas, for example, few papers were online. There are also more subtle changes. It now seems to be rare that a spoken presentation is truly unprofessional; a decade ago many talks were unendurably awful. The growth in the use of good tools for presentations has been a key factor in this development, and the use of overhead transparencies has become archaic.

On the other hand, it now seems common that a talk does not have a clear message and is merely a compilation of clever visuals. Writing style has become less stilted, which is a change for the better, but too many authors are submitting work too early. Today, algorithms are often poorly described; a well-described algorithm has become a welcome, rare exception. The web provides easy access to literature, but perhaps the necessity of using a library imposed discipline, as, increasingly, past work appears to be neglected.

The perspectives of all scientists are shaped by the research cultures in which they work. My research has involved some theoretical studies, but the bulk of my work has been experimental. I appreciate theoretical work for its elegance, yet find it sterile when it is too detached from practical value. While experimental work can be ad hoc, it can also be deeply satisfying, with the rewards of probing the space of possible algorithms and producing technology that can be applied to the things we do in practice. My perspective on research comes from this background, as does the use of experimental work as examples in this book (an approach that is also justified by the fact that such work is generally easier to outline than is a theoretical contribution). But that doesn't mean that my opinions are simply private biases. They are—I hope!—the considered views of a scientist with experience of different kinds of research.

Many people helped with this book in one way or another. For the first edition, thanks were due in particular to Alistair Moffat, who contributed to Chapters 6, 7, 8, 9, and 12; and to Philip Dart, who also contributed to Chapter 12. I remain grateful to both Alistair and Philip for our collaborations. Additionally, I thanked Isaac Balbin, Gill Dobbie, Evan Harris, Michael Fuller, Mary and Werner Pelz, Kotagiri Ramamohanarao, Ron Sacks-Davis, Ian Shelley, James Thom, Rodney Topor, Ross Wilkinson, and Hugh Williams. I also thanked my research students and the students who participated in my research methods lectures. To all these people—thanks again.

For this edition, I thank Timothy A.H. Bell, Bodo Billerbeck, Beverley Ford, Michael Fuller (again), Paul Gruba, Lin Padgham, Jenny Wolkowicki, and the many readers who pointed out mistakes or made helpful suggestions. My children showed remarkable patience and I thank them for their forbearance. The person I thank the most is my wife, Penny Tolhurst.

Justin Zobel
Melbourne, Australia
February 2004

9 Writing up

I used to think about my sentences before writing them down; but . . . I have found that it saves time to scribble in a vile hand whole pages as quickly as I possibly can . . . . Sentences thus scribbled down are often better ones than I could have written deliberately.

Charles Darwin
Autobiography

Science is more than a body of knowledge; it is a way of thinking.

Carl Sagan
The Demon-Haunted World

In every research project, a stage is reached at which it makes sense to begin to write up. A good principle is to begin early: if it is possible to start writing then the writing should start, typically well before the project's half-way mark. Shaping the research and its outcomes into a write-up is an effective way of giving structure to a project, even if the outcomes are not yet clear or months are needed to complete system development.

The task of writing up research is the topic of this chapter: gathering material, organizing it so that the work tells a story, giving this story the structure of a thesis or of an academic paper, and starting to write. The research that precedes the write-up is the topic of Chapters 10 and 11.

The scope of a paper

To begin a paper, the first task is to identify your aims. Write down everything that motivated you to start the research. What did you want to achieve? What problems did you expect to address? What makes the problems interesting? Next, define the scope of the work that you plan to write up. To do so, it is necessary to make choices about what to include, and thus it is necessary to identify what *might* be included. Typically, by this stage your research has become focused on investigation of a small number of specific questions, and you have preliminary experimental or theoretical results that suggest what the core contribution of the work is going to be.

You might start, for example, by asking questions such as:

- Which results are the most surprising?
- What is the one result that other researchers might adopt in their work?
- Are the other outcomes independent enough to be published separately later on? Are they interesting enough to justify their being included?
- Does it make sense to explain the new algorithms first, followed by description of the previous algorithms in terms of how they differ from the new work? Or is the contribution of the new work more obvious if the old approaches are described first, to set the context?
- What assumptions or definitions need to be formalized before the main theorem can be presented?
- What is the key background work that has to be discussed?
- Who is the readership? For example, are you writing for specialists in your area, your examiners, or a general computer science audience?

Other questions are given in the checklist on page 155.

