

Writing a DPhil Thesis

Dr Maggie Charles, Language Centre

maggie.charles@lang.ox.ac.uk

1: ORGANISATION OF THE THESIS

Parts of the Thesis

The thesis usually consists of the following parts in this order.

Title page

Abstract

A short summary of the whole thesis.

Acknowledgements

Thank your supervisor first, followed by any other academic, technical or funding support. Finally, thank your friends and relations.

Table of Contents

This should list ALL parts of the thesis, including Bibliography and Appendices

List of Figures, Tables etc.

Chapter 1, 2 etc.

Introduction, several chapters which form the main body of the thesis, Conclusion. Chapter titles should be informative for the content of each chapter.

Bibliography

The list should include ALL the sources that you refer to (print and electronic), giving author, title and publication details. Formatting should be consistent.

Appendix A, B etc.

These give supporting information for your thesis. They provide extra, non-essential information.

Thesis length can be around 50,000-80,000 words or 150-200 pages (not including Appendices). Theses often have around 8 chapters, but your choice of organisation should follow from the nature of your own material. You will probably need about 2,000 words (4 pages) for the Introduction and rather more, around 3,000- 5,000 (6-10 pages) for the Conclusion, although there is a lot of variability. The rest of the chapters will also vary in length, depending on the number you include, but many will probably be around 12,000 words (20-25 pages). Again the length depends on the content.

Extract A Chapters: Program Synthesis from Domain Specific Object Models

Chapter 1. Introduction	1
Chapter 2. Introducing Booster	3
Chapter 3. A Mathematical Foundation	17
Chapter 4. The Structure of Booster models	37
Chapter 5. The Semantics of Booster	48
Chapter 6. Model completion	74
Chapter 7. From predicates to programs	87
Chapter 8. Discussion	115

The first chapter is a general introduction, while the second narrows this down to the specific topic of the thesis. Chapter 3 provides the mathematical framework for the study. Chapters 4-6 each focus on a different aspect of the specification language 'Booster'. Chapter 8 is the conclusion.

2: INTRODUCTION

The introduction is often one of the hardest parts to write. It may be easier to write it last, as you will then have a clear overview of your whole thesis. It gives the reader some general background to explain and justify the subject of the thesis. It establishes that there is a problem or gap in the research field, indicates how the problem will be dealt with and gives an outline of the thesis structure so that the reader can follow it more easily. There are often three stages (moves) starting with more general information, then establishing the problem or gap and finally narrowing down to specific statements about your own work. These moves may occur more than once.

Introduction

Move 1 Background (establishing the territory)

- Possible Steps
1. Why the area is important
 2. Giving background information
 3. Reviewing previous research

Move 2 Indicating a Problem or Gap in the Research Field (establishing the niche)

Move 3 Presenting the Research (occupying the niche)

- Possible Steps
1. Purposes, aims, objectives
 2. Work carried out
 3. Justification or importance of the research
 4. Structure of the thesis

TASK 1: Investigating an Introduction

- Extract B is part of an introduction to a thesis.
- Identify the 3 moves and the steps that make them up. The moves occur more than once.
- Underline the language used to signal these moves and steps.

Extract B Introduction: On the Specification and Analysis of Secure Transport Layers

Computer and communication security is important, sometimes critically so, to everyone...

The world is becoming strongly dependent on computers, and on distributed communication between remote devices...

All of these day-to-day activities, and many more, rely on security protocols in order to achieve authentication, confidentiality and data integrity between remote computers communicating over distributed, and inherently insecure networks. Finding strategies for discovering attacks against security protocols and for proving security protocols correct is an extremely important area of research.

When security researchers design security protocols, they take for granted the fact that they can build their protocol on top of a reliable transport layer protocol... Naturally, the security protocol is much simpler once these properties can be assumed of the transport layer.

An increasingly popular technique that is being used to simplify further the design of security protocols is to rely on a secure transport layer to protect messages on the network, and to provide some protection against attackers; see e.g. [FRH+08, OAS05b, Vis06, WSF+03]. By depending on a lower level protocol to provide a secure channel between hosts, a security protocol can be designed in a more transparent way...

In such circumstances it is important to understand what is required of the secure transport protocol, and, conversely, what services are provided by different protocols... In order to make the right decision about which secure transport layer protocols to use, it is vitally important to have a good understanding of the properties that they each provide. In order to do this, we require a means to specify these properties precisely.

In 2003, Philippa Broadfoot and Gavin Lowe proposed a layered approach to designing and analysing security transactions [BL03]. With their approach, a security transaction is layered over a lower-level secure transport protocol. The properties of the transport protocol are specified formally, and they can then be modelled abstractly.

There are several advantages to adopting this abstract layered approach... It makes sense... to model the properties of the secure transport layer abstractly rather than modelling the transport layer protocol itself; in order to do this, we require a precise way of specifying these properties.

