

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

John V. Carlis

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

draft number 2.11 ©

October, 2005

Introduction.

My goal in this paper is to codify an effective and efficient way to design and then write a thesis in one draft – yes, really *one* draft – so that I more readily could pass it on to my own students, and share it with other students and their advisors.¹ If you are a Ph.D student who needs to write a thesis or a thesis proposal then this paper ought to interest you, and your advisor too. The design advice presented here also applies to you if you merely (hah) need to write a research paper. To get going here are some questions for you to ponder.

What does “draft” mean? Here 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 alone, 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, producing a thesis is work, but it is merely work, not some mystical thing.

What makes me qualified to give advice about a one-draft thesis? Experience does. I partially did it. Four of my PhD advisees, Leone Barnett, Libby Shoop, Paul Wagner, and Yue (Jake) Chen, have done it. (They, of course, agree with everything I say. ☺)

What’s my story? While writing my thesis and, at the same time, working as a professor (yes, it was stressful), several streams of thought happily converged. I read lots about writing, and was particularly struck by the forward to the John McPhee Reader, which described his disciplined, design-before-write way of work². I read Richard Mager’s Preparing Instructional Objectives, which has this message that translates to writing: teaching (writing) is about them, the students (readers), not the teacher (author). 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. By then I already believed that software should be designed and not hacked, and, coming to the same conclusion about a thesis, I decided to try to transfer software notions to writing. Since a paragraph is a unit of development, like a software procedure, I, following McPhee, chose to design everything 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.

What can motivate you to follow my advice? Five 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. Your thesis most likely is, as it is for most students, the largest writing project

¹ I was telling some “Bushies,” mentees of mine in the University of Minnesota’s Early Career Faculty Development Program (funded initially by the Bush Foundation), about my one-draft notions when Patricia Schaber said “Come talk to my department about it.” That invitation motivated my writing this paper, and to my leading a number of workshops. So, thanks Patricia.

² McPhee won the Pulitzer Prize for “Annals of a Former World,” Farrar Straus & Giroux, 1998. Besides writing great books, he teaches non-fiction writing at Princeton. Read his books.

you have undertaken, or perhaps will ever undertake, and you do not have the experience to say “no problem.” Second, you may find writing painful; many students (and professors) do. Sadly, after years of hard work on nifty ideas, they struggle to produce a barely-acceptable thesis. So you ought to be willing to listen if doing so will ease your pain. Third, *ceteris paribus*, your committee and the field at large will consider your thesis as having higher quality substance if it is better written. So you will 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 Ph.D., no faculty job. Fifth, your future success in any job is impelled, or impeded, by your communication skills, or lack of them. So it behooves you to improve your writing -- and your speaking too.

If, after you have striven mightily and have something to write about, you follow my design way of work, can you produce a one-draft Ph.D. thesis? No matter what your field, my answer is emphatically “yes,” and this happy thought ought to make you willing to read on. Begin by studying the summary figure below. It briefly states what you need to do to obtain that success, and the five key notions you should understand and employ. It also delineates the organization of the paper. (An addendum for advisors follows.)

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

To Succeed:		<i>Key Notion</i>
1. Believe	In a one-draft thesis.	<i>Gumption</i>
2. Contribute	To the field.	<i>Destination</i>
3. Avoid	The fatal, wrong audience error	<i>Sketch</i>
4. Acquire	Writing skills.	<i>Paragraph</i>
5. Design	A thesis.	<i>Avoid False Progress</i>

What’s the catch? Well, you cannot just casually read this paper, and then, poof, you make a thesis appear. If you are going to succeed you will have to think hard about writing ideas that may be new to you. You will need to practice. You may need to break some old, bad habits. You will need to become disciplined, because just winging it doesn’t work for a thesis. Educating your advisor may be your hardest task. However, the benefits accruing to you will far outweigh the costs.

Does this way of work apply to course assignment or other kinds of papers? Yes! Change these terms: (thesis, field, and contributions) into these: (paper, audience, and messages).

1. Believe In a one-draft thesis.

Gumption

In order for you to be willing to try to write a thesis in one draft, you must believe it can be done at all, and, more to the point, that *you* can do it. To acquire that faith consider the following questions.

Can it be done? Absolutely, since four of my students did it.

What specifically should you believe? Three things. First, believe that re-drafting is both costly and avoidable. By writing only one draft, you can save at least several months worth of effort. Second, believe that one draft is preferable to several drafts. It indicates good organization that yields a more readable thesis, and does not indicate of lack of effort – beware, you may have to fight your own or your advisor’s preconceptions about this point. Hopefully, you will come to consider re-drafting to be a consequence of an incomplete or flawed design. Third, and most important, believe that indeed *you* can write a one-draft thesis. Become convinced you that if you can write a paragraph — and surely you can or can learn to do so — then you can write a one-draft thesis, one paragraph at a time.

What benefits will accrue to you if you do believe? It will raise your gumption level³, because when you understand what is going on and why, then you will have more gumption. With low gumption you can worry and perhaps do no writing — those with lots of gumption do not suffer much from writer's block — or do wasteful things in the forlorn hope that good things will occur. Also, the way of work you will learn will help you with other writing tasks, such as writing research papers and proposals for funding.

Did you have doubts about my message because this is not the first draft of my one-draft missive? (I do not trust people who do not take their own advice.) Well, once your thesis gets approved it gets archived and does not change, although you may extract portions of it for papers or massage it into a book. On the other hand, I wrote a tutorial, not a thesis, and have extended my original version by about 25% based on feedback and requests from workshop students and advisors. However, since I did plan for success, the additions consisted largely of pain-free fleshing out of my original design.

³ See R. Pirsig’s “Zen and the Art of Motorcycle Maintenance,” Morrow, 1974. It’s a great book; philosophy with a plot and characters, and relates to this paper and your thesis. Read it.

2. Contribute To the field.

Destination

You need to understand that the purpose of a thesis is to defend your claimed contributions to your field. In other, more dramatic words: *No contributions, no thesis!*

What constitutes a contribution? Well, in spite of its importance, many students, and some advisors, misunderstand the meaning of “contribution.” To your committee, who represent the other⁴ experts in the field, you will say “Look, here is something new and valuable, that these people care about, and I can prove it is good.” Your immediate 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 new, valid and significant. (And maybe even useful. ☺) 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. To help you grok⁵ contribution, next is more detail about three important terms: novelty, consequence, and evidence.

Novelty. 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 used X in a subtly different fashion” You should expect to invent new vocabulary, and then precisely define it – every thesis makes fine distinctions among terms. You demonstrate expertise by picking an appropriate vocabulary to focus on, and organizing it. The literature review then serves as the substrate upon which you show the goodness of your additions. In your proposal a literature review helps you describe your potential contributions. (More to come on this topic in section 5.)

What if your intended collection of contributions defines a new field or connects two or more existing fields? Well, most of the research ferment is in the cracks or chasms between fields, so such a situation is not only not bad, it is terrific news, because then you are more likely to claim more significant consequences. However, it will make your literature review dicier to write.

Consequence. To persuade your committee about consequence, you state who cares about your novel contributions, and what impact your contributions have on them. Of course *you*

⁴ By the time you defend your thesis you must be an expert. So act the part.

