



A Guide to Writing a Successful Paper on an Analysis of Algorithms and Data Structures

This guide describes how to explain your research in a persuasive, well-organized paper modeled on those published in computer science journals. For details on how to choose a project or conduct research on this topic, see [COMP 482 Project: Analysis of Algorithms and Data Structures](#)

To illustrate specific elements of this type of paper, this guide refers to a published sample (SAGAHYROON, A and ALOUL, F.A., 2007. Using SAT-based techniques in power estimation. Microelectronics Journal 38, 706-715.). Access the full paper via the Rice proxy at http://www.elsevier.com/wps/find/journaldescription.cws_home/405904/description#description.

Leading through Excellent Papers

As a member of a research community in either academia or industry, you will need to publish your ideas and research work. A research paper places your work in context, helping readers to understand why the study was undertaken, what others have done in this same area, and the important question (hypothesis) that you examined. Your paper allows you to convince the readers that your conclusions and recommendations, backed by your evidence, are sound.

Meeting Readers' Expectations through Structure

Readers expect certain questions to be answered as they read papers:

- Why is this work important (introduction)?
- How was it done (methods)?
- What were the results (results)?
- What do those results mean (discussion)?

A traditional scientific research paper is structured to answer these questions. However, computer scientists often choose to further subdivide the content, using headings that signal the technical content that they discuss in each subsection. Note the descriptive headings that Sagahyroon and Aloul use to organize their paper:

1: Introduction

2: Boolean SAT and integer linear programming

3: Problem description and formulation

3.1 Estimation of weighted maximum switching activity

3.2 An illustrative example

3.3 Maximum power-up current estimation

3.4 An illustrative example

4: Experimental results

4.1 Weighted maximum switching activity results

4.2 Maximum power-up current estimation results

5. Conclusions

This guide discusses how to prepare each of the traditional sections of a research paper and includes questions for you to consider as you write. Read each section of your draft to determine how well it answers these questions. Next, revise the draft so that it is clear, complete, and concise. The abstract, which appears first in publication, should be written last, and advice about it appears after information on the conclusion.

Introduction

The Introduction formally invites your readers into your paper by providing background information necessary for understanding your research. Your first sentences should signal the focus of your paper and engage the reader's attention, as in the following sentences from Sagahyroon and Aloul:

The last few years have seen a remarkable growth in the use of Boolean satisfiability (SAT) models and algorithms for solving various problems in electronic design automation (EDA). This is mainly [because] SAT algorithms have seen tremendous improvements in the last few years, allowing larger problem instances to be solved in different applications domains The growing integration scales in VLSI and the recent surge in the deployment and initialization of portable electronic devices have brought power dissipation to the forefront as a major design concern.

Once the reader's attention is engaged, most Introductions move from what is *known* about the topic to what is *not known* or *in question*, and then to the *specific question* that your research answered. As you write, include references to credible, peer-reviewed sources to support your claims. Your introduction should answer the following questions:

- What algorithms or data structures did you compare? How are they used? How have others studied them in the past?
- What question did you answer, or what problem did you solve through this comparison?
- Why is this question or problem important?
- How have others answered this question or solved this problem in the past? Why have these approaches been tried? To what extent are the approaches adequate?
- What was your specific hypothesis (i.e., what did you expect to happen)?
- How did you test this hypothesis?
- What insights does your comparison reveal (i.e., why is your specific approach useful)?

At the end of the Introduction, readers expect to find a concise forecast of the research work reported. This forecasting statement or paragraph offers the main point or claim of your research and previews what the reader will learn in the remainder of the paper. This forecast does the following:

- States the specific problem your research explored
- Highlights unique aspects of your approach
- Reviews, briefly, your most important findings (and why they are important)
- Provides a roadmap for the rest of the paper.

Consider the forecasting statement associated with the Sagahyroon and Aloul paper. Notice how these authors employ a consistent, easy-to-follow structure that leads readers through the main points of their paper.

