Design:
The Key to Writing
(and Advising)
a One-Draft Thesis

John Carlis

Computer Science and Engineering
University of Minnesota
{carlis@cs.umn.edu}

Summary

You succeed by seeing the whole thesis as flesh on the bones of the design, with contributions the *raison d'etre* of the design.

To Succeed:

1. Believe	**In a one-draft thesis.**	*Gumption*
2. Focus on	**Contributions to the field.**	*Destination*
3. Avoid	**The fatal writing error: wrong audience.**	*Sketch*
4. Acquire	**Draft-writing skills.**	*Craft*
5. Follow	**The "Design a Thesis" way of work.**	*Don't Hack*

1. Believe **In a one-draft thesis:** *Gumption*

To begin consider the following questions.

What does "draft" mean? A draft is a completed work, something that you give to others for review. I do not consider small scale revising, say, editing within a paragraph while leaving the structure the same, as re-drafting. Do not misunderstand; "one draft" does not mean that you just start from a blank slate with final words of truth and beauty rolling off your fingertips. No, writing is work, but it is merely work, not some mystical thing.

What can motivate you to follow, or at least consider, my advice? Four things can. First, since a thesis is territory unknown to you, you should be wary, if not scared, and eager for advice on how to succeed. The thesis for most students is the largest writing project they will ever undertake. So I doubt that you have the experience to say "no problem". Second, for many students it is a painful process. Sadly, despite years of hard work, they struggle to produce a barely acceptable thesis. Third, *ceteris paribus*, your committee will consider your thesis as having higher quality substance if it is better written. So you benefit by writing well. Fourth, the thesis is more important than other writing that you have done — your career hangs in the balance, because: no thesis, no faculty job.

What makes me qualified to give advice about a one-draft thesis? Experience does. I partially did it. Two of my advisees (Leone Barnett and Libby Shoop) have done it. Others are in the middle of doing so.

What's my story? While working on my thesis, and at the same time a professor at Syracuse, several streams of thought happily converged. I read the John McPhee Reader, which had a forward describing his disciplined, design-before-write way of work[1]. I read Richard Mager's Preparing Instructional Objectives, which has this message: teaching (writing) is about them, the students (readers), not me. I read several books on writing. I taught software development using Yourdon and Constantine's "Structured Development" (Addison-Wesley, 1975). Convergence began when the strong parallels in their content struck me.

I already believed that software should be designed and not hacked, and came to the same conclusion about my thesis. I decided to try to transfer software notions to writing. Since a paragraph is a unit of development, like a software procedure, I chose to design down to the paragraph topic sentence level before writing. And it worked! Now, as an advisor, I coach my students so that they too can successfully write a one-draft thesis. My goal here is to codify this effective and efficient way of work, so that I more readily can pass it on to my own students, and to share it with other students and their advisors.

What should you believe? Three things. First, believe that re-drafting is both costly and avoidable. By writing only one draft, you can save several months worth of effort. Second, believe that a one-draft thesis is preferred, and not an indicator of lack of effort. You might have to fight your own or your advisor's preconceptions about this point. As you will see, I consider re-drafting to be a consequence of incomplete design. Third, and most important, believe that indeed you can write a one-draft thesis.

These beliefs will raise your gumption level[2]. When you understand what is going on and why, then you will have gumption. If not, then you can worry, or do wasteful things in the vague hope that good things will occur.

---

[1] McPhee won the Pulitzer Prize for "Annals of a Former World", Farrar Straus & Giroux, 1998. Besides writing great books, he teaches writing at Princeton.

[2] See R. Pirsig's "Zen and the Art of Motorcycle Maintenance", Morrow, 1974. It's great book; philosophy with a plot and characters. Read it.

2. Add Contributions to the field. *Destination*

You need to understand that the purpose of a thesis is to defend your claimed contributions, that is, your additions to your field of study. In other, more dramatic words:

- No contributions, no thesis!

Where do contributions come from? Your abiding interest in a topic, and hard work enable you to have something interesting to say to others. To obtain contributions you might state a scientific hypothesis that you test, or prove a theorem, or engineer a solution to a problem, inquire into the state of some phenomena, or use some other scholarly mechanism.