⁵ Grok was coined by Robert Heinlein in “Stranger in a Strange Land,” Ace Books, 1961. It means “deeply understand”.

care, but that is immaterial. If you produce something new, but it impacts nobody, 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 all this work," because the committee's reaction will be "So what?" One student I met wrote a 50,000 line program for his thesis, certainly a massive effort, but it seemed more like contract programming than research. So, unlike that student, be sure that you remember to distinguish between means and ends, and that it's the ends that get judged as contributions, whereas means are a way to get evidence — unless you have invented novel means, in which case they become contributions. Also do not stop at saying "I learned this," because *your* learning something is not a contribution to the field if the field already knows what you just learned — we call that education.

Instead, say things like "My explanatory model for ... is better by ...," "My new method yields 50% more ... than the best previous method," or "My new operators enable users to state ..., which, previously, they were unable to state." In each case you state how your work improves the state of the field.

Evidence. To persuade your committee to accept your claims about the impact of your novel contributions, you need to provide evidence. If you invent or discover 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, strengths and limits of your contributions.

Every thesis has limits. No, admitting to a limitation does not mean that you do not graduate until you remove it. Do not try to hide limits; embrace them because your text about them is often very interesting to your readers — and to future students who will extend your work, saying something like "XXX stopped here, but my work goes beyond."

Using contributions. How should you use contributions? From beginning to end, keep the contributions in the forefront of your thinking. Keep a copy in front of your eyes -- literally. (Paste them on the wall.) 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 hear a clear explanation from you and then smile on your work.

Starting. When should you think about contributions? Well, here is a major mistake:

- You first write the bulk of the thesis, and then try to determine what contributions you can claim.

That's the tail wagging the dog, and will cause you great trouble. Instead, you should start thinking about contributions much earlier than draft-writing time. In your proposal you must have target contributions. If you do not then your committee will balk. (If your department does not require a proposal, write one anyway and ask your committee for feedback.) In pre-proposal pondering you should try out possible contributions, perhaps

with a five-or-so-page white paper. Even earlier, you can imagine a potential “BIG WOW!,” 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 redirect your work. Also, at each point you can assess risk: if something goes wrong, will you still graduate? Earlier, you can begin by acquiring expertise in the field, studying what has been done, and practice uttering informed opinions. Hopefully, you get ideas (lots of theses start with “Oooh, what if I changed ...,” or “Nobody has ...), or even get mad (others germinate from “What a mess!,” or “I don’t like ...”).

Where do contributions come from? Your abiding interest (passion!) in a topic and hard work enable you to have something interesting and verifiable to say to others. Do not expect your advisor to assign you a topic for your thesis or to tell you what to claim.

In short, you think about potential contributions from the very beginning as you acquire expertise and interest. As you proceed you change the phrasing from “What could I perhaps contribute” while musing to “I intend to contribute” in the proposal to “I have contributed” in the thesis.

Contributions are not an afterthought, but are the *raison d’etre* of your thesis. So you should formulate your contributions early because when you know what contributions you intend to defend you can more easily do four important things:

- *Know 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, and, happily, not after.
- *Protect yourself from your advisor.* A student/faculty candidate, interviewing one 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 goofiness happen to you — schedule doesn’t matter, substance does.
- *Decide what and in what depth you must write about, and what you may skip.* Of course, you will have lots to say, but, 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. 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.” Do you think he wasted his time? I do.
- *Decide what other work you need to do, and not do.* Suppose you want to claim a particular contribution, but cannot prove it yet. You must drop the claim, or weaken it, or do more work to gather evidence, or move it to a future work section. Do not spend effort on something just because you *could* do it. For example, a student did irrelevant work, spending 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. He was poorly advised.

Destination. Think of contributions as the destination. Research explores new territory (or re-envision old 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 away from your destination. You wander if you say to yourself "this is a thesis day," but get distracted and work on notions that do not advance your contributions.

Suppose you are driving, and come to a fork in the road. If you don't know where you are going then you cannot sensibly decide whether to turn left or right. A thesis writer without contributions faces the same dilemma.

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 will 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 your research making a contribution to others. 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. So you can't tell when you will be done. Instead reply with "Look, here is something new and valuable about XXX, that these folks in the field care about, and I can (or will) prove it with this evidence."⁶

Use this notion of destination when you are in pre-proposal or proposal stage. One of my students (Paul Wagner) wanted to work on a particular topic but he could not find an appropriate dataset on which to test his notions. Because he knew early that he would not be able to get the needed evidence to support his intended contributions, he wisely modified them, saving himself (and me) a lot of pain.

Use destination to help you think about what needs to be done (by you or the field) as follow on work. Your thesis destination is not the end of the road. Every thesis opens up new vistas for exploration.

Distance. So, since the road never ends but your thesis must, an important question is "How much of an advance of the field is enough for a thesis?" Here too this knowledge is part of the expertise that you acquire. Stating your intended endpoint is not some easy pronouncement, but is something you work out with your advisor and committee (and with reviewers when you try to publish or get funding). Focus on articulating your contributions comparing them to the field as a way of calibrating how far you intend to go.

There are no firm, universal rules about just how much of a contribution is needed for a thesis. Disciplines differ, as do advisors, but here are three things to not do. Don't just 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 think

⁶ One workshop attendee said that for her this was the most valuable paragraph in the paper. So you may want to re-read it.


in terms of moving out from where you are or how long you have been moving, which puts the focus on you, and not on the field where it belongs. Don't merely propose to make three, smaller but published, weakly- or un- connected contributions, and claim that they make a thesis. A thesis is not a grab bag, so avoid this mistake. However, in some fields or with some advisors the thesis is just a slapping together of three papers, co-authored with the advisor, and not an integrated work. Be sure that you ask so that you avoid misconceptions.

If you are thesising (Yikes, is this an acceptable verb?) but have not read any theses, then you are driving blind, and are likely to wander. To scout out the structures in the territory, read several theses – not just one – in your field, by graduates of your department, and by advisees of your advisor. Read some closely, especially if the content will be in your literature review. Skim⁷ several others to get a sense for the kind and amount of contribution that students have done and faculty have smiled on. Why read theses? Well, if I told you to write a sonata and you had never heard one, of course you would seek examples so that you could acquire a sense of sonata-ness before you compose.

⁷ Advice: allow yourself at most two hours to skim a thesis; read the front, back, & table of contents; look for a succinct statement of contributions. If you do this several times you will start to grok the notion of thesis.

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:



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 craft and polish a well-designed product for them to read. However, such effort is a waste, a fatal error, if *you* are the audience, as is the case early on for sketches and later for blueprints.

Sketch. To help you avoid the fatal error, think of early 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.
- Discarding sketches — keep the ideas, not the medium on which you jotted them down.

Why sketch? Well, when you have only a vague idea of your destination (aka contributions), you can 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. Do not make the mistake of thinking that sketching is a sign of weakness — Leonardo Da Vinci sketched. Do not avoid sketching in order to save time -- you will cost yourself time because you will discard what you pretended was a final product.

Blueprint. Sketching is important, but because a thesis is big, you need to design a “blueprint” for the thesis — still writing for yourself, but planning for the audience’s reading. Think of it as coming up with a house’s blueprint, which is a more formal design than a back-of-the-envelope sketch, but is still clearly not a house. A house’s 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.