A valuable exercise at this stage is to speculate on the format and scope of the results. Early in the investigation, decisions will have been made about how the results are to be evaluated—that is, about which measures are to be used to determine whether the research has succeeded or failed. For example, it may be that network congestion is the main respect in which the research is expected to have yielded improvements in performance. But how is network congestion to be measured? As a function of data volume, number of machines, network bandwidth, or something else? Answering this question suggests a form of presentation into which the experimental results can be inserted: a graph, perhaps. The form of this graph can be sketched even before any coding has begun, and doing so identifies the kind of output that the code is required to produce.

Consider a detailed example: an investigation of external sorting in database systems. In this task, a large relation—tens of millions of records, say, constituting several gigabytes—must be sorted on a field specified in a query. An effective sorting method is to sort the relation one block at a time, storing the sorted blocks in a temporary file then merging them to give the final result. Costs include processing time for sorting and merging, transfer time to and from disk, and temporary space requirements. The balance between these costs is governed by available in-memory buffer space, as large blocks are expensive to sort but cheap to merge. The specific research question being investigated is whether disk costs can be reduced by compressing the data while it is sorted.

Speculation about how compression might affect costs suggest how the work should be measured. For small relations, compression seems unlikely to be of help—compressing and then decompressing adds processing costs but does not provide savings if all the data fits in memory. For large relations, on the other hand, the savings due to reduced disk traffic, increased numbers of records per block, and use of less temporary space may be significant. Thus it seems likely that the savings due to compression would increase with the size of relation to be sorted, suggesting use of a graph of data volume against sorting time for fixed block size. Note too that the question of what to measure identifies an implicit assumption: that the data was uncompressed to begin with and is returned uncompressed. All of these decisions and steps help to determine the paper's content.

The content of a paper is to a significant extent determined by the readership. You may be reporting a particular piece of work, but the way it is reported is determined by the characteristics of the audience. For example, a paper on machine learning for computer vision may have entirely different implications for the two fields, and thus different aspects of the results might be emphasized. Also, an expert on vision cannot be assumed to have any experience with machine learning, so the way in which the material is explained to the two readerships must be based on your judgement, in each case, of what is common knowledge and what is unfamiliar. The nature of the audience may even determine the scope of what can be reported.

Making choices about the content of a paper places limits on its scope; these choices identify material to be excluded. Broadly speaking, many research programs are a cycle of innovation and evaluation, with the answers or resolution of one investigation creating the questions that lead to the next. An advance in, say, string sorting might well have implications for integer sorting, and further work could pursue these implications. But at some point it is necessary to stop undertaking new work and write up what has been achieved so far. The new ideas may well be exciting—and less stale than the work that has been preoccupying you for months—but they are likely to be less well understood, and completing the old work is more important than trying to include too many results. If the newer work can be published independently, then write it up separately. A long, complex paper, however big a breakthrough it represents, is hard to referee. From an editor's perspective, accepting such a paper may be difficult to justify if it squeezes out several other contributions.

Another element in the process of developing a paper is deciding where the work might be published. There are many factors that should be considered when making this decision, such as relevance to your topic and how your work measures against the standard for that forum. In particular, the venue partly determines the scope of a paper. For example, is there a page limit? Are there specific conventions to be observed? Are the other papers in that venue primarily theoretical or experimental? What prior knowledge or background is a reader likely to have? Do the editors require that your code be available online? If you select a particular forum but haven't cited any papers that have appeared there, you may have made the wrong choice.

Once the material for a paper has been collected it has to be organized into a coherent self-contained narrative, which ultimately will form the body of the write-up. Turning this narrative into a write-up involves putting it in the form of an academic paper: including an introduction, a bibliography, and so on. These issues are discussed later.

Telling a story

A cornerstone of good writing is identifying what the reader needs to learn. A paper is a sequence of concepts, building from a foundation of knowledge assumed to be common to all readers up to new ideas and results. Thus an effective paper educates its readers. It leads readers from what they already know to new knowledge you want them to learn. For this reason, the body of a good paper—everything between the introduction and the conclusions—should have a logical flow that has the feel of a narrative.

The narrative told by a paper is a walk through the ideas and outcomes. It isn't a commentary on the research program or the day-to-day activities of the participants, nor is it meant to be mysterious. Instead, it is like a guided tour through a gallery, in which each room contains something new for the readers to comprehend. There is also an expectation of logical closure. The early parts of the paper's body typically explain hypotheses or claims; the reader expects to discover by the end whether these are justified.