There are many secure transport protocols that provide the same or similar properties... In order to compare and equate secure transport protocols in this way we require a framework in which we can specify their properties.

The aim of this thesis, therefore, is to improve our understanding of the security guarantees that can be provided by secure transport protocols...

We describe how one can build abstract models of the secure transport protocol properties in Casper, and we use the simulation relation to prove the soundness and completeness of these models... Finally, we present a new methodology for designing security protocols which is based on our secure channel properties... This new approach to protocol design simplifies the design process and results in a simpler protocol.

Thesis Overview

This thesis is split into eight chapters. In the next chapter, we give an overview of the necessary background material for the later chapters. We describe what a security protocol is, how attacks against security protocols can arise, and how they can be fixed...

Thesis Overview (continued)

In Chapter 3 we describe our research into specifying secure channel properties...

In Chapter 4 we describe several useful results about secure channel properties...

In Chapter 5 we examine the possibilities for chaining secure channels. We consider chaining channels in two different ways...

In Chapter 6 we present our abstract CSP models of the secure channel properties from Chapter 3...

In Chapter 7 we report on our analysis of two single sign-on protocols for the internet...

Finally, in Chapter 8 we conclude by summarising the contributions of this thesis. We describe several possible extensions to the specification and modelling work, and we discuss how this future work might be researched. We discuss possible extensions to the channel properties themselves, and interesting areas for studying the interactions between different secure channels...

3: THE BODY OF THE THESIS

These are the main chapters of your thesis and will therefore vary substantially according to the exact type of work you are carrying out. They are likely to be the longest chapters, although there is a very wide range in length (e.g. 3,000– 24,000 or 11-42 pages). Each chapter is organised into sections and sub-sections, again of varying length.

Extract C Chapter Structure: On the Specification and Analysis of Secure Transport Layers

Chapter 5 Chaining secure channels	101
5.1 Simple proxies	102
5.1.1 Secure channels through simple proxies	103
5.1.2 Simple chaining theorem	106
5.1.3 An automated proof of the simple chaining theorem	108
5.2 Multiplexing proxies	116
5.2.1 Secure channels through multiplexing proxies	120
5.2.2 Chaining theorem	122
5.2.3 An automated proof of the chaining theorem	124
5.3 Related work	130
5.4 Conclusions	131

The chapter usually begins with its own introduction, which links it back to earlier chapters and sets out the way in which the material is organised. It usually ends with its own conclusion, which summarises the results and may point forward to the next chapter. These parts of the chapter are important to ensure that the reader follows the long text easily and efficiently.

Extract D Chapter 5 Chaining secure channels

In Chapters 3 and 4 we exclusively described channels that secure point-to-point connections; **in this chapter we examine** the possibilities for chaining secure channels.

We consider chaining channels in two different ways: first, in Section 5.1, through a set of dedicated intermediaries (simple proxies), **and then, in Section 5.2**, through a (much smaller) set of trustworthy (multiplexing) proxies... **In Section 5.3 we discuss** some similar results about chaining secure channels discovered by other researchers. **Finally, in Section 5.4 we conclude and summarise our findings...**

Conclusions

In this chapter we investigated chaining our secure channel properties through a trusted third party (a proxy). **We showed**, in two different cases, that our channel properties are invariant under chaining ...

Finally, we compared our chaining results to similar results discovered and proved by other researchers.

SITUATION – PROBLEM – RESPONSE – EVALUATION

One useful way of structuring the sections of body chapters is to use the following pattern, which is frequent in many types of English informational writing.

The situation provides necessary background information and describes what you are writing about.

The problem is the reason for writing about the topic; it describes the specific nature of the difficulty.

The response gives details of what is done to solve the problem.

The evaluation assesses the effectiveness of the response. The evaluation may be positive or negative.

A negative evaluation can become a new problem which in turn requires a response, thus leading to a new cycle of the pattern.

Extract E Nominal Game Semantics

Chapter 3 Nominal Games

3.2 Innocence

In game semantics for pure functional languages, the absence of computational effects corresponds to innocence in the strategies. Here, although our aim is to model languages with effects, our models will be constructed by use of innocent strategies: the effects will still be achieved, by using monads.

Innocence is the condition stipulating that strategies be completely determined by their behaviour on P-views. In our current setting the manipulation of P-views **presents some difficulties** since P-views of plays need not be plays themselves. For example, the P-view of the following play (where curved lines stand for justification pointers) is $\cdot (\cdot, \cdot) \cdot \alpha$ and violates (NC2)...

We rectify these problems by explicitly imposing innocence on plays too.