You begin thinking about contributions much earlier than draft-writing time. In your proposal you have target contributions. In pre-proposal pondering you try out possible contributions. You imagine potential "big WOWs", that is, if all goes well, what would turn your field on its ear. If you have no chance of a big wow, pause and examine where you are heading. Also, assess risk: if something goes wrong, will you still graduate?

Here is a way to think about contributions that may help you. In the figure below, the striped blob represents unknown territory. You contribute knowledge to the field, represented by the clear box. When you bound your thesis work you decide where to place the box, and its dimensions. So you choose your thesis contributions by coming in from the larger context.

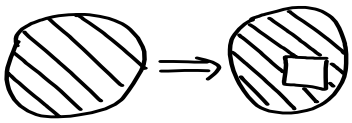

Here, in picture form, are two things you should avoid:

On the left, you propose to take something you have already done (perhaps a master's project), and, without considering the context, say "for my thesis I'm going to add this bit". Don't try to just move out from where you are. On the right, you propose to claim three, small, weakly or unconnected contributions, and claim that they make a thesis. A thesis is not a grab bag, so avoid this mistake.

From beginning to end, keep the contributions in the forefront of your thinking. Put a copy in front of your eyes. Critically review them often. Polish them as you design; it is well worth the effort. Talk to others about them, because, if you can explain your contributions to somebody outside the field, then your committee will be more likely to smile on them.

In spite of its importance many students, and some advisors, misunderstand "contribution". To grok[3] it consider these three important terms: novelty, consequence, and evidence. To your committee, who represent the other[4] experts in the field, you will say "Look, here is something new and valuable, that somebody cares about, and I can prove it". Your goal is to get the committee to sign the paperwork allowing you to graduate. You do that by informing them of what your contributions are, and by persuading them that your contributions are valid and useful.

[3] Grok was coined by Robert Heinlein in "Stranger in a Strange Land", Ace Books, 1961. It means "deeply understand".
[4] By the time you defend your thesis you must be an expert. So act the part.

To persuade them about novelty, you present a literature review, an organized critique of who has done what, their contributions, and the limitations of their contributions. You need to include the relevant literature, but exclude the irrelevant literature — thereby showing that you are an informed expert who can make judgments. You need to present it cogently, succinctly, pithily, and other –ly words to show that you are not a ninny, and can distinguish dross from diamonds. Since you are exploring unfamiliar territory, you should expect to resolve vocabulary mismatches, e.g., "Smith 1990 coined the term X, but Jones 1991 uses X in a subtly different fashion … ". You should expect to invent new vocabulary, and then precisely define it. You demonstrate expertise by picking an appropriate vocabulary to focus on, and organizing it. The literature review also serves as the substrate upon which you show the goodness of your additions. In your proposal a literature helps you describe potential contributions. (More to come on this topic.)

To persuade your committee about consequence, you state who cares about your contributions, and what impact your contributions have on them. Of course you care, but that is immaterial. If you produce something new, but it has no impact, it is not valuable to the field, and, therefore, not a contribution.

Avoid the error of thinking that a contribution is about you — it's about them, the people in your field. So do not stop at saying "I did this", because that alone does not make it a contribution. You continue and say "and here's its impact on them". Furthermore, if you learn something (you improve), it is not necessarily a contribution to them.

To persuade your committee to accept your claims about the impact of your contributions, you need to provide evidence. If you learn something new that will, if true, have impact on the field, but you have no evidence, then you have no contribution. Evidence persuades the committee that you have been scholarly, and have fairly assessed the validity and strength of your contributions.

Do not write the bulk of the thesis, and then try to determine what contributions you can claim. That's the tail wagging the dog, and is a bad mistake.

Knowing what contributions you need to defend helps you in five ways. First, you can figure out when you are done. Quite a few students plead "I'm tired. Can I graduate?", or "I've been at this for six years, can't I stop now?", or some such irrelevancy. You are done when you have defended contributions — not before, not after. Second, contributions help defend you from your advisor. Here is a true story. A faculty candidate, interviewing in January, said "I can be here Fall semester. I want to be done in May, but my advisor said, 'no, not until December'. We compromised on late September". Don't let this happen to you — schedule doesn't matter, substance does. Third, it helps you decide what you need not write about. If the audience, your committee, already knows and believes something, or it is not relevant to the contributions, then you do not need to produce a treatise on it. Fourth, it helps you decide about what and in what depth you must write. Fifth, it helps you decide what other work you need to do. For example, suppose you want to claim a particular contribution, but cannot prove it. You must drop the claim, or weaken it, or do more work to gather evidence, or move it to a future work section.