There are several common ways for structuring the body of a paper, including as a chain, by specificity, by example, and by complexity. Perhaps the most common structure is the first of these alternatives, a *chain* in which the results and the background on which they build dictate a logical order for presentation of the material. First might come, say, a problem statement, then a review of previous solutions and their drawbacks, then the new solution, and finally a demonstration that the solution improves on its predecessors.

The "compression for fast external sorting" project suggests a structure of this kind. The problem statement consists of an explanation of external sorting and an argument that disk access costs are a crucial bottleneck. The review explains standard compression methods and why they cannot be integrated into external sorting. The new solution is the compression method developed in the research. The demonstration is a series of graphs and tables based on experiments that compare the costs of sorting with and without compression.

For some kinds of results, other structures may be preferable. One option is to structure by *specificity*, an approach that is particularly appropriate for results that can be divided into several stages. The material is first outlined in general terms, then the details are progressively filled in. Most technical papers have this organization at the high level, but it can also be used within sections.

Material that might have such a structure is an explanation of a retrieval system. Such systems generally have several components. For example, in text retrieval a parser is required to extract words from the text that is being indexed; this information must be passed to a procedure for building an index; queries must likewise be parsed into a format that is consistent with that of the stored text; and a query evaluator uses the index to identify the records that match a given query. The explanation might begin with a review of this overall structure, then proceed to the detail of the elements.

Another structure is by *example*, in which the idea or result is initially explained by, say, applying it to some typical problem. Then the idea can be explained more formally, in a framework the example has made concrete and familiar. The "compression for fast external sorting" could also be approached in this way. The explanation could begin by considering, hypothetically, the likely impact of compression on sorting. To make the discussion more concrete, a couple of specific instances—a small relation and a large relation, say—could be used to illustrate the expected behaviour in different circumstances. Given a clear explanation of the hypothetical scenario, you can then proceed to fill in details of the method that was tested in the research.

Another alternative is to structure the body by *complexity*. For example, a simple case can be given first, then a more complex case can be explained as an extension, thus avoiding the difficulty of explaining basic concepts in a complex framework. This approach is a kind of tutorial: the reader is brought by small steps to the full result. For example, a mathematical result for an object-oriented programming language might initially be applied to some simple case, such as programs in which all objects are of the same class. Then the result could be extended by considering programs with inheritance.[12]

Some other structures are inappropriate for a write-up. For example, the paper should not be a chronological list of experiments and results. The aim is to present the evidence needed to explain an argument, not to list the work undertaken.

Most experiments yield far more data than can be presented in a paper of reasonable length. Important results can be summarized in a graph or a table, and other outcomes reported in a line or two. It is acceptable to state that experiments have yielded a certain outcome without providing details, so long as those experiments do not affect the main conclusions of the paper (and have actually been performed). Similarly, there may be no need to include the details of proofs of lemmas or minor theorems. This does not excuse you from conducting the experiments or convincing yourself that the results are correct, but such information can be kept in logs of the research rather than included in the paper.

The traditional structure for organizing research papers can encourage you to list all proofs or results, then analyze them later; with this structure, however, the narrative flow is often poor. It usually makes more sense to analyze proofs or experimental results as they are presented, particularly since experiments or theorems often follow a logical sequence in which the outcome of one dictates the parameters of the next.

When describing specific results, it is helpful, although not always possible, to begin with a brief overview of whatever has been observed. The rest of the discussion can then be used for amplification rather than further observations. Newspaper articles are often written in this way. The first sentence summarizes the story; the next few sentences review the story again, giving some context; then the remainder of the article presents the whole story in detail. Sections of research papers can sometimes be organized in this way.

[12]Structuring by complexity is good for a paper but, often, inappropriate for ongoing research. It is not uncommon to see a paper in which the authors have solved an easy case of a problem, say optimizations for iteration-free programs, motivated by hopeful claims such as "we expect these results to throw light on optimization of programs with loops and recursion". All too often the follow-up paper never appears.

Organization

Scientific papers follow a standard structure that allows readers to quickly discover the main results, and then, if interested, to examine the supporting evidence. Many readers accept or reject conclusions based on a quick scan, not having time to read all the papers they see. A well-structured write-up accommodates this behaviour by having important statements as near the beginning as possible. You need to:

- Describe the work in the context of accepted scientific knowledge.
- State the idea that is being investigated, often as a theory or hypothesis.
- Explain what is new about the idea, what is being evaluated, or what contribution the paper is making.
- Justify the theory, by methods such as proof or experiment.

