DRAGONS—©CAN STOCK PHOTO/DVARG. CAP AND DIPLOMA—IMAGE LICENSED BY GRAPHIC STOCK.

Finding a research problem: Tips for new Ph.D. students

Kundan Kandhway

By the time Ph.D. students complete their degree, they are expected to have three to four major contributions in their dissertation. However, the foundation stone of their thesis is laid at the very beginning of their degree. The area in which the student is supposed to work is usually known beforehand, and it is the same as the research area of his/her advisor. But more often than not, students are faced with the challenge of finding a suitable problem on which to start their research work. Skills to achieve this in an organized manner, keeping future implications in mind, are seldom taught to students. This article is to help new Ph.D. students find an interesting and suitable research problem at the start of their degree.

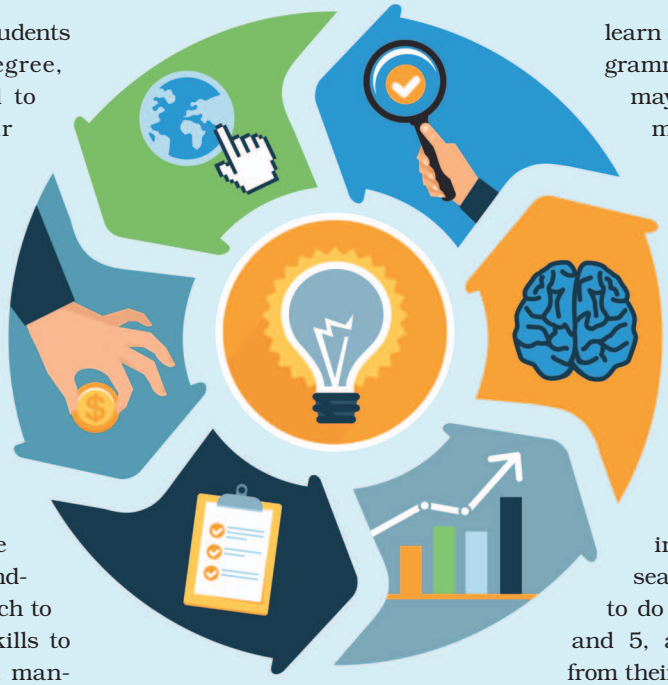

IMAGE LICENSED BY INGRAM PUBLISHING

Inception to conclusion

From the beginning to the end of the program, a Ph.D. student goes through the following steps:

1) Decide on the broad area on which to work (e.g., complex networks, wireless communications, or speech processing). Usually you will join a research group based on your area of interest.
2) Choose a subfield within that broad area (e.g., cooperative communications or cognitive radio within wireless communications; epidemiology or error/attack tolerance of networks within complex networks).
3) Find a specific unsolved problem in the subfield.
4) Propose solutions to solve the problem. You may be required to learn new mathematical and programming tools. The solution may involve designing, implementing, and experimenting with (sub)systems.
5) Write up the research for publication in conferences and journals.
6) Present the findings at conferences (and also during comprehensive exams and thesis defense).

Steps 3–6 are repeated until there is enough substantial new material in the dissertation. All research students are expected to do the major work in steps 4 and 5, although help is available from their advisor and other collaborators at important junctures.

This article focuses on steps 2 and 3, which may be a challenge for some new students (partly due to a lack of resources catering to this—specific to electrical and computer engineering) and because, if not done properly, all subsequent efforts in solving the problem and writing the results are wasted.

Various scenarios

Depending upon the university and research culture of the group you join, you may face the following scenarios with respect to the expectation of finding your research problem.

Digital Object Identifier 10.1109/MPOT.2014.2313072
Date of publication: 6 May 2015

The hunt for a problem should begin early, and you should start familiarizing yourself with the area of your interest as soon as possible.

- Sometimes you will join a specific project, and the problem is predefined. It is in your interest to do enough digging before accepting the offer. An example is an industry-sponsored project that has deliverables and deadlines. If you have an option of being a part of a big project, you can choose the part which interests you the most (e.g., mathematical modeling and analysis, simulation/programming, hardware implementation).
- The advisor gives you the first problem but expects you to find future problems based on your interests and comfort.
- The advisor asks you to find a problem on which to work. This gives you a very good opportunity to find something that interests you and is aligned to your career goals. Since the problem is your brainchild, you will be more motivated to solve it and excited to see the results. Also, you lead from the outset and can use the necessary literature review that is completed at future stages (possibly even beyond your days as a student).