Why is ignoring this audience distinction fatal? Well, you care about cost, and you can change a sketch or a blueprint or model cheaply (in money, time, effort, or emotion), but it will cost you dearly to change the house. Moving a window on a sketch costs pennies, perhaps dollars on a blueprint, but on a house perhaps thousands of dollars. If you try to write without a design you will end up rewriting, tossing away expensive words — so don't do it. Happily, when you work from a blueprint, you can concentrate, focusing your attention on building one small piece, and ignoring almost everything else. Likewise, redrafting is expensive, but modifying your design is relatively cheap. Furthermore, the distinction's importance grows nonlinearly with size (a thesis is indeed big), and with your inexperience.

Duh! Why do folks fatally err? Well, of course nobody confuses the cardboard model or the blueprint of a house with the house. Everybody agrees that it is a waste of time (foolish) to make the model's toilet flush. However, many people forget, when writing, that designing is not the same as drafting — because both activities use the same material, namely, words. Because of this error I do not use the term "writing to learn" as Zinsser⁸ does, preferring "sketching to learn" instead. My advice then is: don't fool yourself, and commit the fatal error of confusing a sketch or the blueprint (design) with the house (draft). (Expect more on this topic in a bit.)

You would not build a house without sketches and then a blueprint, so do not write a thesis without a design. But how detailed a design? The next two sections address that question.

⁸ William Zinsser, "Writing to Learn," Harper & Row, 1988.

4. Acquire Draft-writing skills.

Paragraph

Are you already a good writer? If you are, well that's terrific, but you should expect that writing your thesis will challenge you anyway. If you now are a poor writer, then you better exert yourself to get better, and do so right away. Don't wait until you are partway through a bad draft to begin remedial training. Beware, some students are unaware of how poor their writing skills are, while others think it is no big deal⁹ — it is — and still others excuse themselves by saying that they didn't get the writing gene. Don't whine, get to work and nurture the skills that you need.

By the way, do not be smug if English is your native language. I have seen plenty 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 — Jake Chen, who was relatively new to English, wrote a fine, one-draft thesis. After all, you don't need to be a Hemingway to write a thesis. You need proficiency in a *craft*.

Paragraph. The core of that craft is the notion of a paragraph. The good news here, and it is really good news, is: if you can take a paragraph topic sentence, PTS¹⁰, and craft a paragraph of sentences, then you can write a thesis! The potentially-bad news for you is that many students struggle to write a paragraph, and few think explicitly in terms of paragraphs.

To succeed you need to be able to:

- Write a good a PTS.
- Develop one PTS into one paragraph.
- Connect paragraphs together to tell a story to the audience.
- Recognize structural errors: poor or missing PTS, and a mismatch between a PTS and its paragraph, especially thoughts that do not support the PTS and belong elsewhere.
- Use paragraph as a unit of thought.

The last bullet is, admittedly, vague, but you can acquire this skill. You can become better at it just by trying, without any particular training. (Expect more on this in the next section.)

⁹ A student said to his committee "I have good ideas, I don't need to express them well." Not surprisingly, his committee frowned on his thesis and sent him off to re-write.

¹⁰ 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.

Furthermore, 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 your ideas to flow to the audience within and between paragraphs. For example, in the paragraph below that begins with "Third," look at how the controlled repetition of "just" glues things together.

Of course, you need more than paragraphs, but the few pages here cannot even begin to address the language skills that you have been acquiring since infancy. However, here are several prime actions that you can take to improve your mastery of the writing 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. Try these books¹¹:

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. [out of print]

Rudolph Flesch, "The Art of Plain Talk," Collier Books, New York, 1951. [out of print]

George Gopen, "Expectations; Teaching Writing from the Reader's Perspective," Perason Education, 2004, and "The Sense of Structure, Writing from the Reader's Perspective," Pearson Education, 2004.

Among other things, these books provide some valuable larger scale templates¹² to mimic. These templates help you avoid spewing words, giving you organizing principles, so that you can craft paragraphs, sections and chapters. 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. (I really like these books but wish they had explicitly included ways to expose a complex network of thoughts, which is what you must do in your thesis.) Flesch invented a "readability formula" that some word processors implement. It's a rough measure, but with it you can gauge how difficult others will find your writing. He also explains live, crowded, and

¹¹ And 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 (Tuchman won a Pulitzer Prize for "Guns of August," but suffered terribly in her writing process.) ; 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.)

¹² These templates strongly parallel software/database templates (or shapes or patterns).

empty words. Gopen sees problems with the Flesch quotient. (I do too, but enjoying saying "Flesch quotient.") Gopen distinguishes between a paragraph's topic and its point, and delineates richer structural possibilities than one PTS per paragraph. Gopen emphasizes that readers expect certain things to appear in certain places. So where you place words dramatically affects how readers are likely to perceive them.¹³ Get these books, put their notions in your writing craft toolbox, and use them.

Second, practice smaller-scale basics. Whatever your current level of skill, act like a baseball player during spring training. Practice fundamentals until you are sharp, smooth and habitual – you need 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 in just the right order to convey precisely what you intend. Doing so can make you aware of how incomplete and imprecise your current thinking and phrasing is. (It still happens to me all the time.) Finally, exercise judgment; do not blindly follow rigid writing rules, because good writing is about readers receiving what you intended to convey.

Of course you must have command over sentences, including even smaller scale basics: verb, voice, and person. Write sentences 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 "the transference of Julia had no complications." Avoid passive voice when you can, 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, often 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 seldom happens because I write not about me but about the topic for the audience. (So, dissect — see below for more on dissect — this paper focusing on person. I did so myself and edited out some gratuitous "I"s.)

Try this: take some finished writing of yours and highlight the verbs. If you see lots of "is" or other forms of "to be," or "has," then others will perceive your writing as weak.

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. For example, when a few paragraphs back you read "students who have put their theses at risk," you may have noticed the imprecision – there is no thesis done by a group, and only very rarely does one student have theses (plural). Level distinctions form a significant share of the skills I teach in my books.¹⁴ 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

¹³ For example, shortly before writing this paragraph I heard a radio DJ say "The first one who correctly calls in and identifies the artist of this song will win ..."

¹⁴ John Carlis and Joseph Maguire, "Mastering Data Modeling," Addison-Wesley, 2001.
John Carlis with Scott Krieger, "Mastering Database Analysis," Addison-Wesley, 2006 (in press).

limiting factor in database work is not technology, it is communication.¹⁵ Good level skills will help you with writing your thesis.

Third, just do it. Pretend 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 much as you would if training for a marathon. Don't just write something once and then move on to something else. Write, or at least design, the same content several different ways and then critique their strengths and weaknesses. Write the same content for different audiences and likewise critique. Don't just write words, read them aloud to yourself to hear their rhythm.

Fourth, 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 excellent book "Gigabit Networking" (Addison-Wesley, 1993). Focus on its design, selecting a chapter and reading just the paragraph topic sentences (each begins a paragraph). Largely, those PTSs tell the chapter's story, but, of course, in less detail. Each paragraph fleshes out the topic sentence, and nothing else.

Fifth, read a bunch of theses. Find well-written ones, not considering their relevance to your field. Dissect the writing, and mimic relevant structures. 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 fully determine the writer's way of work that produced a well-written piece by dissecting it, there's hope — read the next section.

¹⁵ For example, the query narrative "Find students who take evening and chemistry courses" can be interpreted, sensibly, a surprisingly large number of ways.

5. Design A thesis.

Avoid False Progress

Here are the parts of my “Design a Thesis” way of work.