Theses, journal articles, and conference papers have much the same organization when viewed in outline. There are distinctions in emphasis rather than specific detail. For a thesis, for example, the literature review may be expected to include a historical discussion outlining the development of the key ideas. There is also an expectation that a thesis is a completed, rounded piece of work—a consolidation of the achievements of a research program as well as a report on specific scientific results. Nonetheless, these forms of write-up have similar structure.

A typical write-up has most of the following components:

Title and author

Papers begin with their title and information about authors including name, affiliation, and address. The convention in computer science is to not give your position, title, or qualifications; but whether you give your name as A. B. Cee, Ae Cee, Ae B. Cee, or whatever, is a personal decision. Use the same style for your name on all your papers, so that they are indexed together. Include a durable email address or web address.

Also include a date. Take the trouble to type in the date rather than using "today" facilities that print the date on which the document was last processed, or later you may not be able to tell when the document was completed.

The front matter of a paper may also include other elements. One is acknowledgements, as discussed on page 26, which may alternatively follow the conclusions. Another element is a collection of search terms, keywords, or key phrases—additional terminology that can be used to describe the topic of the paper. Sometimes these keywords must be selected from a specific list. In other cases, the conventions for choosing such terms are not always clear, but in general it is unhelpful to use words that, for example, are a description of the experimental methodology: don't write "timing experiments", for example. Use words that concern the paper's principal themes.

Abstract

An abstract is typically a single paragraph of about 50 to 200 words. The function of an abstract is to allow readers to judge whether or not the paper is of relevance to them. It should therefore be a concise summary of the paper's aims, scope, and conclusions. There is no space for unnecessary text; an abstract should be kept to as few words as possible while remaining clear and informative. Irrelevancies, such as minor details or a description of the structure of the paper, are inappropriate, as are acronyms, abbreviations, and mathematics. Sentences such as "We review relevant literature" should be omitted.

The more specific an abstract is, the more interesting it is likely to be. Instead of writing "space requirements can be significantly reduced", write "space requirements can be reduced by 60%". Instead of writing "we have a new inversion algorithm", write "we have a new inversion algorithm, based on move-to-front lists".

Many scientists browse research papers outside their area of expertise. You should not assume that all likely readers will be specialists in the topic of their paper—abstracts should be self-contained and written for as broad a readership as possible. Only in rare circumstances should an abstract cite another paper (for example, when one paper consists entirely of analysis of results in another), in which case the reference should be given in full, not as a citation to the bibliography.

Introduction

An introduction can be regarded as an expanded version of the abstract. It should describe the paper's topic, the problem being studied, references to key papers, the approach to the solution, the scope and limitations of the solution, and the outcomes. There needs to be enough detail to allow readers to decide whether or not they need to read further. It should include motivation: the introduction should explain why the problem is interesting, what the relevant scientific issues are, and why the solution is a good one.

That is, the introduction should show that the paper is worth reading and it should allow the reader to understand your perspective, so that the reader and you can proceed on a basis of common understanding.

Many introductions follow a five-element organization:

1. A general statement introducing the broad research area of the particular topic being investigated.
2. An explanation of the specific problem (difficulty, obstacle, challenge) to be solved.
3. A brief review of existing or standard solutions to this problem and their limitations.
4. An outline of the proposed new solution.
5. A summary of how the solution was evaluated and what the outcomes of the evaluation were.

An interesting exercise is to read other papers, analyze their introductions to see if they have this form, and then decide whether they are effective.

The introduction can discuss the importance or ramifications of the conclusions but should omit supporting evidence, which the interested reader can find in the body of the paper. Relevant literature can be cited in the introduction, but unnecessary jargon, complex mathematics, and in-depth discussion of the literature belong elsewhere.

A paper isn't a story in which results are kept secret until a surprise ending. The introduction should clearly tell the reader what in the paper is new and what the outcomes are. There may still be a little suspense: revealing what the results are does not necessarily reveal how they were achieved. If, however, the existence of results is concealed until later on, the reader might assume there are no results and discard the paper as worthless.

Body

The body of a paper should present the results. The presentation should provide necessary background and terminology, explain the chain of reasoning that leads to the conclusions, provide the details of central proofs, summarize any experimental outcomes, and state in detail the conclusions outlined in the introduction. Descriptions of experiments should permit reproduction and verification, as discussed in Chapter 11. There should also be careful definitions of the hypothesis and major concepts, even those described informally in the introduction. The structure should be evident in the section headings. Since the body can be long, narrative flow and a clear logical structure are essential.

The body should be reasonably independent of other papers. If, to understand your paper, the reader must find specialized literature such as your earlier papers or an obscure paper by your advisor, then its audience will be limited.

