

**SOME THOUGHTS ON WHAT IT
TAKES TO PRODUCE A GOOD PH.D.
THESIS**

**Avi Kak
Purdue University**

December 14, 2012
7:46pm

This presentation is dedicated to the fact that

..... in the best of relationships between a professor and a graduate student, the student strives to show the professor why the professor is not as smart as the professor considers himself/herself to be, and the professor strives to show the student why the student is actually smarter than the student believes himself/herself to be.

DIFFERENT PHASES OF PH.D. RESEARCH

- finding a good problem
- staying on top of the literature
- getting plugged into the broader research community
- communication of research results through oral presentations and writing

FINDING A GOOD PROBLEM

This is probably the most stress inducing phase of the whole program.

How does one go about finding a good problem?

- Ask your major professor
- Ask your office mate
- Read the tea leaves
- Ask your mom
-

FINDING A GOOD PROBLEM (contd.)

In my mind, the correct answer is: None of the options listed on the previous slide.

To appreciate my answer, you have to get to the bottom of what engineering research is fundamentally about.

WHAT DOES RESEARCH IN ENGINEERING REALLY MEAN

Basically, engineering research is about

..... observing and understanding the world around you with regard to how things work now and how they could be made to work better

..... discovering the current best practice in your area and pushing that to a higher level of performance

..... bringing together two hitherto disparate threads of engineering knowledge and creating a new thread for study and analysis

HOW DOES ONE DISCOVER A GOOD PROBLEM

In order to discover a good problem

..... you have to first push yourself to the current state of the art, before you can advance the state of the art.

Are there any strategies for rapidly pushing oneself to the current state of the art?

STRATEGY FOR DISCOVERING THE STATE OF THE ART

In engineering research, I believe that the best strategy is to actually try to do a state-of-the-art experiment.

STRATEGY ... (contd.)

If you want to discover a good problem to work on in any area of engineering, there does not exist a faster way to get to the state of the art than creating your own implementation for a core problem in that area.

For example, let's say you want to understand the state of the art in information retrieval from large software libraries.

There are probably a couple of hundred papers now that have been published on the subject of information retrieval from software libraries. These papers use a variety of methods that range from static source code analysis to the construction of statistical models of the source code libraries using techniques developed by folks in information retrieval from large text corpora.

STRATEGY ... (contd.)

You could spend a couple of years trying to read all these papers, but by the time you are done, there will be another 100 papers to read.

If, instead of chasing at the outset all the papers that are out there, you create your own retrieval engine, you are much more likely to get a good feel for the state of the art (even if your own implementation is rather crude compared to the best out there).

The process of creating your own implementation will give you deep intuitions that would be hard to acquire by just reading the literature.

MAKING EFFICIENT YOUR PROBLEM DISCOVERY PHASE

It is much more efficient if the problem discovery phase is experiment-driven as opposed to literature-driven. **What you read in the literature should be dictated by your current experimental obsession, as opposed to the other way around.**

STAYING ON TOP OF THE LITERATURE

This is probably the most traumatic phase of the whole program.

Much technical literature is poorly written, designed more to hide than to reveal, designed more to obfuscate than to clarify, designed to gain short term recognition, etc.*

In other words, much technical literature is written with motives that are less than noble.

How does one cope?

*If you believe that scientists and researchers are a nobler breed than most, you are mistaken.

MY STRATEGIES FOR COPING WITH THE LITERATURE GLUT

Every engineering contribution is based on assumptions about the real world. When I look at a new paper, my first attempt is to quickly extract those assumptions. If I find those assumptions excessively unrealistic, I do not pay much further attention to the paper.

I read papers to seek out their limitations. But some authors do a great job of hiding the limitations.

Sometimes I discount papers if I have already written the authors off in my mind.

GETTING TO THE BOTTOM OF A RESEARCH PAPER

In a face-to-face interaction (even by e-mail sometimes), people are more likely to tell you about the limitations of their work, limitations that they did not mention in their written papers.

In any case, if your overall research effort is experiment driven (as opposed to literature driven), you are much more likely to spot the limitations in the papers written by other people.

GETTING PLUGGED INTO THE BROADER RESEARCH COMMUNITY

This is probably the most frustrating phase of the whole program.

Every research area has its in-group. People who are already on the inside make it difficult for people from the outside to break in.

Research communities operate like small tribes, with each tribe considering itself to be the keeper of the truth.

Each tribe creates its own rules for what separates mediocrity from excellence. Unfortunately, there is much subjectivity in the application of the rules, with favors implicitly granted to those on the inside.

GETTING PLUGGED IN (contd.)

There is another important reason for getting plugged into a research community:

The research program of the funding agencies is determined to a great extent by the collective debate that takes place within research communities. So if you want funding for your research, you must become a part of the debate.

Moreover, everyone needs recommendation letters when you are trying to get a new job, or when you come up for tenure, promotion, etc.

PROGRESSIVE STEPS FOR BREAKING INTO A NEW RESEARCH COMMUNITY

- Actively participating in conferences and workshops in order to become noticed.
- Expressing verbal interest in other people's work at conferences and workshops and following that up with e-mail interaction.
- Forming friendships and collaboration with researchers from other institutions.
- Volunteering to help out with workshops.
- Volunteering to organize workshops.

PROGRESSIVE STEPS (contd.)

- Volunteering to help out with conferences.
- Volunteering to organize conferences.
- Volunteering to help out with journal refereeing.
- etc.

PROGRESSIVE STEPS (contd.)

Here is a guideline for those Ph.D candidates who want to work in universities:

During the last third of a Ph.D program, about a third of your mental focus should be on the research world outside.

During this period a lot of your energy has to go into forming friendships (they will be your future collaborators) with people on the outside.

COMMUNICATING RESEARCH RESULTS THROUGH ORAL PRESENTATIONS AND WRITING

Ability to express ideas precisely and unambiguously is a key to success in all human endeavors, particularly so in research.

WRITING IS CENTRAL TO GOOD RESEARCH

A Ph.D. is, as the degree says, a doctorate in philosophy, a doctorate in ideas, a degree that requires that the chosen ideas be articulated precisely and with rigor.

Moreover, writing imposes a discipline on thinking. Every time you write something down, you are committing yourself to a position. The act of making that commitment forces you to examine with care what it is that you are writing.

POST PH.D LIFE

Post Ph.D. life in industry

Post Ph.D. life in a research university

Post Ph.D life in a non-research university

There is world of a difference between these three lifestyles. This is not to say that any one particular post Ph.D existence is better than the other two.

If you are not sufficiently self-aware to know where you belong among these three possibilities, you could end up frustrated and with a lot of emotional and other problems down the road.

The more than 40 Ph.D.'s that I have produced populate all three categories.

FINALLY ...

Practically all research consists of **incremental** advances over the state of the art.

Even the most stunning developments are stunning only from the outside. To those on the inside, they are practically always incremental.

In rare cases, when they seem stunning on the inside, it is because someone injected a brand new approach (from what was until then an unrelated area) into a problem domain.