- A. Remember the target audience (you or them), while you write three products (contributions, blueprint, draft).
- B. Work backwards from contributions to the field, not forward from what you know or what you did.
- C. Choose vocabulary, both technical content and glue.
- D. Use figures both to find and show organizing principles, and to present details.
- E. Design a story, considering many exposition strategies, depict it in a story tree augmented with disciplined connections.
- F. Finish your design down to paragraph topic sentences before writing *any* of the draft. That is, *don't hack*.
[Wait! This statement may seem extreme to you, but read on before concluding it's a bunch of hooey.]
- G. Draft one paragraph at a time, blissfully ignoring notions outside of that one paragraph.
- H. Edit in the small, not in the large. That is, produce a thesis in one draft — and smile.

The blueprint for your thesis consists of: contributions, vocabulary, figures, a story tree, and PTSs. Armed with a blueprint you will avoid false progress, and will more efficiently produce a better thesis than if you had hacked.

A. Audience and products. To avoid the fatal audience error, remember that you always write using words, but the audience and required prettiness varies by product. The table below summarizes the product-audience pairing. Yes, of course, your sketched or polished contributions and 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 “Look, read this. I have created something new and valuable, that somebody (the field) cares about, and I can prove it.”

		Audience	Pretty
Products	Contributions	<i>You</i>	<i>No</i>
	Blueprint	<i>You</i>	<i>No</i>
	Draft	<i>Them</i>	<i>Yes</i>

B. 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, which is a lot. Remember, it's not about you; it's about them.

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.

C. Vocabulary. Vocabulary is an important structural component of your thesis. Because you live 24/7 with your thesis, you can easily forget that they (the committee, who represent the field) do not, and then be surprised when they don't use vocabulary as you do, or don't know what your vocabulary means. Also, it requires discipline to not let undefined or loosely defined terms lay about. Therefore, you should craft precise definitions for them early in design, and feature them prominently in your thesis.¹⁶ One kind of vocabulary is the content terms that come from the field, or that you invent. (I defined "draft" early in this paper.) 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 your understanding with instances." These phrases 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 papers and books about computer science 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 the audience misunderstanding you, and thinking less of your contributions. (Here, my way of work has "parts" and not components, steps, stages, pieces, elements, etc.) Also, by paying attention to vocabulary you can minimize "drift." In something as big as a thesis, written over several months (or years☺), style and vocabulary drift will naturally occur. You need to take care to avoid it.

Since it behooves you to work hard at vocabulary, keep a running glossary with content and glue term-definition pairs. Refer often to your glossary, refine it as needed, and as you design the thesis, consider where each definition will appear. You should expect that definitions will not just pop out in final form, so sketch to try out various definitions.

Finally, terminologically speaking, you should consider the possible advantages accruing to most of those who generally avoid the overuse of somewhat qualifying words. That is, don't use weasel words (such as those in the preceding sentence.).

¹⁶ Voltaire supposedly said "Before you talk to me, define your terms." "Supposedly" means I don't have the reference. Sniff, sniff -- sorry, your committee will frown on such loose, "supposedly said" quotes.

D. Figures. To help the reader perceive the meat of your thesis you should liberally use figures (including graphs and tables.) Everybody uses figures for raw experimental data and statistical results (and I'll ignore such figures here ¹⁷), but figures can do much more, showing the audience what is coming and how pieces relate. For example, before you begin to design your literature review build a table. In the Appendix Figure A1 shows a portion of the thesis literature review summary figure by Libby Shoop. Each row is about a relevant research project. Each column is about a relevant property. Each cell in the matrix describes something about a project – property pair. You, of course, will have to decide what is relevant for your work – and doing so is good for you. The terms you use should be part of your content vocabulary. Also, pay attention to the ordering of columns and rows. Pick orderings that make it easier for you to tell your story. For example, you might say something like "The circled terms in the lower right portion of Figure 2-1 highlight the differences among ..."

Other, big-picture figures might help both you and your readers. It gives the reader a 2-D picture to go with your words, which are inherently linear. For example, Figure A2 shows the field before and after Leone Barnett's thesis. Figure A3 shows a system diagram depicting what Jake Chen built and how it fits in bioinformatics. Each figure became an organizing principle for at least a thesis chapter. The summary figure at the beginning of this paper helped me write it and hopefully has helped you read it.

When you construct a figure, good things happen:

- You will more readily see what is unclear, what you need to think more about before writing.
- You will see more to write about than you would have otherwise.
- You will make your points with more consistent vocabulary.
- The figure will suggest an organizing principle for the story contained in the figure.
- The audience will think more highly of your work than they would have otherwise.

Spend time polishing your figures; it is well worth the effort. Figures A4 and A5 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 ink¹⁸ and is easier to follow because related things appearing closer to each other allows shorter, no-bend lines. Figure polishing is helped by writing. So when you put a figure 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

¹⁷ Check out R. Harris, "Information Graphics: A Comprehensive Illustrated Reference," Management Graphics, 1996. I find it a valuable resource.

¹⁸ E. Tufte, "The Visual Display of Quantitative Information," Graphics Press, 1983. Another valuable resource.

the story. That is what happened with Figures A4 and A5. When Brian Bailey and I talked about A4, we found, to his surprise, that it did not directly support the paper's story, and that the story was incomplete. We fussed with A4 saying "What does this symbol connote," "What do we want to say here," and "What's missing." We fussed with discipline, and Figure A5 was the much better result.

Read other theses looking for figures that you can mimic. Pay attention to how well integrated each is with the text.

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 can work but I prefer big picture figure explanations in the text.

E. Story tree. So, you've been diligent, and have contributions, vocabulary and some figures, and, probably, a pile of notes to yourself, but how do you turn that into a thesis?

Your thesis undoubtedly contains an intricately connected network of topics, which you must organize so readers can read linearly and grok it. They read imperfectly, so they need introductions, look-aheads, look-backs, transitions, and summaries in addition to the topics proper. Somehow you must take them on a tour through your network, and you necessarily will refer to a topic more than once, gradually exposing them to intricacies. To accomplish this you need to invent a story tree, including paragraph topic sentences, and disciplined connections.

Story tree. What you need to invent is the heart of your design, a story tree, that is, a detailed depiction of the pieces and subpieces of your thesis, the order in which they appear, and the smaller signposts that guide the reader along the way. You decide the order in which the readers will see the topics, how much content to present each time a topic is touched, what you say again, but with a different emphasis or elaboration, what guideposts you add to help the audience.

Why use the term "story tree"? Well, "outline" is too vague, and could lead you to merely sketch out a few points and begin hacking away at a draft. A "table of contents" is just the major branches of the story tree that you end up with. However, "story tree" puts the emphasis on designing for readers. After all, you tell a "story" to others, and, therefore, you think about what they already know and believe, and how you are going to get them to your destination for them — understanding and accepting your contributions. "Tree" connotes that i) you organize your story into pieces with subpieces (limbs), but unlike an oak, a thesis story tree's limbs get linearized, and ii) you prune your tree, tailoring it to your audience, and not saying all that you know. By using "story" you switch your thinking from yourself to the audience – a story is for the listeners, not the teller.

What does a story tree look like? It can take many forms: on index cards as McPhee works, yellow stickies, a word processor's outline mode, a big piece of paper with words and lines on it, or whatever works for you. It's the content, not the form that matters. Its form can change as you move from sketches of it to more fleshed out versions.

PTSs. How detailed a design? This is crucial: **Design down to paragraph topic sentences.** [More on PTSs in a bit.] You might prefer to keep the PTSs separate from the rest of the topic tree, because they can get bulky. You should feel free to augment your PTSs or any part of the design with sketches noting the main points you intend to make, or key vocabulary terms. After all, you are the design's audience.