Think of the contributions as the *destination*. Research explores new territory. When you reach the destination you are done. While you still have to choose a path to get there, and that will take work, at least you will not wander. You wander if you work on notions that do not relate to your contributions. For example, I knew a student who added forty pages to his thesis saying "yeah, I know this is not germane, but I want others to read it, so I put it in anyway". Another student spent a year devising evidence for his novel software, but, sadly, he picked a sample problem that could not possibly have shown the goodness of an intended contribution.

Think in terms of destination when somebody says "tell me about your thesis". If you reply with "I'm studying XXX; it's fascinating and I'm learning a lot", then you are thinking as a student, not as a research scholar. You are thinking in terms of your own movement, and not about the territory — and what's new to you may not be new territory to the field. Yes, of course, you indeed should have an abiding interest in the topic. Don't work on it unless you find it interesting enough to continue working on it for years after you finish the thesis. Yes, you will learn, but that's irrelevant for research. If you reply with "I'm hoping to find something new about XXX", then you are thinking about new territory, but only about direction, not how far in that direction. (Use this notion of destination when you are in pre-proposal or proposal stage.) Remember, you will say "Look, here is something new and valuable, that somebody cares about, and I can prove it".

3. Avoid The fatal writing error: wrong audience *Sketch*

Not all writing is the same. In particular, writing for yourself qualitatively differs from writing for others. In stick figure form:

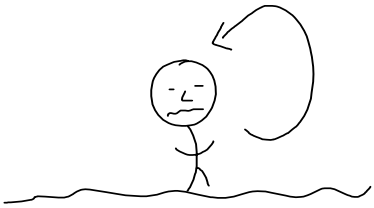

is not

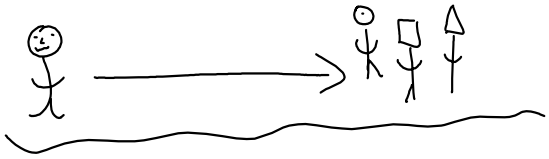

.

When you write for them (the audience), you put forth effort to make your product good: neat, complete, proffed (sic), unified, cogent, etc. That is, you draft a well-designed product for them to read. However, if you are the audience such effort is a waste.

An analogy can help make this clear. Everybody understands that the cardboard model, or the blueprint, of a design is not the house, and that it is a waste of time (foolishness) to make the model's toilet flush. However, it is easy to forget, when writing, that designing is not the same as drafting — because both activities use the same material, namely, words. So, don't fool yourself, and commit the fatal error of confusing the model with the house.

Why is this distinction important? Well, you care about cost, and you can change the model cheaply (in time, effort, or money), but it will cost you dearly to change the house. Moving a window on a blueprint costs pennies, but moving a window on a house costs thousands of dollars. Likewise, redrafting is expensive, but modifying your design is relatively cheap. Furthermore, the distinction's importance grows exponentially with size (a thesis is indeed big), and your inexperience. A thesis is a really big writing task, and, I hope, your first one — one Ph.D. is plenty.

You would not build a house without a blueprint, so do not write a thesis without a design.

To help you avoid the fatal error, think of writing for yourself as *sketching*. If you do, then you expect several things:

- Changing sketches — do not try to get it right the first time.
- Incomplete sketches — do not feel driven to finish them.
- Messy sketches — do not waste time making them pretty.
- Discarded sketches — keep the ideas, not the medium on which you jotted them down.

Why sketch? Well, you can have only a vague idea of your destination (aka contributions), and you sketch to generate ideas. You can sketch to try out a collection of terms, or an exposition strategy to see what its strengths and weaknesses are, so you can compare it to other things you have sketched.

When you design the thesis you still are writing for yourself, so think of it as coming up with a blueprint, which is more formal than a sketch, but is still clearly not a house. A blueprint shows the important design decisions that affect the building of the house, and shows enough detail so that competent workers can build the house. A thesis's detailed outline is like a house's blueprint. You design to a level of detail where you, as a competent writer, can confidently proceed to construct the draft. But how detailed? I claim that you should design down to the paragraph

topic sentence (PTS) level[5]. This key to success is important enough to say another way: don't write the draft if you have a fuzzy design. (More to come on this topic.)

