

How To Write a Good (no, Great) PhD Dissertation

Priya Narasimhan

Assistant Professor
Electrical & Computer Engineering
Carnegie Mellon University
Pittsburgh, PA

priya@cs.cmu.edu



Carnegie Mellon

Carnegie Mellon

Something About Me First

■ Background

- ▼ ~10 years of experience in fault-tolerant distributed systems research
- ▼ Both at the enterprise/desktop world and the embedded world
- ▼ Established a fault-tolerant industrial standard
- ▼ Previously CTO & VP of a company that I helped to start (based on my PhD research)

■ Research

- ▼ Dependable distributed systems

■ Teaching at CMU

- ▼ 18-349: Introduction to Embedded Systems
- ▼ 18-846: Fault-Tolerant Distributed Middleware Systems
- ▼ 18-549: Embedded Systems Design

2

India

- Born in India
- Part of my primary schooling
- My undergraduate education

Zambia

- Was mostly raised in Zambia
- Completed high school
- Finished my MS and PhD
- Best PhD thesis award
- Met my husband here ☺
- Started and ran a company

Santa Barbara

- Joined CMU in 2001
- Great research group

Pittsburgh

How I Got Here

3

Santa Barbara to Pittsburgh, huh?

- Why leave a well-paying industrial job?
 - ▼ Money is not the only motivator in life
 - ▼ Wanted to be in a job where I would learn EVERY day
- Academia is an exciting place to be in
 - ▼ Great research, great people, great environment
- Why did I want a research/academic life?
 - ▼ Like to challenge myself all the time
 - ▼ Love to build new things and solve new problems
 - ▼ Love to teach and share with others what I know
 - ▼ Chance to work with top-notch students from whom I learn every single day
 - ▼ And, of course, I love to talk ☺



Why Do Research?



- Thrill of finding out something that no-one else has done before you
 - ▼ Being a pioneer
 - ▼ Becoming a world-class expert in cutting-edge topics
- Going to conferences
 - ▼ Location, location, location,
 - ▼ Great motivator ☺
- Matters both in industry and in academia

5

So, What's a Ph.D., Anyway?

- Significant and substantial piece of ground-breaking work
 - ▼ Can be in one single field or inter-disciplinary
- It's about defining a hypothesis and providing arguments to substantiate or refute that hypothesis
- Evidence that **you** can do independent research that matters
 - ▼ Your work must somehow make a difference in your field
- Must have 2-3 key research ideas that you should be able to articulate at the drop of a hat
 - ▼ Definitely do not settle for less because of a rush to graduate
- Your Ph.D. dissertation is a significant piece of independent writing that you want to be proud of, for years to come
 - ▼ Don't write something in a hurry that you will cringe at later

6

Analogy –Breaking the “Sound Barrier”

- For many years, aircraft were not able to break the sound barrier
 - ▼ “... many people thought that any plane trying to fly faster than the speed of sound would break apart once it reached the "sound barrier"—and indeed, many planes that hadn't been properly designed for such high speeds were destroyed as they neared Mach 1.”
- The sound barrier was broken by a Bell X-1 piloted by Chuck Yeager on October 14, 1947
- The Bell X-1 included **only a few major architectural innovations** over prior aircraft that enabled it to fly at supersonic speeds
 - ▼ A unique .50 caliber bullet shape
 - ▼ Strong, super-thin wings
 - ▼ An adjustable horizontal stabilizer



*<http://www.californiasciencecenter.org/Exhibits/AirAndSpace/AirAndAircraft/BellX1/BellX1.php>
<http://www.nasm.si.edu/exhibitions/gal100/bellX1.html>

7

What’s a Great Ph.D. Dissertation?