Connections -- with discipline. What about connecting the content in one part of the story tree with content elsewhere? You can think about at least these kinds of connections: I for introduction for a subtree or a sequence of siblings, T for transition between topics, S for summary for a subtree or a collection of siblings, and C for crossing (perhaps grafting is a better term).¹⁹ A crossing shows where, elsewhere in the tree²⁰, it crosses to.

Building a story tree helps you discipline yourself. You can't just say anything anywhere – it will drive your audience away. You won't put connections everywhere, but you explicitly choose them during design, and not yield to temptation by putting in gratuitous connections where it happens to feel appropriate as you write. 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 connections, 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 and its 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 connections 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.

As a good example, consider the quote below from Richard Coe²¹ and the subtree (two levels plus PTSs) for its part of Coe's story.²² He wrote a ten-paragraph subsection entitled "Subject Matter," which is the fifth of seven siblings in an "Academic Discourse" section. It begins with a one sentence introduction, I, paragraph which lists six terms (now in bold) characterizing good academic writing. Each term appears in the first sentence of a following paragraph – in the same order as in the I paragraph. Two of the terms take two paragraphs each to explain, but reading is easy because of the absence of the next term clearly signals a continuation of the preceding paragraph's topic (for paragraphs 4 and 8), or a transition (paragraph 6). Most paragraphs are inward looking, but one (paragraph 7) has a crossing connection – to some earlier place in the tree. But that crossing is in the middle of the passage and, naturally, it seems to me, Coe de-emphasized it by placing it in parentheses. Also, armed with a connections mindset, I noticed that Coe chose to have no summary paragraph.

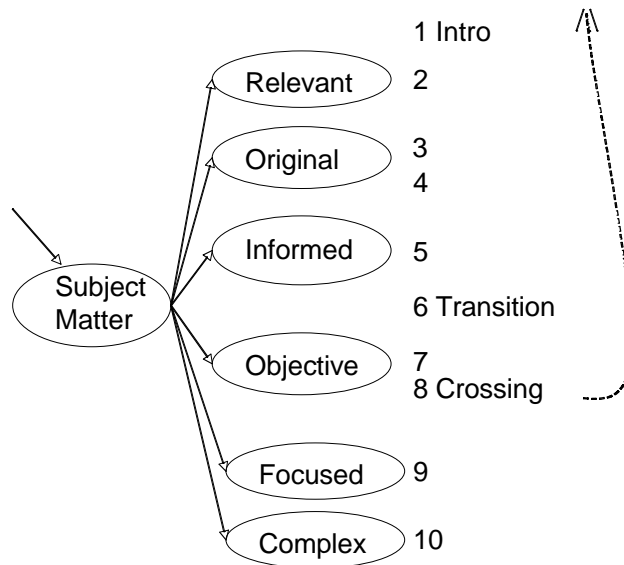
¹⁹ You might find it useful to employ a finer taxonomy of components. For example, see the rhetorical structure theory at <http://www.sfu.ca/rst/01intro/definitions.html>.

²⁰ Yes, eager computer scientist, I know. The ordering and connections make the story tree not a pure tree. Please forgive me. I choose simplicity over pedantry.

²¹ Pages 443-443 in Richard M. Coe excellent book, "Process, Form and Substance, a Rhetoric for Advanced Writers" 2nd Edition, Prentice Hall, 1990.

²² Beware! I made this story tree pretty to aid your reading, but my own version was messy.

1. Good academic writing is not only rigorous but **relevant, original, informed, objective, sharply focused**, and sophisticated in its treatment of **complex** subject matter.
2. The subject matter of a piece of academic writing must be **relevant** to an intellectual problem ...
3. Because the purpose of academic discourse is to generate knowledge, academic writing should be "**original**" ...
4. The challenge must be carried out in a particular way. You are allowed to express any opinion you like, but –and this can be a big “but” – you must express and support that opinion in a proper way. ...
5. Opinions must be “**informed.**” ... based on facts.
6. For this reason, ... ceases to be merely personal or subjective. ...
7. Certain key terms, such as *objectivity*, represent values generally shared by the academic community. It is important to find out precisely what objectivity means, in practice, in your particular field (see pages 320-22 on operational definition).
8. Like almost all ... technical terms. ...
9. Academic writing usually has a sharp and narrow **focus.** ...
10. Academic discourse tends to concern itself with **complexities.** ...



Coe was not rigidly consistent in style and terminology. Of the six terms, in their respective paragraph topic sentences, he a) quoted two, italicized one, and chose no special formatting for three; and b) exactly repeated three and modified three. Because of its size, in your thesis you should seek consistency, and employ parallel construction where you can.

Also, in “Blink” Malcolm Gladwell²³ weaves together disparate threads in his writing, but does so with discipline. He crosses, C, the threads mainly at the beginning and end of subsections. For example, on page 117 he ended a section on improvisational comedy with a crossing paragraph, tying it to previously described unorthodox techniques in a military war game.

²³ Malcolm Gladwell, "Blink," Little, Brown, 2005.

Do you have to be particularly quick-witted or clever or light on your feet to play that scene? Not really. It's a perfectly straightforward scene. The humor arises entirely out of how steadfastly the participants adhere to the rule that no suggestion can be denied. If can create the right framework, all of a sudden, engaging in the kind of fluid, effortless, spur-of-the-moment dialogue that makes for good improv theater becomes a lot easier. This is what Paul Van Riper understood in Millennium Challenge. He didn't just put his team up onstage and hope and pray that funny dialogue popped into their heads. He created the conditions for successful spontaneity.

Not surprisingly, examples from a book for advanced writers, and a book by an award-winning writer exhibit the product of advanced writing skills.

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. Cohesion 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.) Crossings lower cohesion and increase coupling – and make your writing, and, later, your audience's reading harder

How should you go about designing a story tree? Well, any real design process 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, you think along several avenues.

- You think about topic relationships (a capitalized letter stands for a (small or large) topic):
 - "A must be explained before B." For example, I needed to define "draft" up front, before using it, so that I would reduce your uncertainty about this paper's prime term.
 - "A, B and C are a set of components of D." For example, this "way of work" has eight components, called "parts," and I chose a sensible order for them, but other orderings would have worked too.
 - "D and E are nearly separable but both have X in common," etc. For example, PTSs are a component both of a story tree and of writing craft in general. I thought PTSs are important enough to put the material about them in section 2 and not subordinate it here.

Such thinking reduces the number of possible story orderings.

- You think about what they are ready to receive²⁴. You cannot just spew words – storytellers weave threads. So, you expose some content and enough of its context

²⁴ This notion is consistent with learning theory. Read Ausubel (Ausubel, D., Novak, J., & Hanesian, H. (1978). Educational Psychology: A Cognitive View (2nd Ed.). New York: Holt, Rinehart & Winston), Perry (Perry, W.G. (1999). Forms of Ethical and Intellectual Development in the College Years. San Francisco: Jossey-Bass Publishers – first published in 1979, but so good it was revived and has a fine forward), and Caine (Caine, Renate Nummela and Caine, Geoffrey. (1994). Making Connections: Teaching and the Human Brain. Menlo Park, CA: Addison-Wesley Publishing Company). Each is a fine work, but I *really* value Perry.

so that they can get some of your message. You return to that content again, and, crucially, the readers are different.