In some disciplines, research papers have highly standardized structures. Editors may require, for example, that you use only the four headings Introduction-Methods-Results-Discussion. This convention has not taken hold in computer science, and in some cases such a structure impedes a clear explanation of the work. For example, use of fixed headings may prohibit development of a complex explanation in stages. In work combining two query resolution techniques, we had to determine how they would interact, based on a fresh evaluation of how they behaved independently. The final structure was, in effect, Introduction-Background-Methods-Results-Discussion-Methods-Results-Discussion.

Even if the standardized section names are not used, the body needs these elements, if not necessarily under their standard headings. Components of the body might include, among other things, background, previous work, proposals, experimental design, analysis, results, and discussion. Specific research projects suggest specific headings. For the "compression for fast external sorting" project sketched earlier, the complete set of section headings might be:

1. Introduction
2. External sorting
3. Compression techniques for database systems
4. Sorting with compression
5. Experimental setup
6. Results and discussion
7. Conclusions

The wording of these headings does not follow the standard form, but the intent of the wording is the same. Sections 2 and 3 are the background; Section 4 contains novel algorithms, and Sections 4 and 5 together are the methods.

The background material can be entirely separate from the discussion of previous work on the same problem. The former is the knowledge the reader needs to understand your contribution. The latter is, often, alternative solutions that are superseded by your work. Together, the discussion of background and previous work also introduce the state of the art and its failings, the importance and circumstances of the research question, and benchmarks or baselines that the new work should be compared to.

A body that consists of descriptions of algorithms followed by a dump of experimental results is not sound science. In such a paper, the context of prior work is not explained, as readers are left to draw their own inferences about what the results mean.

In a thesis, each chapter has structure, including an introduction and a summary or conclusions. This structure varies with the chapter's purpose. A background chapter may gather a variety of topics necessary to understanding of the contribution of the thesis, for example, whereas a chapter on a new algorithm may have a simple linear organization in which the parts of the algorithm are presented in turn. However, the introduction and summary should help to link the thesis together—how the chapter builds on previous chapters and how subsequent chapters make use of it.

Literature review

Few results or experiments are entirely new. Most often they are extensions of or corrections to previous research—that is, most results are an incremental addition to existing knowledge. A literature review, or survey, is used to compare the new results to similar previously published results, to describe existing knowledge, and to explain how it is extended by the new results. A survey can also help a reader who is not expert in the field to understand the paper and may point to standard references such as texts or survey articles.

In an ideal paper, the literature review is as interesting and thorough as the description of the paper's contribution. There is great value for the reader in a precise analysis of previous work that explains, for example, how existing methods differ from one another and what their respective strengths and weaknesses are. Such a review also creates a specific expectation of what the contribution of the paper should be—it shapes what the readers expect of your work, and thus shapes how they will respond to your ideas.

The literature review can be early in a paper, to describe the context of the work, and might in that case be part of the introduction; or the literature review can follow or be part of the main body, at which point a detailed comparison between the old and the new can be made. If the literature review is late in a paper, it is easier to present the surveyed results in a consistent terminology, even when the cited papers have differing nomenclature and notation.

In many papers the literature review material is not gathered into a single section, but is discussed where it is used—background material in the introduction, analysis of other researchers' work as new results are introduced, and so on. This approach can help you to write the paper as a flowing narrative.

An issue that is difficult in some research is the relationship between new scientific results and proprietary commercial technology. It often is the case that scientists investigate problems that appear to be solved or addressed in commercial products. For example, there is ongoing academic research into methods for information retrieval despite the success of the search engines deployed on the web. From the perspective of high research principle, the existence of a commercial product is irrelevant: the ideas are not in the public domain, it is not known how the problems were solved in the product, and the researcher's contribution is valid. However, it may well be reckless to ignore the product; it should be cited and discussed, while noting, for example, that the methods and effectiveness of the commercial solution are unknown.

Conclusions

The closing section, or summary, is used to draw together the topics discussed in the paper. It should include a concise statement of the paper's important results and an explanation of their significance. This is an appropriate place to state (or restate) any limitations of the work: shortcomings in the experiments, problems that the theory does not address, and so on.

The conclusions are an appropriate place for a scientist to look beyond the current context to other problems that were not addressed, to questions that were not answered, to variations that could be explored. They may include speculation, such as discussion of possible consequences of the results.

A *conclusion* is that which concludes, or the end. *Conclusions* are the inferences drawn from a collection of information. Write "Conclusions", not "Conclusion". If you have no conclusions to draw, write "Summary".