4. CONCLUSION

In the Conclusion you summarise your research and then stand back from it and assess it as objectively as you can. This shows the examiners that you are capable of evaluating your own work according to the standards of the field. It is acceptable to mention areas that were not very satisfactory because this shows what you have learnt from carrying out the research.

There are often three moves, starting with the summary, then giving an evaluation, highlighting both the achievements and limitations and finally the possible future extensions of the work. The Evaluation move may also be carried out as part of the Summary. The Conclusion is the mirror image of the Introduction, in that it starts with the narrow concerns of the research and widens out to more general statements about further work.

Move 1: Summary

Summarises the most important aspects of the work.

Move 2: Evaluation

Indicates what the writer considers to be the most valuable contribution of the thesis and its limitations. Provides an assessment of how far the aims of the thesis have been achieved.

Move 3: Future Work

Gives possible extensions of the research.

Future work may follow from the limitations mentioned in the Evaluation move.

TASK 2: Investigating the Conclusion

- Extract F is part of the conclusion to a thesis
- Underline the language used to signal the three moves.
- What do you notice about the tenses and verb forms used?

Extract F Conclusion: Nominal Game Semantics

Summary

In this thesis we have examined the semantics of nominal computation, that is, of computation capable of creating fresh names, comparing them and passing them around. Following the work on the ν -calculus, a characteristic feature of our approach is the stipulation of names being constants rather than variables. We find this more adequate not only for denotational reasons... but also because it seems more intuitive: names are just like integers, but can only be compared for equality.

The constants-as-names rationale allows for a simple... modelling of names once an adequate framework for such constants has been laid down...

Our denotational models were built in game semantics, and in particular in nominal Games... We have examined names for general references and exceptions. The methodology followed was that of establishing a basic category of games corresponding to a basic nominal calculus, and from that obtain models of the additional nominal computational effects by means of computational monads. Thus, our models differ importantly from previous models of general references and exceptions. While in those models names were semantically modelled as compound objects encapsulating the structure necessary for name manipulation... in our models names are elementary objects manipulated by computational monads. This feature allowed us to obtain fully complete models without the need to add 'bad' constructors in the language.

Extract F Conclusion (continued)**Further directions**

This thesis has taken some basic steps in the modelling of nominal computation which can serve as a stable platform for further research on names. A first further direction is that of characterising the nominal effect... in abstract categorical terms. Here we have pursued this task to some extent by introducing the monadic-comonadic description of nominal computation, but it is evident that the description needs further investigation...

A second direction, which has not been pursued here, is that of decidability of observational equivalence in nominal languages... A major challenge to be faced is that the fully abstract models we have devised lean heavily on semantical quotientings and therefore disallow direct reasoning on strategies...

A third direction would be to examine nominal languages for concurrent computation... It would also be interesting to examine formal properties...

Finally, it would be interesting to examine AJM-games [AJM00] under the lens of nominal sets... We see that tags play the role of names and it seems therefore natural to use atoms as tags. This approach seems more natural than the usual naturals-as-tags and would greatly simplify the presentation of AJM-games, allowing also the nominal consideration of issues arising in Linear Logic and Geometry of Interaction.

5: ABSTRACT

The abstract provides a summary of the whole thesis, usually about 250-350 words. It may contain some or all of the following moves, often in this order. The abstract must be informative and ideally, attractive to the reader. Only use one or two sentences for each move. Keep the language and sentence structure simple, so that the reader can read the abstract and obtain the necessary information quickly and efficiently.

Move 1:	Background to the Thesis
Move 2:	Statement of the Problem/Gap in the Research
Move 3:	Purpose of the Thesis
Move 4:	Work carried out
Move 5:	Results
Move 6:	Conclusions/Implications
Move 7:	Contribution of the Thesis

TASK 3: Investigating the Abstract

- Extract G is the abstract of a thesis
- Identify the moves used in the abstract and underline the language which signals the moves.

Extract G Abstract: Program Synthesis from Domain Specific Object Models

Automatically generating a program from its specification eliminates a large source of errors that is often unavoidable in a manual approach. While a general purpose code generator is impossible to build, it is possible to build a practical code generator for a specific domain. This thesis investigates the theory behind Booster — a domain specific, object based specification language and automatic code generator. The domain of Booster is information systems — systems that consist of a rich object model in which the objects refer to each other to form a complicated network of associations. The operations of such systems are conceptually simple (changing the attributes of objects, adding or removing new objects and creating or destroying associations) but they are tricky to implement correctly. The thesis focuses on the theoretical foundation of the Booster approach, in particular on three contributions: semantics, model completion, and code generation. The semantics of a Booster model is a single abstract data type (ADT) where the invariants and the methods of all the classes in the model are promoted to the level of the ADT. This is different from the traditional view that considers each class as a separate ADT. The thesis argues that the Booster semantics is a better model of object oriented systems. The second important contribution is the idea of model completion — a process that augments the postconditions of methods with additional predicates that follow from the system's invariant and the method's original intention. The third contribution describes a simple but effective code generation technique that is based on interpreting postconditions as executable statements and uses weakest preconditions to ensure that the generated code refines its specification. (274 words)

6: AVOIDING PLAGIARISM

Plagiarism is defined as the use of other people's work and the submission of it as though it were one's own work. It may be intentional or unintentional.