- You think about inform and persuade. You need to inform the audience about the scholarly content of your contributions, connecting to what they know but not boring them with “things that go without saying” – don’t say them.²⁵ You need to noodle about what, if anything, in your contributions is controversial, and conjure up persuasive arguments to overcome their objectives.
- You think about content to include and to exclude. That is, you bound the thesis. What is in is what they need to understand and accept your proposed contributions. Everything else is left out or relegated to an appendix or, these days, a webpage.²⁶

Second, you experiment with organizing principles. For example, a Libby Shoop chapter took the readers on a tour of a system figure (not shown), 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, etc. Let me say this stronger:

- Build figures *early* because they will suggest organizing principles to you.

Third, look at siblings and decide if they are a set or a sequence of topics. For a sequence, usually you will choose to expose them first to last, or the reverse. For a set of siblings, any order will do, 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 about pruning, 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.

Fourth, study your vocabulary; perhaps it is vague or incomplete. Talk it out with your advisor or a friend. When talking with one of my students as she or he struggles with a hard notion, I often will hear a gem. When that happens, I say “don’t talk anymore until you write that down.” From such gems a better story emerges. So you can “talk to learn” too.

F. Finish your design down to PTSs before drafting. 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.**

²⁵ Your contributions may re-envision things that other think of as going without saying.

²⁶ Of course, you can use a footnote (as I am doing now), or an endnote, or collect some loose ends into a section (you will see one in a few pages). Also, of course, watch out for putting your main line topics in a footnote, as I am doing now, since it is likely to be jarring for the reader. If your footnote extends more than a few lines think about why it is not part of your main text. [So, thinking as a story tree designer, would you move this footnote?]

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*.²⁷

False Progress. Following this finish-the-design-first 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.” However, real progress does not occur if you write easy paragraphs, and then 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. You benefit if you were patient, and the PTSs, for the paragraphs that you didn't write, get moved, changed, or discarded.

False progress occurs a lot in software development. I know people who were told by their bosses to start writing code even though the specifications were barely begun. The bosses wanted to get going, but caused problems. Not surprisingly that software was tossed. You may have a similar problem with your advisor. (More on that in a bit.) False progress occurs when you say “I don't know what I have to say so I'll just start writing,” and fall into the fatal wrong audience error, or fail to focus on contributions.

Be careful to not misunderstand this last sentence. You put words on paper to learn, to invent vocabulary, and to determine if your story and figures are mutually supportive. However, when noodling you are writing for *you*, not them, so you should sketch, and perhaps do so on paper, not on a computer, so that you cannot forget that it is a sketch.

You will be willing to be patient (i.e., not draft) 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.

G. Draft one paragraph at a time. If you have a good quality design, with high cohesion and low coupling, then you can, and should, write one paragraph at a time. Doing so has several benefits. Like a carpenter with a blueprint, it allows you to focus on just one thing, and therefore makes you efficient. It raises your gumption level, because, while writing a

²⁷ 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) Both articles address teaching design.

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 keeps you from making gratuitous connections. It enables you to judge the goodness of a paragraph, because you know its context – the intent of the surrounding paragraphs. 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.

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. Of course you should write term-defining paragraphs before any that use the terms, because otherwise you are not grounded and will waffle. Write a section all at once, rather than skipping around. Write paragraphs with relatively uncertain 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.

A Caveat. When can you begin to flesh out the design? In other words, do you really have to wait until the entire topic tree is designed before writing at all? Well, if a portion of the topic tree is unfuzzy and relatively disconnected²⁸ from other parts, then you can draft it, while continuing to develop other parts of the tree, or to accomplish the research for other parts. For example, you might decide to draft your literature review early, because you have a literature review from a proposal, and deem its organization adequate for the thesis. However, be careful. The more you draft early based on an incomplete design, the higher the risk of false progress when you discover its flaws.

H. Edit. After you produce the draft, of course you will edit. However, 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. Furthermore, because your paragraphs are focused, you will have less editing to do than if, without a design, words just spewed forth words.

Some loose ends.

[Oops, no introduction. So, you write one.]

What if your thesis diverges considerably from your proposal? Well, this happens quite often, especially if your proposal presentation was early, so do not worry about conforming to your old content or format. Since your proposal and thesis are different documents and

²⁸ I could have used the term “isolated,” but by using the term “disconnected” I hoped that you would more readily understand that I meant a portion of a story tree with few connections in or out. I followed my own advice of not using gratuitous synonyms.

can be separated by years, do not be surprised if you save little of the proposal's verbiage or structure. Remember, you will be advancing, discovering new things and ways of expressing them.

What if you are in a hurry to finish and feel the urge to start writing? Well, sure, you can write, but if you do not have a design you will finish later, not sooner, because you will be forced to re-draft. Do not write a thousand pages and keep two hundred.

What if part way through drafting you discover that the design, and therefore the draft, is seriously flawed? Well, with 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. You regret your bad planning and vow to do better next time. Do likewise when drafting a thesis.

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

What if your field, or advisor, has a standard template of say six chapters, so you are told to use it. Well, sure, but that's way too coarse an outline, and is like saying to a builder no more than "Build me a three-bedroom, two-bath house." Use that coarse template if it fits, but do not think it will allow you to avoid the hard work. Be willing to fight against that template if it doesn't fit.


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.

Is there a size limit? Well, a thesis has no official page limit, so do not worry about it. However, if you are writing for a journal or conference you will face limits. In the latter case use your topic tree to help you prioritize topics and contributions for pruning.

What about the Abstract? A thesis 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. Don't make the mistake of making your abstract into a table of contents. For example, don't say "I address several issues arising from my experimental data." Such statements are about form not content.

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 works with you along the way, approves the design, and then accepts the draft you write. For that to happen the advisor also must be able to imagine the flesh that you will put on the bones.



When does it end? When you are done drafting and polishing, you must "defend" the thesis to a committee of professors. 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. Your figures will provide an organizing principle for your thesis talk, just as they did for the written thesis.

Conclusion.

You should believe that you can produce a one-draft thesis, because focusing on contributions allows you to first choose vocabulary, craft figures, and grow a story tree down to PTSs, and to then, for an audience of other experts, draft things, one at a time, that you have the skill to draft, namely, paragraphs.

To help you believe in design as a way of work, think about this: While I can readily juggle three rocks and my skilled friend John Riedl can juggle, for a while, six rocks, nobody can juggle a hundred rocks – vertically. However, horizontally (on a table) juggling a hundred is not a problem, because a rock stays where I put it and I can focus on one or a few at a time. Drafting without a design is like vertical juggling, but with a design is like horizontal juggling -- much easier.

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.

Addendum for advisors.

Should you work with your students following this design way of work? Well, of course I think so, but before you will adopt this way of work you need to believe that both advisor and advisee benefit. To help you decide for yourself, consider this: advising students is immensely rewarding but quite taxing too, so if you can advise better with less effort then your satisfaction to pain ratio improves. I find it hard work to give substantive feedback about sketches and designs, but really hard work to do so for hacked-together drafts. With less total effort I give better feedback about the thesis in its various stages of gestation.

Furthermore, when I am in on the thesis design I short circuit avoidable bad surprises. I liken it to influencing an avalanche, where it is easy to correct the flow of pebbles, but heavy lifting to shift boulders.

Here are opposing arguments, only slightly edited quotes from professors, plus my responses.