Finally, if something goes wrong during house construction, you don't hack at the house. Instead, you consult the blueprint, fix the design (perhaps sketching alternative solutions), and then resume construction. Do likewise when drafting a thesis.

[5] Barnett and Shoop both got in the habit of shortening "Here is my detailed design down to paragraph topic sentences" to "Here are my PTSs". So now I use the acronym too.

4. Acquire **Draft-writing skills** *Craft*

Are you already a good writer? If so, well that's terrific, but expect that writing your thesis will challenge you anyway. If not, then you better exert yourself to get better. Don't wait until you are partway through a bad draft to begin remedial training. By the way, do not be smug if English is your native language. I have seen a number of such students who have put their theses at risk due to poor writing skills. On the other hand, a limited idiomatic vocabulary will not prevent you from communicating your contributions. After all, you don't need to be a Hemingway to write a thesis. You need proficiency in a *craft*.

Here, in only a page or two, I cannot even begin to address the language skills that you began to acquire as an infant. However, there are several things that you can do to improve your mastery of the craft.

First, read about writing. Go to the library, and check out a handful from the thousands of books about writing. Read a little bit every day, picking out pearls of advice. I recommend these books[6]:

> Gerald Alred, Charles Brusaw & Walter Oliu, "Handbook of Technical Writing", 6th Edition, St. Martin's Press, 2000.
>
> Robert Ian Scott, "The Specific Writer", Intl Society for General Semantics, 1995.
>
> Rudolph Flesch, "The Art of Plain Talk", Collier Books, New York, 1951.

Second, just do it, pretending that you are indeed writing for the people in your field. Don't just sit and worry. Practice to get better. Don't just sit waiting for the muse to strike. Work, because thesis writing is craft, not art. Don't just sit idly, waiting to begin some massive final push. Practice writing daily just as you would if training for a marathon. Don't just write something once and then move on to something else. Write the same content several different ways and then critique their strengths and weaknesses. Write the same content for different audiences and likewise critique.

Third, practice basics. Whatever your current level of skill, act like a baseball player going to spring training. Review fundamentals until you are sharp and smooth. Use repetition to inculcate those fundamentals. Slow down and look closely at small things. Become sensitive to the nuances of language, selecting exactly the right words to express precisely what you intend. Doing so can make you aware of how imprecise your current phrasing is. (That happens to me all the time.)

What basics? Here are two crucial ones: sentence and paragraph. Many students struggle to express themselves, even at this small scale. (So do professors — writing never does get easy.) Learn to be patient in exposing the audience to your ideas. Don't have words gush forth with no designed flow. Become good at using glue words, connectors that allow ideas to flow within and between paragraphs. For example, in the paragraph two before this one, look at how the controlled repetition of "just" glues things together.

Here is good news: if you can take PTSs, and craft paragraphs of sentences, then you can write a thesis!

Here are several more, even smaller scale basics: verb, voice, and person. Write with powerful verbs in active voice. Don't turn a verb into a noun and then accompany it with a weak verb. For example, don't replace "Julia transferred simply" with "Julia's transference had no complications". Avoid passive voice, because with it you more likely will produce vague statements. Use active voice instead. (You may have to fight with your committee on this one. So prepare yourself to use this argument: passive voice is not more authoritative, it just makes readers work harder.) Don't be afraid of writing in first person; "I" is not forbidden. I use "I" when I need to, but that

---

[6] I also recommend these books: a) Joseph Williams, "Style: Ten Lessons in Clarity and Grace", 6th Edition, Addison Wesley, 1999; b) Edward MacNeal, "Mathsemantics: Making Numbers Talk Sense", Penguin, 1994; c) Barbara Tuchman, "Practicing History", Ballantine, 1981; and d) D.A. Fryxell, "How to Write Fast (While Writing Well)", Writer's Digest Books, 1992. (He likes McPhee's way of work too.)

seldom happens because I write for "them", my audience. (So, dissect — see below for more on dissect — this paper focusing on person.)

To help you avoid muddy writing, pay attention to these pairs of "level" words: type or instance; singular or plural; active or passive; set or sequence; and general or specific. Level distinctions form a significant share of the skills I teach in my books.[7] The books are about computer systems people learning how to precisely communicate with users about what they want to remember in a database and what analyses they want to perform. The books are valuable because the limiting factor in database work is not technology, it is words. I think that good level skills will help you with writing your thesis.