Bibliography

A paper's bibliography, or its set of references, is a complete list of theses, papers, books, and reports cited in the text. No other items should be included. Citation and bibliographies are discussed in detail starting on page 19.

Appendices

Some papers have appendices giving detail of proofs or experimental results, and, where appropriate, material such as listings of computer programs. The purpose of an appendix is to hold bulky material that would otherwise interfere with the narrative flow of the paper, or material that even interested readers do not need to refer to. Appendices are rarely necessary.

The first draft

For the first draft of a write-up you may find it helpful to write freely—without particular regard to style, layout, or even punctuation—so that you can concentrate on presenting a smooth flow of ideas in a logical structure. Worrying about how to phrase each sentence tends to result in text that is clear but doesn't form a continuous whole, and authors who are too critical on the first draft are often unable to write anything at all. If you tend to get stuck, just write anything, no matter how awful; but be sure to delete any ravings later.

Some people, when told to just say anything, find they can write freely—if anything is acceptable, then nothing is wrong. For others, finding words is still a struggle. A last resort is to write in brief sentences making the simplest possible statements.

✓ In-memory sorting algorithms require random access to records. For large files stored on disk, random access is impractically slow. These files must be sorted in blocks. Each block is loaded into memory and sorted in turn. Sorted blocks are written to temporary files. These temporary files are then merged. There may be many files but in practice the merge can be completed in one pass. Thus each record is read twice and written twice. Temporary space is required for a complete copy of the original file.

This text certainly isn't elegant—it is annoying to read and should be thoroughly edited long before the paper is submitted. But it is capturing the ideas, and the writing is proceeding.

A consequence of having a sloppy first draft is that you must edit and revise carefully; initial drafts are often turgid and full of mistakes. But few authors write well on the first draft anyway; the best writing is the result of frequent, thorough revision.

Mathematical content, definitions, and the problem statement should be made precise as early in the writing process as possible. The hypothesis and the results flow from a clear statement of the problem being tackled. Describing the problem forces you to consider in depth the scope and nature of the research. If you find that you cannot describe the problem precisely, then perhaps your understanding is lacking or the ideas are insufficiently developed.

It was said earlier, but is worth repeating: the writing should begin long before the research is finished, and perhaps as soon as it is started. The later the writing is begun, the harder it will be. Delay increases the time between having ideas and having to write about them, increases the number of papers to discuss, and increases the number of experiments to describe. Completing your reading, for example, is a poor reason to defer writing, because reading is never complete. Writing is a stimulus to research, suggesting fresh ideas and clarifying vague concepts and misunderstandings; and developing the presentation of the results oftens suggest the form the proofs or experiments should take. Gaps in the research may not be apparent until it has been at least preliminarily described. Research is also a stimulus to writing—fine points are quickly forgotten once the work is complete. Don't expect the writing to progress steadily, but do expect progress overall. If the writing seems to have stalled, it is time to put other tasks aside for a while.

From draft to submission

There are many approaches to the process of assembling a technical paper. The technique I use for composing is to brainstorm, writing down in point form what has been learnt, what has been achieved, and what the results are. The next step is to prepare a skeleton, choosing results to emphasize and discarding material that on reflection seems irrelevant, and then work out a logical sequence of sections that leads the reader naturally to the results. A useful discipline is to choose the section titles before writing any text, because if material to be included doesn't seem to belong in any section then the paper's structure is probably faulty. The introduction is completed first and includes an overview of the paper's intended structure, that is, an outline of the order and content of the sections. When the structure is complete, each section can be sketched in perhaps 20 to 200 words. This approach has the advantage of making the writing task less daunting—it is broken into parts of manageable size.

When the body and the closing summary are complete, the introduction usually needs substantial revision because the arguments presented in the paper are likely to mature and evolve as the writing proceeds. The final version of the abstract is the last part to be written.

With a reasonably thorough draft completed, it is time to review the paper's content and contribution. Anticipate likely concerns or objections, and address them; if they can't be addressed, acknowledge them. Consider whether extra work is needed to fill a hole. Ask the probing, critical questions that you would ask of other people's work. The burden of proof is on you, not the reader, so be conservative in your claims and thorough with your evidence.

During drafting and revision, ensure that the topic of the paper does not drift. At the start of the writing process, you wrote down your aims, motivation, and scope. Use these as a reference. If you feel that you need to write something that is not obviously relevant to your original aims, then either establish the connection clearly or alter the aims. Changing the aims can affect the work in many ways, however, so only do so with great care.