- One that kicks the door open for another 3-4 future dissertations
 - ▼ Means that you started an area of work or a line of thought that opens many more lines of thought
- High-impact, i.e., “*Look, Ma, I*”
 - ▼ Have left my mark behind in my field
 - ▼ Fundamentally changed the way that something is done today
 - ▼ Introduced a new concept that can be “mined” by others in the future
 - ▼ Solved a problem that has plagued the field for years
 - ▼ Eliminated a fundamental assumption that has been made in the field
 - ▼ Provided strong empirical evidence that the field has been lacking
 - ▼ Will continue to influence the field and be cited by others
- Accessible to others in computer science/engineering
 - ▼ It’s great if you have a wonderful theory/system, but what if no-one except you (and your advisor) understands it?

8

Where Should You Start?

- Papers appearing in the top *conferences* (not necessarily journals) in your field over the past 2-3 years
 - ▼ Look at the best papers in those conferences
 - ▼ Look for taste in research, taste in presentation style, amount of work that it takes to have a best-paper award
- Theses of the ACM Doctoral Dissertation Award winners and honorable mentions over the last 2-3 years
- Your advisor's thesis
 - ▼ Helps you to understand how much "work" your advisor will expect
 - ▼ Ask your advisor what he/she is proudest of and what he/she would do differently, if given a chance to re-write the dissertation
- The theses and recent papers of your committee members
 - ▼ This will help you to understand their outlook on publication and writing

9

Step 1 – Writing Your Dissertation Abstract

- Paragraph 1: What is the problem?
 - ▼ Not more than 3-4 sentences telling the reader what the problem is, in as simple English as possible
- Paragraph 2: Why is the problem hard?
 - ▼ What has eluded us in solving it?
 - ▼ What does the literature say about this problem?
 - ▼ What are the obstacles/challenges? Why is it non-trivial?
- Paragraph 3: What is your approach/result to solving this problem?
 - ▼ How come *you* solved it?
 - ▼ Think of this as your "*startling*" or "*sit up and take notice*" claims that your thesis will plan to prove/demonstrate
- Paragraph 4: What is the consequence of your approach?
 - ▼ So, now that you've made me sit up and take notice, what is the impact?
 - ▼ What does your approach/result enable?

10

Step 2 – Your Thesis Title

YOUR THESIS TITLE

CONDENSING OVER HALF A DECADE OF YOUR LIFE IN ONE SENTENCE.

www.phdcomics.com
JORGE CHAM © 2006

the colon
Can't decide what to title your thesis? Use a colon!

a preposition
A good preposition tells your readers "hey, this is not just a futile exercise"

"Witty catchphrase"

: Length-enhanced superlative verbiage with prolixity

in/of/for

Obscure topic few people care about.

witty catchphrase
Makes people think you're hip and culturally relevant. Only marginally related to the actual thesis? No problem.

the boring stuff
Nothing says "academic rigor" like a long string of dry scientific-sounding terminology and fancy buzzwords.

obscure topic few people care about
Sad, but true.

- Think of what you would want others to "google" your thesis for
 - ▼ Be precise and don't look for "wildcard" words that cover a range of topics
- You should be able to state your contribution/approach/result in no more than 7-8 words (that's the ideal thesis title)
- Look through your dissertation abstract
 - ▼ What are the half a dozen keywords that you would italicize for emphasis?

11

Step 3 – The Dreaded Thesis Outline

WRITING YOUR THESIS OUTLINE

NOTHING SAYS "I'M ALMOST DONE" TO YOUR ADVISOR/SPOUSE/PARENTS LIKE PRETENDING YOU HAVE A PLAN

STEP 1 Aim for a respectable number of chapters:

THESIS OUTLINE

- 1.
- 2.
- 3.
- 4.
- 5.
- 6.
- 7.

chapter #'s

5 = "That's IT??"
6-7 = "Not bad"
8+ = "Are you crazy??"

STEP 2 Fill in the "freebies":

THESIS OUTLINE

1. INTRODUCTION
2. LIT REVIEW
3. METHODOLOGY
- 4.
- 5.
- 6.
7. CONCLUSIONS

You're half way done!