Fourth, learn templates to mimic. To avoid spewing words, you need to own tools, organizing principles, so that you can craft paragraphs, sections, and chapters. The books by Alred et al, by Scott, and by Flesch provide some valuable ones. Alred, et al present a number of methods of development: order of importance, cause and effect, sequential, chronological, increasing importance, decreasing importance, division and classification, comparison, spatial, specific to general, and general to specific. They present four forms of discourse: exposition, description, persuasion, and narration. Scott suggests that you see writing as connected sets of questions and answers. He also has chapters about organizing: within paragraphs, sequences, comparisons, arguments, inferences, and figures of speech. Flesch invented a "readability formula" that some word processors implement. It's a rough measure. He also explains live, crowded, and empty words. Get these books, put their notions in your toolbox, and use them.

Fifth, read good writing, and dissect it. You can become a better writer by looking closely at the structure of something you enjoyed reading. For example, read Craig Partridge's "Gigabit Networking" (Addison-Wesley, 1993). While its content interests me and my students, you can focus on its design, selecting a chapter and reading just the paragraph topic sentences (each begins a paragraph). Those PTSs tell the chapter's story, but, of course, in less detail. Each paragraph fleshes out the topic sentence, and nothing else.

Read a bunch of theses. Find well-written ones, not considering their relevance to your field. Dissect the writing. Find poorly written theses, dissect them, and avoid the errors you discover. Find theses in your field, dissecting to get a sense for what kind of things committees have accepted. However, do not feel bound by what others have done.

While you cannot determine the writer's way of work that produced a well-written piece by dissecting it, there's hope — read the next section.

[7] John Carlis and Joseph Maguire, "Mastering Data Modeling", Addison-Wesley, 2001.
John Carlis with Scott Krieger, Mastering Database Analysis", Addison-Wesley, 2002 (in press).

5. Follow The "Design a Thesis" way of work. *Don't Hack*

Here are the parts of the "Design a Thesis" way of work.

- Remember the audience (you or them), while you write three products (contributions, design, draft).
- Work backwards from contributions.
- Design a story, considering many exposition strategies.
- Finish the design before writing the draft. That is, *don't hack*.
- Draft one paragraph at a time.
- Edit in the small, not in the large.

Missing from this way of work is the bulk of the research, where you work hard so that you have something to write about.

Audience and products. Remember, you always write using words, but the audience varies by product. The table below summarizes the step-audience pairing. Yes, of course, your polished contributions and your design will appear in the thesis, but, when you craft them before writing the draft, you are the audience. With a statement of contributions you say to yourself "This is what I intend to contribute". With a design you say to yourself "This is my plan of how I intend to inform them (my committee) of my contributions and my evidence for those contributions". With a draft you say to them say "Look, read this. I have created something new and valuable, that somebody cares about, and I can prove it".

		Audience
Products	Contributions	*You*
	Design	*You*
	Draft	*Them*

Work backwards from contributions. Working backwards from contributions means three things. First, you articulate your contributions early for yourself. Second, you keep this contribution destination in mind, so you can more readily make design decisions, deciding what to include or exclude (bound the thesis), and deciding how to present your material. Third, you do not work forwards from what you could write. Remember, it's not about you; it's about them.

Design a story. When you design your thesis for yourself imagine that in the draft you will be telling a story to them. The point of the story is your set of contributions. The context for the story is the state of your field without your work. To tell the story you need to work on two pieces: the vocabulary, and a story tree.

Vocabulary. Vocabulary is an important structural component of your thesis. Because you live 24/7 with your thesis, you can easily forget that they do not, and then be surprised when they don't use vocabulary as you do, or don't know what your vocabulary means. Therefore, you should craft precise definitions for them early in design,

and feature them prominently in your thesis.[8] One kind of vocabulary is the content terms that come from the field, or that you invent. In my database work I teach people two phrases to help them precisely define vocabulary: "What do we mean by one of these?", and "Anchor understanding with instances". They might help you too.