For a novice writer who doesn't know where to begin, a good starting point is imitation. Choose a paper whose results are of a similar flavour to your own, analyze its organization, and sketch an organization for your results based on the same pattern. The habit of using similar patterns for papers—their standardization—helps to make them easier to read.

The practice of building a file of notes as you proceed is invaluable. Keep a dated log with records of the following:

- Meetings.
- Decisions.
- Ideas.
- Expectations of outcomes.
- Papers you have read.
- Sketches of algorithms.
- Code versions.
- Theorems.
- Experiments.
- Sketches of proofs.
- Outcomes.

Expect the log to be a mixture of a written notebook and data kept electronically. In its raw state, the content of a file of notes is not suitable for inclusion in a paper, but the themes and issues of the paper can be drawn from the file, and it serves as a memory of issues to discuss and material to include.

In computer science, most papers are co-authored. The inclusion of several people as authors means that, in principle, all these people contributed in some non-trivial way to the intellectual content of the paper. In many cases, it also means that the task of writing was shared. There are a range of strategies for co-authoring, which vary from colleague to colleague and paper to paper. It is not unusual, for example, for an advisor to use a student's thesis as the basis of a paper, in which case both advisor and student are listed as authors. In this process, the advisor may well dramatically revise the student's work, if only because a typical paper is much shorter than a typical thesis.

In cases where scientists are working more or less as equals, one strategy is to brainstorm the contents of the paper, then for each author to write a designated section. Another strategy—my preferred model for collaboration—is to take turns. One person writes a draft, the next revises and extends, and so on, with each person holding an exclusive lock on the paper while amending it. With this approach, the final paper is likely to be a fairly seamless integration of the styles and contributions of each of the authors (especially if each author contributes to revision of the other authors' work). In contrast, the strategy of writing sections separately tends to lead to papers in which the authorial voice makes dramatic shifts, the tables and figures are inconsistent, and there is a great deal of repetition and omission.

Taking turns is effective, but it does have pitfalls, and agreed ground rules are needed to make it work. For example, I rarely delete anything a co-author has written, but may comment it out; thus no-one feels that their work has been thrown away. Another element of successful co-authoring is respect. Accept your colleagues' views unless you have a good reason not to.

Co-authoring is a form of research training. It is an opportunity for advisors to learn in detail where their students are weak as scientists, while a paper that has been revised by an advisor is an opportunity for a student to contrast an attempt at research writing with that of an experienced scientist. An advisor's revision of a student's draft can involve a great deal of work, and may be the most thorough feedback on writing that the student receives during the course of a research program.

Prepublication

Traditionally, prior to a paper appearing in a refereed venue it might have been made available as a manuscript or technical report. These forms of publication once had the advantage of making the work available quickly—a particular concern if there is likely to be a substantial delay between submission and publication. (In some journals, the delay is years.) Departments prided themselves on the quality of their technical report series. However, this form of publication has withered away as the web has grown in importance; some academic institutions and large corporate research labs still publish significant numbers of reports, reflecting perhaps internal publication-approval processes, but these are the exception.

The web allows academics to readily publish their own work, independent even of the structures imposed by their departments. Most computer science researchers have websites on which they list their publications, and many researchers additionally list papers that are not yet published (and may never be published). An alternative is to place such papers in public archives.[13] Additionally, some research groups maintain topic-specific repositories.

Web publication has a range of advantages. Most importantly, it makes the work available immediately. While there is still an expectation that the work follows the conventions of a scientific paper, additional material can easily be included, such as links to data and source code. Many researchers access papers exclusively—both from academic publishers and from individuals—through the web, and there is growing acceptance amongst publishers that preliminary versions of papers are made publicly available by their authors.

The papers in most journals and conferences are available on the web via their publishers' websites. The fact that a paper is available through such a website tells the reader that the work has probably been refereed, that is, independently assessed by other scientists. (The quality of the refereeing varies from one conference or journal to another. When making an assessment of a paper, consider the reputation of the venue in which it appeared and issues such as those raised in Chapters 10 and 12.) While work that is published by an author on the web is immediately available, the lack of refereeing means that readers cannot be as confident of its validity.

Theses

A thesis (or, in some universities, a dissertation) is how research students present their work for examination. A thesis may have longer-term importance as a description of significant research results, but your primary goal should be to produce a piece of work that the examiners will pass.