The primary objective of this paper is to demonstrate the validity of ILP solvers, SAT-based and generic-based, in addressing the two [power] estimation problems discussed above.	<i>Summarizes the authors' specific claim or contribution.</i>
In this work, we first show how to model the two problems as SAT, i.e., 0-1 ILP, instances; and secondly we evaluate the performance of the latest SAT and generic ILP solvers in handling these instances.	<i>Summarizes the methods used and specific comparisons made.</i>
We show that generic ILP solvers are likely to achieve better results than SAT-based and random-based approaches when solving the proposed power problems.	<i>Describes the results obtained. Note that this summary could have been strengthened by including an explanation for why the generic ILP solvers perform better (though this paper actually makes no claims about why, but simply demonstrates the viability of generic ILP).</i>
The paper is organized as follows. Section 2 includes background information on SAT and ILP. Section 3 describes the estimation problems and their formulation using 0-1 ILP. Experimental results are presented in Section 4, and the paper is concluded in Section 5.	<i>Provides a roadmap for the rest of the paper.</i>

Methods (Algorithms and Implementation)

This section tells how you conducted your project. It should be detailed enough to guide someone who wants to reproduce your study. It should not, however, be structured as a chronology of what you did (for example, "First we tried this, but that didn't work, so then we did this."). If you encountered problems, first discuss the methodology you finally settled on.

Insert notes about problems and adjustments you made, when these are relevant to the final outcomes, in the Results/Discussion section.)

The Methods section of a computer science paper is usually divided into two parts. In the first part, identify the algorithms or data structures you investigated:

- Summarize the algorithms generally for your audience
- Highlight features relevant to your project
- Refer readers to your references for further details
- Use code snippets, descriptions, or diagrams to describe an algorithm or an optimization you employed. (Note: Do not provide your code as part of your paper. Further, do not spend a lot of time describing the algorithms, since the algorithms won't be your new contribution.)

The second part explains how you implemented the algorithms in your research.

Another way to think of the content of the Methods section is to imagine the questions that the two parts of the Methods section answer.

Part 1

- What algorithms or data structures did you select? Who created them? What is their asymptotic behavior? What other specific characteristics are worth noting for this study?
- What programming language and platform did you use? What impact did these choices have on your project?

Part 2

- How specifically did you implement the algorithms?
- How did you handle instrumentation code? Why?
- Did you perform any optimizations? Why or why not?
- How did you choose to test and benchmark your code?
- What inputs (data) did you select to test your implementations? Why?

Results / Discussion (Benchmarking and Analysis)

Unlike many articles in the traditional sciences, computer science research articles tend to combine results and discussion. Journals adopt this structure because the raw data are less important than the analysis. Provide representative examples that illustrate trends or outcomes that are important to your hypothesis. Focus most of your time, however, on the analysis of the results and what they reveal about the algorithms you are studying.

You will summarize the results of your benchmarking tests primarily in figures. **You will be graded on the quality of your figures.** Consider the following advice as you sort through the raw data:

- Select what is important to the question. Choose a representative subset (and be sure to justify your choice in the text). Don't display all of the data you collected. It's up to you to select and represent results.

- Make your data readable and understandable in tables and graphs, using colors, shading, or other means to highlight the points you plan to make.
- Plan to meet the needs of two kinds of readers. While many scientists will read your paper from start to finish, some scientists gauge whether or not to read an entire paper by first skimming the abstract and then examining the figures. To capture the interest of “skimmers,” ensure that your figures highlight your key findings and phrase the captions so that readers can follow your argument. Use informative titles that enable a reader to make sense of the figure without reading the paper’s text. Read through the figure titles, one after the other, and see if the titles capture the main points of your argument. To avoid confusion, check for consistency between the phrasing in the titles and the words in your figures’ axes or your tables’ column headings.
- In addition to employing informative captions, explain all your tables, diagrams, and figures in the text. Insert tables, diagrams, and figures **in the body of the paper for submission in this course**. Introduce a figure before it appears; discuss it afterward.
- Take a look at examples of good and poor graphs from papers from earlier years at the Cain Project Web site:
<http://www.owl.net.rice.edu/~cainproj/courses/comp482.html>

The analysis of the results is the most important part of your paper because it presents your original work and explains its significance. Your analysis confirms or denies your study’s basic hypothesis, interpreting your numerical results to demonstrate how they answer the questions you set out in the introduction. The analysis must support your conclusions.