STEP 3 Make up titles for the "meat" chapters:

6. LIT REVIEW
3. METHODOLOGY
4. (THAT STUFF YOU DID YOUR FIRST YEAR)
5. (STUFF YOU'RE SUPPOSED TO BE DOING NOW)
6. (MAKE STUFF UP)
7. CONCLUSIONS

(It'll be years before you actually have to work on that later chapter, and by then your thesis topic will have changed anyway)

STEP 4 Voilà! You just bought yourself another two years



www.phdcomics.com

12

The Thesis – Introduction

- Motivate the problem and state your hypothesis
 - ▼ Tell a story, and tell it well
 - ▼ Use plenty of concrete examples (or a running example) and figures
 - ▼ Quote data sources, e.g., industry analysts, market surveys, case studies
 - ▼ People often (and naturally) make up their mind within the first few pages
 - ▼ Introduce all your terminology here – especially, acronyms you plan to use often
- Do
 - ▼ Provide a concrete problem definition, accessible to a computer-literate person, without “dumbing down” the problem to people in your field
 - ▼ Provide a concrete list of your thesis’ contributions
- Don’t
 - ▼ Oversell your thesis or its claims – be honest and you will be respected
 - ▼ Use hyperbole (e.g., “highly reliable”, “extremely efficient”)
 - ▼ Try to confuse the reader with big words – plain, simple English is best
 - ▼ Try to sound like your thesis covers your entire field (unless it does, of course!)

13

The Thesis – Related Work

- Know the key papers and players in your field
 - ▼ Survey your field in depth, knowing the seminal *and* the recent work
 - ▼ When you read a paper in your field, make sure to cover all the papers referenced in that paper as well
 - ▼ End-result should look like an *ACM Computing Surveys* article
- Do
 - ▼ Mention all of the related work in your thesis
 - ▼ Acknowledge the role that each paper has played in evolving your field
 - ▼ Be constructively critical of where the shortcomings of each paper are, to the extent that it justifies your approach
- Don’t
 - ▼ Belittle a paper just for the sake of “showing off” how much you know
 - ▼ Make the criticism harsh/personal, just because you do not like the author or the author has disagreed with you or disapproved of your approach

14

The Thesis – The Meat

- This is the part where you reel off a sequence of chapters, each with a unique result or building-block
- Ask your advisor for presentation style and help with outline
- Think of this as a sequence of 2-3 *distinct top* conference papers that you have published
 - ▼ There should be a natural progression from one chapter to the next
 - ▼ Keep in mind that you are still telling a story
- Use figures, and plenty of them
 - ▼ They draw the reader in and make the thesis more interesting
 - ▼ Can convey a lot more information than text, sometimes
- Ways to present your data
 - ▼ *Visual Display of Quantitative Information*, by Tufte

15

The Thesis – The Meat

- Do be clear and candid about
 - ▼ Your *assumptions* – yes, every one of them
 - ▼ Your *limitations* – yes, every one of them
 - ▼ *Requirements* of your solution/approach – both mandatory and optional
 - ▼ *Constraints* under which your solution will work
 - ▼ Above all, *why* these assumptions, limitations, requirements & constraints
 - ▼ A concrete *validation* plan for your hypothesis – experiments, simulation, theorems, proofs, etc.
 - ▼ *Scope* – what's part of your thesis and what is definitely not
- Don't
 - ▼ Expect to shoe-horn all of the work that you did during your graduate research career into your thesis
 - ▼ Present a set of scattered, unrelated results that don't add up to a whole
 - ▼ Tout all of the advantages of your approach repeatedly
 - ▼ Conjecture wild promises from your results (i.e., stay factual throughout)

16

The Thesis – The Meat

- Presenting your approach or methodology well
- Tell the reader why you picked this approach
 - ▼ Did you know that it would work? Did you have a basis for knowing this?
 - ▼ What was your overall philosophy in your approach?
- What other approaches did you consider and discard?
 - ▼ Where did they fall short? How were they inappropriate?
- What interesting negative or counter-intuitive results do you have?
 - ▼ For instance, are there instances of where your hypothesis breaks down?
- Two questions that are almost always part of any Ph.D. defense
 - ▼ How do you know that you are done? When is the problem solved?
 - ▼ If you had to do it all over again, what would you do differently?