A second kind of vocabulary is glue, the words, sentences, and paragraphs that make your writing flow for the reader. Poor glue words cause trouble. For example, in works I read, a writer often will use glue words such as method, methodology, process, and so on. Some writers intend them to be synonyms, and use the different words for variety. Others mean them to be different things, but don't explicitly say so. My advice: a) don't use synonyms for variety, pick one and stick with it, because the repetition tells the reader that you are referring to the same notion again; and b) explicitly define terms, because otherwise you risk them misunderstanding you, and thinking less of your contributions. Also, by paying attention to vocabulary you can minimize "drift". In something as big as a thesis, written over several months (years☹), style and vocabulary drift will naturally occur. You need to take care to avoid it. Finally, you should consider the possible advantages accruing to these who generally avoid the overuse of limiting, qualifying, or similarly self-protective words. That is, don't use weasel words.

It behooves you to work hard at vocabulary, because otherwise you risk readers thinking less of your contributions.

Story tree. The heart of your design is a story tree, a detailed outline depicting the pieces and subpieces of your thesis, the order in which they appear, and the smaller signposts, the glue that guides the reader along the way.

Why must you concern yourself with designing a story tree? Well, your thesis undoubtedly contains an intricately connected network of topics, but readers read linearly. Somehow you must cover the network, and you necessarily will have some redundancy (but, of course, not too much). With the design you specify the order in which the reader will see the topics, and how much content you intend to present each time a topic is touched.

A powerful way that can help the reader perceive the network is the liberal use of figures and tables.[9] For example, before you begin to design your literature review build a two dimensional table with, say, each row about a paper (or subfield), each column about a property of the research in the literature, and each cell about what is memorable in a paper about the property. Besides judiciously selecting the papers and properties, you need to decide the order of rows and columns, and the possible content of a cell. In the Appendix Figures A1 – A3 show literature review tables by Libby Shoop, Elizabeth Stuck, and Leone Barnett. (Figure A-2 has a final row depicting Stuck's work — a good way to highlight her dominating the competition.)

Three other large-scale figures that might be useful for you are: the field before and after your thesis; a system diagram depicting what you built, and a figure showing the intricate associations among concepts. (See Figures A4 – A6.) Read other theses looking for useful figures and tables that you can mimic.

When you have such a table or figure good things happen. You will readily see what you need to write about and how to organize it. You will see more to write about than you would have otherwise. The readers will think more highly of your review than they would have otherwise. Spend time polishing your figures and tables; it is well worth the effort. Figures A7 and A8 show an intermediate and a polished version of a figure from a research paper. The latter more clearly distinguishes between active (oval) and passive things (box). It also has less ink[10] and is easier to follow because related things appearing closer to each other allows shorter, no-bend lines.

When you put a figure or table in your thesis, be sure to write about it, that is, make the figure part of your story. For each one, finish this statement: "Here is what you should get out of this figure …". Don't depend on the reader working hard to determine what is important. It is your job to do that work. As you try to tell a figure's story, you probably will find ways to improve both the figure and the story. That is what happened with Figures A7 and A8.

---

[8] Voltaire supposedly said "before you talk to me, define your terms". ("Supposedly" means I don't have the reference.)

[9] R. Harris, "Information Graphics: A Comprehensive Illustrated Reference", Management Graphics, 1996. (A valuable source.)

[10] E. Tufte, "The Visual Display of Quantitative Information", Graphics Press, 1983. (Another valuable source.)

By the way, fields vary in where they tell a figure's story. Some, by convention, explain a lot in captions, while others don't, and, instead, tell the story in the body of text. Either way works fine.

What does a story tree look like? On the left below is a partially delineated skeleton of a story tree with a maximum depth of 4, and four PTSs. In a real one you would put words with each oval, and where each "PTS" appears. While you surely would use at least a coarse outline anyway, remember that this story tree, when complete, is detailed down to PTSs.

A story tree has another component: glue. On the right below is the same skeleton augmented with a few possible places for glue: I for introduction, T for transition, S for summary, and C for crossing. A crossing shows where, elsewhere in the tree, it crosses to.) You won't write glue everywhere, but you should explicitly decide, and not put glue where it happens to feel appropriate as you write.