Try to think about what your reader wants to know about your results. Here are some questions to guide you:

- What, specifically, did you learn from comparing these algorithms or data structures?
- What do your results say about the problem or question you were investigating?
- Was your hypothesis confirmed or disproved?
- Are the results what you expected?
- If you obtained anomalies or other unexpected results, can you explain them? If not, how could you set about in the future to identify what caused them?
- How do your results compare to past findings? Are they consistent? Different? Why?
- How would you respond to objections or questions that other researchers might have about your methods, results, or interpretations?

Conclusion

The conclusion answers the readers’ question: “So what?” It should give your readers points to “take home” from your paper. It should state clearly what your results demonstrate about the hypothesis you were testing in the paper. It should also generalize your findings, perhaps suggesting how others can use them in future research. All generalizations should be supported by your data, however; the discussion should have proved these points, so

that when the reader gets to the conclusion, the statements are logical and seem self-evident. No new evidence should be introduced in the conclusion.

Abstract

An abstract appears at the beginning of a paper, but you should write it last because it summarizes your whole paper. Readers should be able to read only the abstract and still understand your primary points. Your abstract should briefly answer the following questions:

- What question did you answer or what problem did you solve through this comparison?
- Why is testing a hypothesis about this question or problem important (why should your readers care about the work you conducted)?
- What methods did you select? Why did you select these particular approaches?
- What results did you obtain?
- What is the significance of the results?
- What do they tell about the comparison you made?
- Was your hypothesis confirmed?
- Are there any surprises or anomalies in the results and, if so, how might you explain them?
- What other interpretations might be considered?
- What additional research could be conducted based on your findings?

Acknowledgements

One of the ways that you demonstrate your professional identity is by acknowledging the help others have given you. For example, a professor or researcher may have shared a data set or testing facilities that made your project possible. Friends may have provided access to a manuscript that has not yet been published, or they may have answered a question through email correspondence. Another undergraduate may have run hours of tests for you or loaned you the use of a computer. Thanking such people demonstrates your ethical awareness that research is often, perhaps even usually, possible because of many generous individuals and organizations.

References

All facts, algorithms, figures, or tables taken from print or online sources should be cited and the sources listed at the end of the document in a section titled "References." An article that has been published in a peer reviewed journal and located in an on-line repository (such as JSTOR) is acceptable, but you must state the date on which you retrieved it and give the URL for its permanent location. Web sites posted by individuals are not acceptable references. Do not cite the course textbook for background.

All references should be in Association for Computing Machinery (ACM) style, which is summarized below. Full details on ACM style and style templates can be found at <http://oldwww.acm.org/pubs/> (scroll under Policies and Procedures for guidelines on plagiarism and electronic submission of articles using LaTeX or Microsoft Word).

In-text citation style: To note a parenthetical reference, enclose the last name of the first author and year of publication [Burando 2007]. When there are two authors, both last names and the year of publication are included [Burando and Lee 2007]; when there are more than two authors, cite the last name of the first author followed by et al. [Burando et al. 2007]. If multiple citations are needed, list all in square brackets separated by semi-colons [Burando 2007; Burando and Lee 2007; Burando et al. 2007]. When a citation is part of a sentence, the name of the author is NOT enclosed in brackets, but the year is: "So we see that Burando et al. [2007]..."

Journal reference style: ABDELBAR, A.M., AND HEDETNIEMI, S.M. 1998. Approximating MAPs for belief networks in NP-hard and other theorems. *Artificial Intelligence* 102, 21-38.

Book reference style: GINSBERG, M. 1987. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, Los Altos, CA.

Article in a book or collection: GREINER, R. 1999. Explanation-based learning. In *The Encyclopedia of Cognitive Science*, R. WILSON AND F. KEIL, Eds. MIT Press, Cambridge, MA, 301-303.

Conference proceedings: MAREK, W., AND TRUSZCZYNSKI, M. 1989. Relating autoepistemic and default logics. In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Canada, May 1989, H. BRACHMAN AND R. REITER, Eds. Morgan Kaufmann, San Mateo, CA, 276-288.