17

The Thesis – Wrapping Up

- Future Work
 - ▼ Don't view this necessarily as a list of the limitations of your thesis
 - ▼ Think of what you would do if you had an extra year in your Ph.D.
 - ▼ Don't worry – this is not for your advisor to hold your feet to the fire
 - ▼ Think of 2-3 other follow-on Ph.D. dissertations that you can envision
- Conclusions
 - ▼ Be reflective and honest
 - ▼ What were the lessons learned?
 - ▼ What were the overall insights?
 - ▼ Did you solve the problem completely? How much progress have we made in your field because of your work
 - ▼ Don't bore the reader with a cut-and-paste of your Introduction chapter

18

Within Each Thesis Chapter



■ Introduction

- ▼ What is this chapter all about?
- ▼ What sub-problem or issue is this chapter addressing?
- ▼ How does this chapter fit within the overall “story” of the thesis?

■ The Meat

- ▼ Rigorous approach to sub-problem, or detailed explanation of issue
- ▼ Assumptions underlying sub-problem, or complete description of issue
- ▼ Validation: System design, theory, implementation, graphs, references,

■ Summary

- ▼ Repeat the highlights of the chapter
- ▼ Transition sentence that acts as a “teaser” for the next chapter, and how the next chapter fits with the current one

19

Getting Through the Tunnel



■ Recruiting thesis committee members

- ▼ You want people who are “in your corner”
- ▼ Also look for letter-writers for jobs
- ▼ Look for people who expect high-quality work of themselves and their students – they will give you their best

■ Talk to your thesis-committee members often

- ▼ Provide them with a detailed status update at least once a semester/quarter
- ▼ Avoids any last-minute surprises or miscommunication about expectations

■ Thesis draft – be professional

- ▼ Provide a thesis draft to your thesis committee members at least 1-2 months in advance of the defense
- ▼ Use spell-check tools before you hand off your first draft – nothing worse than a sloppily written thesis that the author forgot to proof-read
- ▼ Give committee members a time-frame for receiving subsequent drafts

20

Baby Steps ... Take Them Today!

- When you walk out of here **today**
 - ▼ Write down a tentative thesis title
 - ▼ Yes, I mean that, even if your thesis is murky in your mind
 - ▼ Write down a first cut at your dissertation abstract
 - ▼ Yes, I mean that, even if you have not done the work yet!
 - ▼ Why? It's a start and a commitment to something bigger
- Write down the road-map of your thesis **today**
 - ▼ What is the ideal thesis that *you* would wish for?
 - ▼ What results would it contain?
 - ▼ How would you evolve the story from start to finish?
- All of this will change, but do it today so that you have something inspirational to start changing
 - ▼ Also, great for putting a smile on your face years from now
- *"A man may write at any time, if he will set himself doggedly to it"* [Samuel Johnson]



21

Being Able to Summarize Your Thesis


- Have a slide for each of the following
 - ✓ Problem statement or Hypothesis
 - ✓ "Why is this a Hard Problem?"
 - ✓ Approach or Methodology
 - ✓ "Why is this Innovative?"
 - ✓ Assumptions and Constraints
 - ✓ Initial Results (Promise of Great Things to Come)
 - ✓ Validation Plan
 - ✓ Limitations and Applicability
 - ✓ Expected Contributions
 - ✓ "What's beyond this thesis?"
 - ✓ Roadmap of Thesis