In academic writing, it must be clear to the reader at any given point whose 'voice' is speaking. This is called **transparent source use**.

- **The responsibility of transparency in writing** means that the writer must use appropriate signals so that an experienced academic reader can understand the actual relationship between the source text and the new one.
- **The responsibility of transparency in language** means that language which is not signalled as a quotation is understood to be original to the writer. If the content is marked with a citation, the language is understood to be paraphrased, i.e., **substantially and independently reworded**.

The process of signalling the presence of another writer's voice in your text is accomplished through **citation**. Citation includes both **paraphrase with reference** and **quotation**.

A **paraphrase** is a **substantial rewording** of an idea from another text. The issue is not really how much has to be changed, but rather that the paraphrase should be **composed autonomously**.

A **quotation** gives the **exact words of the original**. It is marked by quotation marks and a reference to the original source with a page number. Use ellipsis (3 dots): ... to show that something has been omitted and brackets: [added material] to show that something has been added.

Depending on the type of text cited (i.e., book, conference paper, chapter in an edited volume, journal article, website etc.) the elements of a **reference** include:

author; date; title of the book or the article; title of the journal or other work; place of publication; name of the conference; date of publication; page numbers; website address and date of access.

We cite

- to give credit to the person(s) whose work we are referring to.
- to strengthen our arguments and to make it more likely that our points will be accepted.
- to discuss current issues in the field.
- to position our work within the field.
- so that the reader can locate and read the original source.

Unsuccessful Source Use occurs when **the relationship between the source and the new text is not transparent**.

The Consequences of Inappropriate Source Use

- The benefits of citation for the cited writer, the thesis writer and the reader disappear.
- The result can look like plagiarism.

What You Can Do

1. Think about the ways in which your research owes a debt to other, earlier research and acknowledge those debts.
2. Select the right signals to achieve transparency of source use.
3. Be meticulous about taking notes.
4. Know *why* you want to cite each source.

7: EDITING AND REVISING

Use the following points and questions to help you revise and edit your writing.

Communication

1. Think about what your reader already knows and what you need to explain. Are your explanations clear and understandable for someone who is not familiar with your work?
2. Are there any parts of the thesis where you should have included more information or explanation?
3. Are there any parts of the thesis which are redundant or which should be cut down?
4. The reader may need help in order to follow the flow of your ideas. Have you previewed the contents of each chapter? Have you ended each chapter with a brief review of what you have said?

EDITING AND REVISING (continued)

Ideas and Organisation

1. Have you introduced your subject appropriately? Have you developed it in a well-organised and logical way? Have you come to a conclusion?
2. Look at the connections in your text. Are the relationships between ideas clear and logical? Have you explicitly signalled the connections to the reader through the linking words you have chosen?

Language

1. Are there any parts of the text where you think grammar problems could interfere with understanding?
2. Are your sentences all full sentences with a subject and a verb?
3. Check that subjects and verbs agree, both singular or both plural.
4. Check your use of vocabulary. Look for unnecessary repetition, misused words, inappropriate colloquialisms and informal language.
5. Check your punctuation. Look for incomplete sentences and missing or wrongly used commas.
6. Check your paragraphing. Have you used paragraphs to break up the text according to the sequence of your ideas? Make sure paragraph breaks are at appropriate points.
7. Check for typographical errors and use a spellchecker.

8: HINTS ON WRITING

1. Plan how you will organise your thesis. Divide each chapter into sub-sections. This will break up the writing task into parts that are short and manageable.
2. Start with the part you find easiest, perhaps a description of what you have done. Then move on to the more difficult parts as you become more familiar with your topic and experienced in writing.
3. Get your ideas down on paper first, before you revise and edit to improve the language.
4. Don't expect to get the whole thesis perfect at the first attempt. Ideas change as your work progresses and multiple revisions are the norm in academic writing.
5. If you have a problem or a block, don't keep attempting to solve it. Rather, leave it and return to that point at a later stage or try to approach it from a different angle.
6. Print a draft version of each chapter when it is as good as you can make it. Then leave it for a time (overnight is ideal), before reading it to check for problems or parts that are unclear. Looking at your writing in print enables you to notice points that you may overlook when reading on screen.
7. When revising, the aim is to be critical, to read your thesis as though it was the work of someone else. Of course, it is also useful to exchange theses with another student and get their comments.