The questions that examiners respond to are much the same as those a referee would ask of a paper. That is, the examiners seek evidence of an original, valid contribution developed to an appropriate standard. However, it is a mistake to view a thesis as no more than an extended paper. A paper stands (or

[13] The role of these archives is shifting. Originally, the main advantage of adding a paper to an archive was that it then became searchable; at that time, the major web search engines did not index formats such as PostScript. Today, their role is increasingly to ensure permanence—the content of a paper in an archive cannot be changed as easily as can that of a paper on an individual's website, for example—and to promote rapid dissemination of new work, for example through mailing lists.

does not stand) on the strength of the results. A thesis passes (or fails) on the strength of your demonstration of competence; even if good results are not achieved, the thesis should pass if you have shown the ability to undertake high-quality research. Questions that examiners might be asked to address include whether you have demonstrated command of the fundamentals of the discipline, whether you have the ability to correctly interpret results, and whether you have sufficiently strong communication skills.

A particular element of theses that is often weak is the analysis of the outcomes. All too often the discussion can be summarized as "the code ran", "it seems plausible", or "look at the pretty feature". To a greater degree than in a paper, it is necessary to probe why the outcomes occurred or what factors or variables were significant in the experiments. The guidelines to examiners issued by many universities state that the candidate must demonstrate critical thinking. Application of critical thinking and skeptical questioning to the work is an excellent way of persuading an examiner that the candidate understands their own methods and results; many of the questions explored in Chapter 10 concern critical thinking and skeptical examination of research.

Examiners are unlikely to be impressed by students who make grandiose claims about their work. Many researchers—and not just students—are reluctant to admit that their discoveries have any limitations; yet one of the clearest demonstrations of research ability is to ask incisive questions. Was the algorithm an improvement because of better cache use or fewer CPU cycles? What else would explain these results? In what circumstances is the theorem not applicable? A thesis with negative results can, if appropriately written, demonstrate the ability of the candidate just as well as a thesis with positive results. The outcomes may be less interesting, but the capability to undertake research has still been shown.

Examiners are also unlikely to be impressed by a student who accepts the word of established authority without question, or rejects other ideas without giving them due consideration, or appears reluctant to suggest any change or to make unfavourable comment. If you have a relevant point to make, and can defend it by reasonable argument, then make it. Be thorough. A PhD is an opportunity to do research in depth; shortcuts and incomplete experiments suggest shoddy work.

Issues such as whether results have been critically analyzed are of importance in papers, but there is a different emphasis for theses—it is you, not the research, that is the primary object of scrutiny.

For an extended research degree such as a PhD, another difference between a thesis and a paper is that the former may report on a series of more or less independent research discoveries. In contrast, a typical paper concerns a single consistent investigation. A thesis may, moreover, include work drawn from multiple papers. For this reason, there is more variation in structure from thesis to thesis than from paper to paper. An example of the problems faced in organizing a thesis is how to consolidate descriptions of new algorithms. It may make sense to bring all of them into a single chapter and then comparatively evaluate them in subsequent chapters, or it may be preferable to describe them one by one, evaluating each in turn. Factors to consider in choosing an organization include how cohesive the algorithms are (for example, whether they address the same problem) and whether an explanation of one algorithm is meaningful if the previous one has not yet been evaluated.

As the scope of a thesis is more substantial than that of a paper, the introduction may need to be broad in topic and conversational in tone. It could introduce a whole area rather than a single problem, for example, if the thesis happens to concerns a range of topics. Another reason to develop a substantial introduction is that a thesis is a more thorough, detailed document than is a paper. Why was the problem worth investigating in depth? How do the parts of the investigation relate to each other? What are some practical, concrete ways in which the outcomes of the work might be used? Running examples may be outlined in the introduction, to give unity to the thesis overall. The role of a thesis's introduction is, however, much the same as in a paper. As in the introduction of a paper, theory, jargon, and notation are inappropriate.

Take the time to learn about thesis writing as soon as possible. Browse other theses, from your own institution, from other institutions, and from other disciplines. Form views about the strengths and weaknesses of these theses; these views will help to shape your own work.

A writing-up checklist

- Have you identified your aims and scope?
- Are you maintaining a log and notebook?
- Does the paper follow a narrative?
- In what forum, or kind of forum, do you plan to publish?
- What other papers should your write-up resemble?
- Are you writing to a well-defined structure and organization?

- Have you chosen a form for the argument and results?
- Have you established a clear connection between the background, methods, and results?
- How are results being selected for presentation?
- How do the results relate to your original aims?
- Have you used any unusual patterns of organization?
- Have the results been critically analyzed?
- Are the requirements for a thesis met?
- Do you and your co-authors have an agreed methodology for sharing the work of completing the write-up?