22

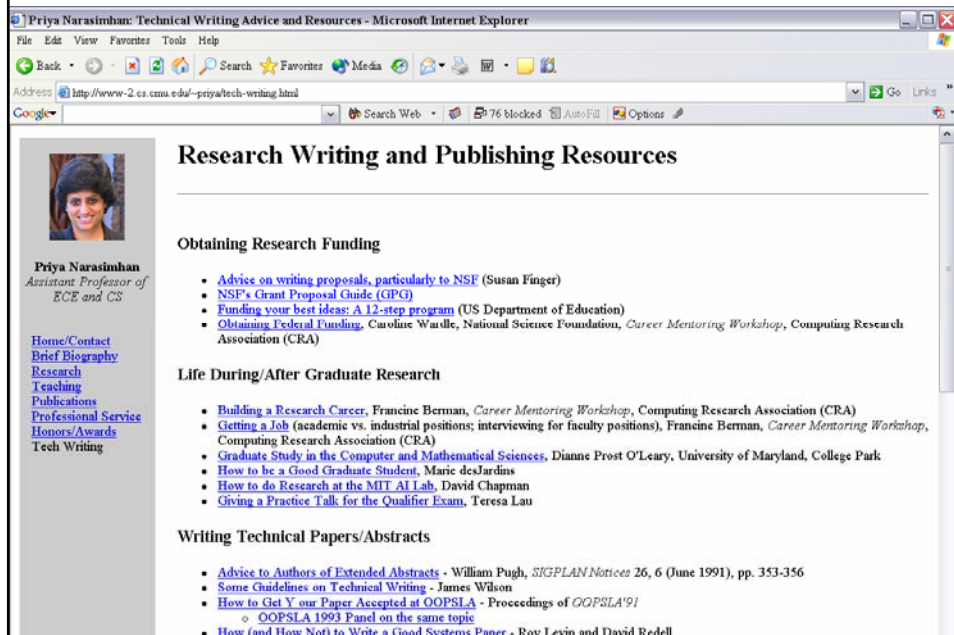
Being Able to Summarize Your Thesis

- Be prepared to answer the following questions for yourself
 - ▼ Why will your work change the world (or at least your field)?
 - ▼ What would you do differently, if you had a do-over?
 - ▼ How will you know you are done? What does success mean?
 - ▼ How much of your work can be generalized?
 - ▼ What part of the work was research and what was engineering?

23



Some Resources



Priya Narasimhan
Assistant Professor of
ECE and CS

[Home/Contact](#)
[Brief Biography](#)
[Research](#)
[Teaching](#)
[Publications](#)
[Professional Service](#)
[Honors/Awards](#)
[Tech Writing](#)

Research Writing and Publishing Resources

Obtaining Research Funding

- [Advice on writing proposals, particularly to NSF](#) (Susan Finger)
- [NSF's Grant Proposal Guide \(GPG\)](#)
- [Funding your best ideas: A 12-step program](#) (US Department of Education)
- [Obtaining Federal Funding](#), Caroline Wurdle, National Science Foundation, *Career Mentoring Workshop*, Computing Research Association (CRA)

Life During/After Graduate Research

- [Building a Research Career](#), Francine Berman, *Career Mentoring Workshop*, Computing Research Association (CRA)
- [Getting a Job](#) (academic vs. industrial positions; interviewing for faculty positions), Francine Berman, *Career Mentoring Workshop*, Computing Research Association (CRA)
- [Graduate Study in the Computer and Mathematical Sciences](#), Dianne Prost O'Leary, University of Maryland, College Park
- [How to be a Good Graduate Student](#), Marie desJardins
- [How to do Research at the MIT AI Lab](#), David Chapman
- [Giving a Practice Talk for the Qualifier Exam](#), Teresa Lau


Writing Technical Papers/Abstracts

- [Advice to Authors of Extended Abstracts](#) - William Pugh, *SIGPLAN Notices* 26, 6 (June 1991), pp. 353-356
- [Some Guidelines on Technical Writing](#) - James Wilson
- [How to Get Your Paper Accepted at OOPSLA](#) - Proceedings of *OOPSLA'91*
 - [OOPSLA 1991 Panel on the same topic](#)
- [How \(and How Not\) to Write a Good Systems Paper](#) - Roy Levin and David Redell