"The student is responsible for the thesis. My role is to judge." At its worst, this utterance reflects laziness. But defining our role is a valid discussion issue, and professors do vary widely in how involved they are in an advisee's thesis. However, surely we are not merely neutral judges. I know of no student who needed no advising, but I do know professors who have stood aside and let students suffer for want of good advice. I view the thesis as shared responsibility, otherwise why would we include a "students advised" section in our resumes or yearly bragging reports?

"By gum, I suffered and it was good for me." Shame on you! Having high standards and inflicting pointless pain (perhaps by neglect) are not the same. The former is laudable; the latter is not.

"I won't know if I like the thesis until I see it." Well, experience tells me just the opposite. By giving feedback early and often – pre-draft – I influence the bones and then I get flesh that I like. I also understand what I get better because I've had a longer learning period. When I am not in on the design and get mush to read I struggle to determine what is good and bad in it.

"I can't just approve the draft; the student expects a critique." Of course, every advisor (and committed committee member) has to put his or her mark on the thesis. Otherwise we shirk our duty. But that can be done along the way and need not wait until the end.

"No, No! Write, revise, write, revise, etc – that is how it should be done." I disagree because I think it wasteful to pretend to write for them. Instead, I think "sketch, revise, sketch, etc.," designing a blueprint is better for both student and advisor.

"I do not know how to advise this way." We all are lifelong learners, and for the most part, the first step is the hardest. So begin with something small, and expect rapid improvement.

"I do not know how to work this way, and I dread showing my inadequacy." This is a real fear felt by many but few express it. Be brave – your students are.

Finally, I found that my own writing has improved since I started advising about writing. Such improvement is consistent with the notion that to really learn something, teach it. So, "try it, you'll like it."

Appendix

Reference	Computation Type	Time Complexity	Space Complexity	Alignment Type	Typical Use	Scoring Measure	Allow For in Scoring Measure:			Gap Weighting Function	Ability to Assess Statistical Significance
							indels/gaps	gap penalty for unmatched ends	other		
						(similarity or distance)	yes	yes		(e.g.: single, multiple, affine linear, affine concave)	
Needlman & Wunsch, 1970	dynamic programming	O(MN)	O(MN)	global	alignment of two presumably homologous sequences	similarity	yes	yes		single indels each penalized	Difficult
				local	alignment of common subsequences of two sequences which may have the same functional site	similarity	yes	no		single indels each penalized	Difficult
Sankoff, 1972	dynamic programming	O(MN)	O(MN)	local	alignment of common sub-sequences	similarity	yes	no		given number of indels treated as one and given different penalty	Difficult
Sellers, 1974	dynamic programming	O(MN)	O(MN)	global	alignment of two presumably homologous sequences	distance	yes	yes		single indels each penalized	Difficult
Waterman, Smith, & Beyer, 1976	dynamic programming	O(MN ²)	O(MN)	global	alignment of two presumably homologous sequences	distance	yes	yes		multiple - allow indels of length greater than 1 ("gaps")	Difficult
Smith & Waterman, 1981b	dynamic programming	O(MN ²)	O(MN)	local	alignment of common sub-sequences	similarity	yes	no		affine linear function for gaps	Difficult
Gotoh, 1982	dynamic programming	O(MN)	O(MN)	global or local	alignment of two presumably homologous sequences OR alignment of common subsequences	similarity	yes	set parameters to determine		affine linear function for gaps	Difficult
Taylor, 1984	dynamic programming	O(MN)	O(MN)	global or local	alignment of two presumably homologous sequences OR alignment of common subsequences	distance	yes	set parameters to determine		allowed lower penalty for large gaps	Difficult
Waterman, 1984a	dynamic programming	O(MN) (est.)	O(MN)	global or local	alignment of two presumably homologous sequences OR alignment of common subsequences	similarity	yes	set parameters to determine		affine concave function for gaps	Difficult

Figure A-1 Part of Libby Shoop's literature review summary table.

Figure 1.1. Big Picture

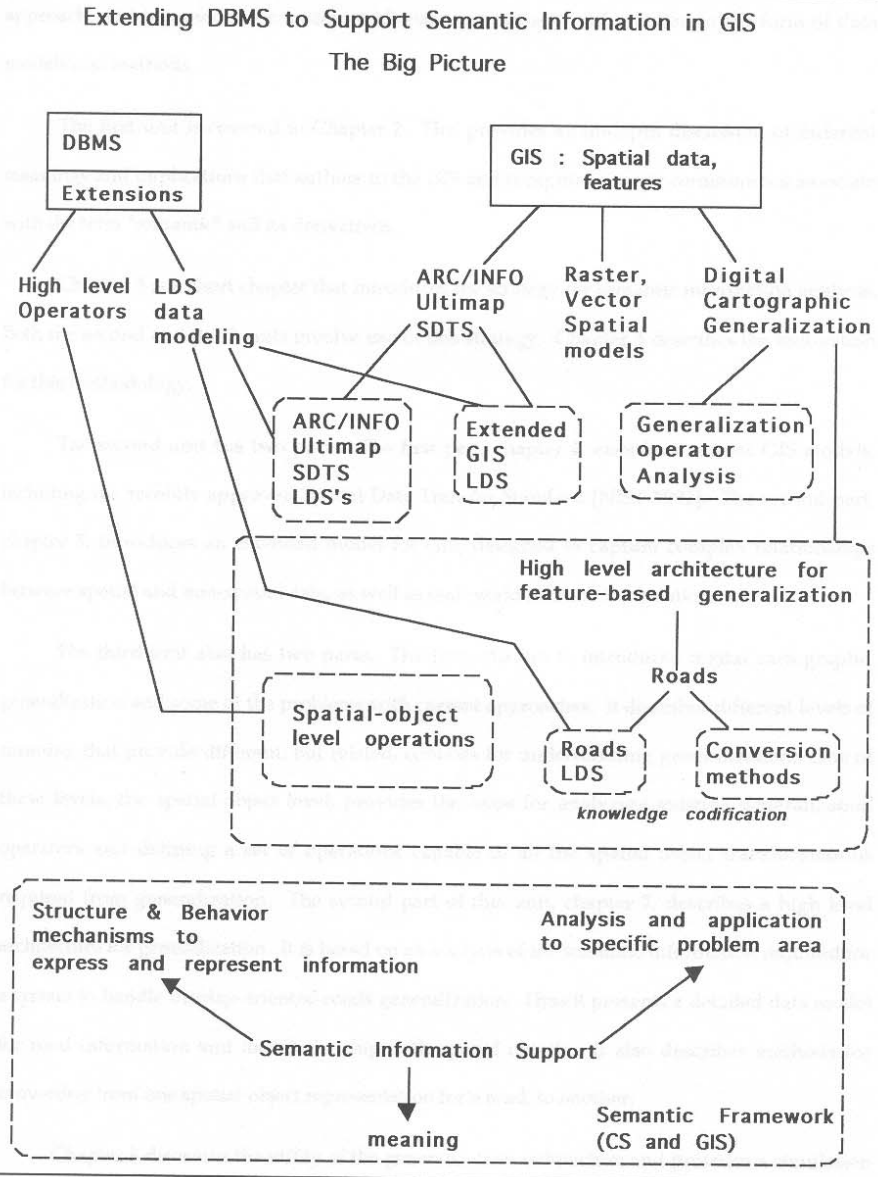


Figure A-2 Leone Barnett's before and after figure.