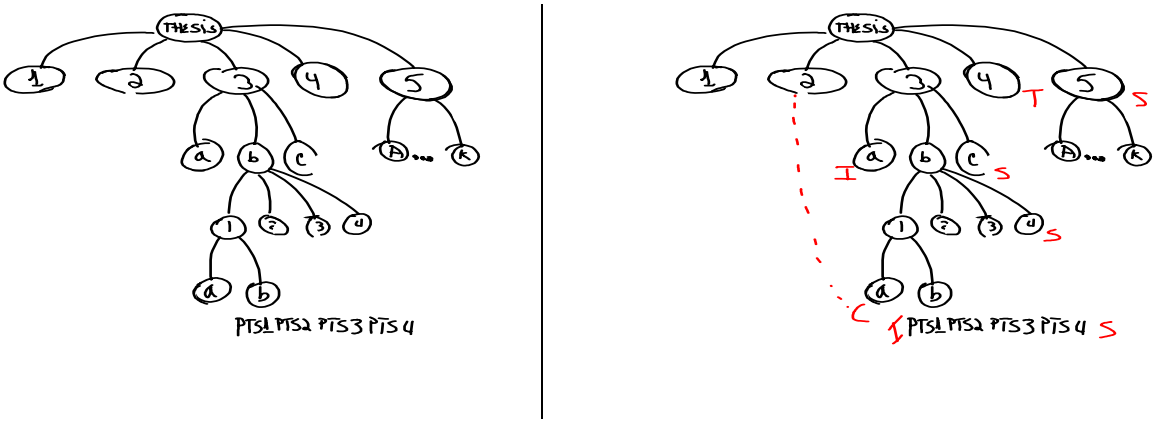

To help you judge the goodness of a story tree, use two notions that come from software: coupling and cohesion. Coupling measures the degree of interdependence between two pieces. Higher coupling means that if you change one piece then you more likely will have to change the second piece. Coupling is the degree of relatedness of the notions within a piece. A piece with high cohesion is about one thing. (Look up "unity" in a writing handbook.)

Building a story tree helps you discipline yourself. You can't just say anything anywhere. Instead, you strive for high cohesion and low coupling. The trick is to a) make each piece inward looking, that is, independent (or nearly so) of the other pieces, and b) limit connections between pieces, especially Crossings. When thinking about glue, look at how two story tree pieces are related. They can have a sibling, a parent/child, or a cousin relationship. (The sections of a chapter are siblings; a section of a chapter have a child/parent relationship; a paragraph in section 3.b.1.b has a cousin relationship with anything in chapter 2.) You should expect to write I, T, or S glue showing how siblings or a parent-child pair connect. However, limit cousin connections — these are Crossings. You will need crossings, but don't put them among siblings; put them before or after.

How should you go about designing a story tree? Well, design is messy, creative, and iterative, not top down, bottom up, or some other mechanical process. Don't worry, there are many paths to success, and, (I am serious about this) since you are an unusually smart person, you will find a way.

While there is no one best way to design, here are some things to try. First, experiment with organizing principles. For example, for Figure A3 Libby Shoop took the readers on a tour of the figure, going from the upper left, down, right, and up. Instead she could have reversed that order, or started in the middle and moved out from there. Second, look at siblings and decide if they are a set or a sequence of topics. For a sequence, usually you will choose to first to last, or the reverse. For a set of siblings, any order will due, but you may order them most to least important, or the reverse. You might also choose an order that allows you to define terms first. Third, think about the balance of the tree. If one branch is large relative to others there may be no problem. However, you may decide to prune the large branch, and put its detail in an appendix. Fourth, to help you decide, keep the reader in mind and refer to your contributions. Fifth, after you design, put the story down, let it cool off, and come back to it with fresh eyes.

Finish the design. Here is a crucial rule:

- **Don't write the draft until your design is done.**

Said a second way:

- **Design the hard parts before you write the easy parts.**

Said a third way:

- ***Don't hack*.**

"Hack" is an interesting word. A hack is a) a person hired to do routine, often dull writing, or b) a bad thing (a hack - a failure), or c) a good thing but only in the short term (e.g., a hack job - duct taping a plastic sheet over a car's broken window). When you hack you cobble together a quick, temporary fix to a small problem. This rule adheres to a long held principle for software programmers: "design, design, design, then code once". In other words: *don't hack*.[11]

Following this rule allows you to avoid the waste that comes from *false progress*. Many students, even those who know better, violate the rule in order to show "progress". (This happens a lot in software development too.) However, real progress does not occur if you write easy paragraphs, and then either discard them when they do not fit, or keep them when they do not truly fit in the next draft. What does occur is either wasted effort, or diminished thesis quality. Since either outcome hurts, you should be patient, and not opt for false progress. If a paragraph is easy to write now, it will be easy to write later. It doesn't hurt if you were patient, and the PTSs, for the paragraphs that you didn't write, get moved, changed, or discarded.