It's the Little Things That Count

- Ask your peers and colleagues in your research group for feedback
- Ask your advisor for continuous feedback
 - ▼ Your advisor is your biggest champion, your biggest promoter
 - ▼ This person *wants* to see you succeed and will rave about you and your work enthusiastically to everyone
 - ▼ Discuss your thesis layout, problem definition, goals of the month, etc.
 - ▼ For sticky issues, your advisor will find you the right "connections" to fill in the gaps in your thesis
- *"We are like dwarfs on the shoulders of giants, so that we can see more than they, and things at a greater distance, not by virtue of any sharpness on sight on our part, or any physical distinction, but because we are carried high and raised up by their giant size."* [Bernard of Chartres, 1130 AD]
 - ▼ Acknowledge the people who helped you, even if it was a little thing
 - ▼ This is the one opportunity you have to thank everyone publicly for their support and their help during a critical time in your life

Attending Conferences



- Important for networking and making contacts
 - ▼ Introduce yourself to professors and other students
 - ▼ Use business cards, if you have them
 - ▼ Volunteer for service in some conferences
- Why should you care?
 - ▼ Build a community around yourself outside your university
 - ▼ People should get to know **you**, not just your advisor
 - ▼ Letters of reference, thesis committee members, job interviews, etc.
- There might be apparent discouragements on the road
 - ▼ Faculty at conferences might seem unapproachable
 - ▼ You might not (yet) have a track record in research
- Don't stand around, waiting for people to talk to you
 - ▼ Aggressiveness is a great quality – don't be shy

27

How To Do Your Best to Fail Your Defense

Source: Doug Comer's "How to Escape at the Last Minute"



- Redefine basic terminology or define new, absurd terminology
 - ▼ For example, say "During this examination, the term *yes* will sometimes mean *no*, *maybe* will sometimes mean *never*, and *no* will sometimes mean *maybe*"
 - ▼ When the astounded panel asks, "Are you serious?" stare at them with a deadpan expression and answer *yes*. Let them try to decide what it means
- Argue with the examiners. It helps if you can slander each of them independently. At the very least, question their credentials: "Who do you people think you are, anyway?"
- Listen carefully to the questions, and take advantage of the wording
 - ▼ If an examiner asks, "Can you outline for us the exact procedures used in your research, and note any exceptional or unusual techniques?"
 - ▼ Simply mutter, "yes" and wait
 - ▼ When the examiners become impatient, point out that they have no grounds for being upset because you have correctly answered the question

28

How To Do Your Best to Fail Your Defense



Source: Doug Comer's "How to Escape at the Last Minute"

- **Stick to the subject at hand, but give as many low-level technical details as possible**
 - ▼ Get down to the bits and stay there
 - ▼ Avoid all concepts and summaries
 - ▼ Use tables of numbers instead of graphs
 - ▼ Introduce long, complex equations without balanced parentheses
- **Add a dozen caveats to each answer**
 - ▼ Begin by saying, "Of course, my answer depends on the communication system available in the country, the supplier used for spare computer parts, a local power company can provide uninterrupted electric power, the probability that cosmic rays from outer space strike anyone or anything involved, the number of days lost because someone is sick, and the stability of world economic markets.."
 - ▼ The beauty is that most of what you say is true
 - ▼ If the panel asserts that you don't need to worry about all that, ask them, "You mean you don't care about a stable world economy? What kind of human beings are you, anyway?"

29

Final Thoughts

- **You are in some of the best universities in the world**
- **Make full use of the opportunities at your university**
 - ▼ Don't be intimidated – every professor was once a student, too
 - ▼ Be aggressive – what do you have to lose? Nothing!
- **Look around – your fellow-students might be**
 - ▼ Tomorrow's ground-breaking researchers
 - ▼ Tomorrow's CEOs and Vice-Presidents
 - ▼ Tomorrow's professors
 - ▼ Get to know them today!
- **Learn at least one new thing every day**
 - ▼ That's at least 365 more new things you will know every year!



30

And Don't Forget – “*The Key to Life is Balance*”