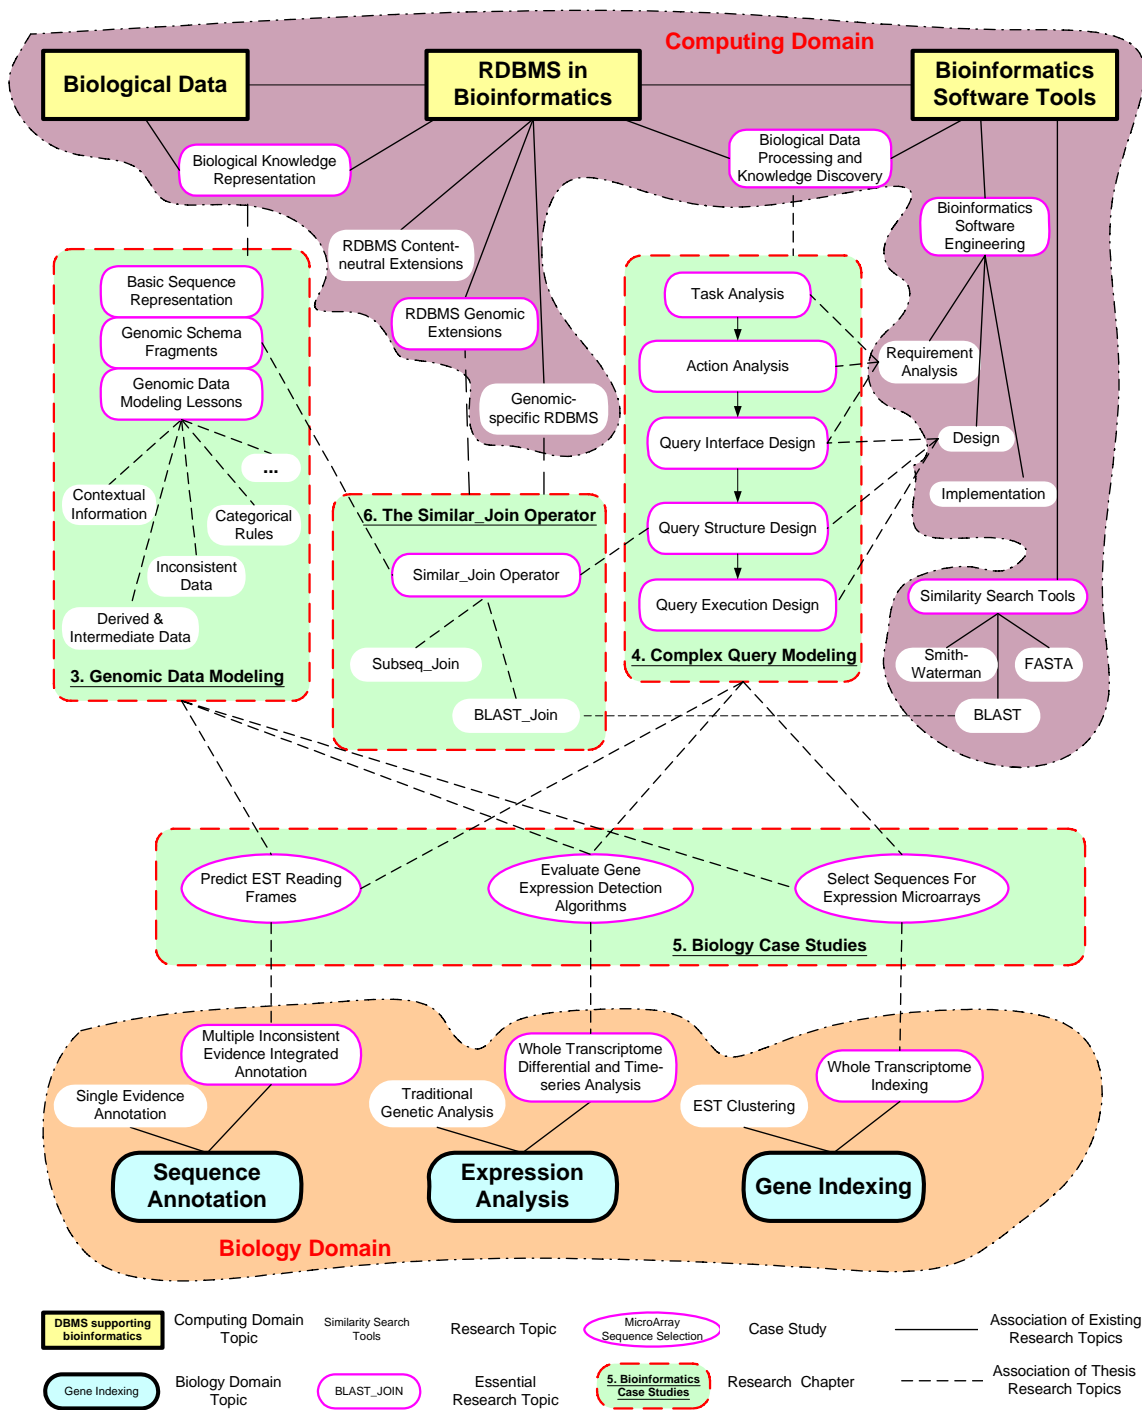


Figure A-3 Jake Chen's Thesis Overview Diagram

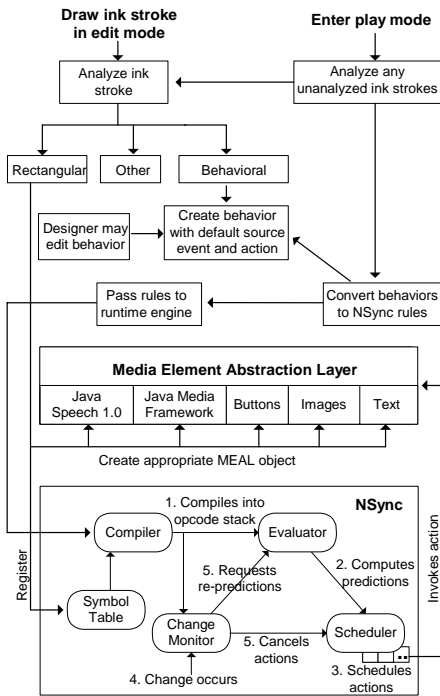


Figure A-4 “Intermediate”

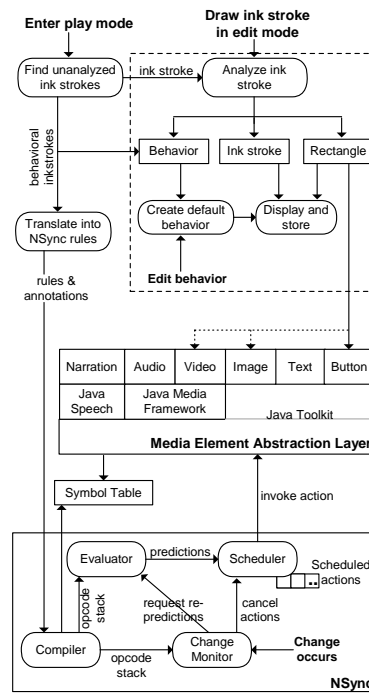


Figure A-5 “Polished”

Figure A-5 is from Brian Bailey, Joseph Konstan, and John Carlis, “DEMAIS: Designing Multimedia Applications with Interactive Storyboards,” Proceedings ACM Multimedia 2001.