You will be willing to be patient if you can "see". See what? See the flesh that will be on the bones of your design, that is, the paragraph that delineates the idea of each PTS. What enables you to see? Mastery of your craft does.

Where does your advisor fit in this way of work? In the figure below, your advisor (hatted) is privy to your sketches, and can (and should) give you feedback much earlier than the draft. Ideally your advisor approves the story tree, and then accepts the draft you write from it. For that to happen the advisor also must be able to see the flesh that you will put on the bones. [An aside to advisors: I work this way, and get a better draft because I was in on the design. Try it, you'll like it.]

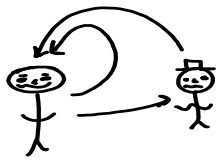

Draft one paragraph at a time. If you have a quality design, with high cohesion and low coupling, then you can, and should, write one paragraph at a time. Doing so has several benefits. It allows you to focus on just one thing, and therefore makes you efficient. It raises your gumption level, because, while writing a *thesis* is a daunting task, writing a *paragraph* is do-able by a regular person. It keeps you from babbling, because all you do is develop the PTS. It serves as an alarm mechanism, because, if you are thinking about parts of the thesis other than the paragraph you are currently writing (plus, perhaps, its nearby paragraphs), then either you need a rest, or you have a flawed design that you need to fix.

[11] See John Carlis, "Wizards See and Dance, While Hackers Grope and Grunt", National Teaching and Learning Forum, April, 2000, (www.ntlf.com). Also see Patrick Starr and John Carlis, "Introducing the Design Process to Beginners: The Spiral Model", ASEE 2000, Learning and Teaching in the 21st Century, September, 2000. (www.cs.umn.edu/~carlis)

In what order should you write the paragraphs? Good news: many different orders will get the job done —after all you created a quality design. I suggest that you write term-defining paragraphs before any that use the terms. Write a section all at once, rather than skipping around. Write paragraphs with relatively uncertain terms or PTSs early, so, if needed, you can revise the design. (Don't let this happen often.) Consider writing the meat of a section before writing the transition paragraphs. Write introductory paragraphs late.

By the way, when you write the thesis' abstract, don't make the mistake of writing a table of contents. The abstract is for people who may not read the rest of the thesis. It should be specific, stating the problem, your approach, and your contributions. A good abstract can motivate potential readers to dive into the thesis, because the impact of your contributions excites them.

Edit. When you produce the draft you will edit. Here is really terrific news: because you designed before drafting, most of your editing will be in the small, that is, within a paragraph, and not in the large, that is, you will not need to make structural changes. So your editing is merely polishing, not rebuilding.

What can go wrong and what you can do about it.

What if you are stuck, and don't know how to tell a portion of your story? Here are several things you could try. Draw a picture or several. Study your vocabulary; perhaps it is vague or incomplete. Talk it out with your advisor or a friend. As I talk with one of my students when she or he is struggling, I often will hear a gem. When that happens, I say "don't talk anymore until you write that down". From gems a better story emerges.

What if you just hacked and wrote an unacceptable product? Instead of hacking again, back up and design, finding a story.

What if, when you write the draft, you find that your PTSs are not really PTSs? You might find yourself writing several paragraphs for a nominal PTS. If so you may just have been lazy, and did not finish your design. You might have PTSs not leading to paragraphs at all. If so you may need some remedial training to learn basic paragraph craft. Remember, make paragraph your unit of thought.

What if your advisor approves your design but then wants major changes to the draft anyway? (To avoid getting into this pickle, be sure that you and your advisor share expectations.) Sorry, I don't have a magic pill for this problem. The good news is that a quality design will limit the effort to make those changes.

Conclusion.

When you are done, you must defend the thesis. However, you don't have to get defensive. You can turn your defense into a celebration, welcoming you to the community of scholars. How? Have good contributions, a well-written thesis, and a thesis talk that focuses on contributions.

Finally, "joy" is important. Completing a Ph.D. is a long haul, and you need joy to sustain you. If you work on a topic that you have an abiding interest in, and follow the "Design a Thesis" way of work, then you can achieve success with joy.