Features of good problems

Understanding the features of a good problem is necessary to identify good ones. Some of them are listed below.

1) Personal interest: Does the problem interest the student who will work on it? It is of paramount importance because a disinterested student will lose the motivation to keep going for months when the problem is being solved.
2) Impact on society: Dies the area to which the problem belongs have an impact on society (e.g., integrating renewables into the smart grid to reduce carbon emission, automatic route guidance to avoid road congestion).
3) Interesting to industry: Does industry find the problem compelling? For example, cognitive radio that promises to mitigate the issues caused by limited bandwidth or a design that cuts down costs (an efficient circuit design with fewer components).
4) Impact on the research community: What type of impact will the resolution of the problem bring? Welcome problems include those that are recognized as difficult by the community, one that is solved with fewer modeling assumptions, or a system design that performs significantly better than its predecessors (e.g., a prototype of a full duplex communication system that can transmit and receive signals simultaneously with a high data rate).
5) Timely: Is the problem currently of interest to the research community or will it be of interest in the near future? If an increasing number of papers are being published in a specific field within the last few months (which can be determined from, for example, Google Scholar or the monthly Publications Contents Digest by the IEEE Communications Society), a problem in such a field is timely.
6) Aligned to your career goals: Is the problem relevant to your interests? If you have a career goal in mind (e.g., joining industry after graduation), the chosen problem should be one that the industry cares about.
7) Enlarges the skill set: Does your problem require you to learn new mathematical or software tools that will be useful in the future? If you used optimization techniques during your master's research (or an internship), then you may want to select a problem that will let you learn, for example, queuing theory.
8) The bigger picture: Does the problem lead to one or two further problems that will allow you to use some of the expertise gained from solving the initial problem and analyzing the subsequent ones? This also leads to the development of a coherent and organized dissertation and keeps you focused.

A good research problem should interest you. Its importance can be easily justified, and it should be worthy of the time and effort spent on its analysis and solution. In addition, the solution to the problem should add enough value to the existing body of literature.

Strategies for finding good problems

Some of the strategies for finding research problems and areas/subfields of current interest includes

- The hunt for a problem should begin early, and you should start familiarizing yourself with your area of interest as soon as possible. If your advisor works in multiple areas, talk about your interest and his/her expectations. If you are funded by a project that was approved based on a proposal, the project proposal might be available from your advisor. It will provide useful hints for finalizing the problem. Recent papers authored by your advisor may also prove useful.
- Talk to senior Ph.D. students, post doctoral candidates, and faculty members to learn their opinion about which areas are currently of more interest to the research community and what they think about the future of those areas. Invited talks and Ph.D. defense organized in the department provide good learning opportunities.
- Attending a conference or a workshop early in the Ph.D. process, perhaps as a student volunteer, will help significantly. Such an opportunity is generally not advertised, but can be

learned through your advisor. Talking to others face to face about what interests them in the area and their problem is inspiring and educational in making informed decisions.

- The problems that crop up every now and then in the news that have societal impact are worth solving. An example includes communication network design that is either robust to natural disasters, like hurricanes and cyclones, or can be resurrected with less effort. If it is known that the government is making significant investments in a sector, such as renewable energy, over the next few years, it may be worth investigating problems in those areas. A field trip may reveal the problems that some of the systems (like renewables) face after deployment without spending time on experimentation.
- Recent white papers from industry, areas in which investments are made (often made public in tech news), and ongoing work in research labs (that can be learned from company Web sites or recent publications from researchers) reveal the kind of problems in which industry is interested.
- Some of the fields in electrical and computer engineering progress very fast so having the latest information is important prior to the selection of the problem. It takes a month or two for conference proceedings to show up in the databases like IEEE *Xplore* but the information can be borrowed from someone who has attended the conference. Journal Web sites often provide "early access" to articles that are accepted but not yet assigned to an issue. Many research communities are increasingly using preprint archives (like arXiv.org) for manuscripts submitted for publication. These archives can be searched for the latest information. Advanced search options on Web sites that index scholarly work (e.g., Google Scholar) are also very useful in searching recent literature.
- Calls for papers published in reputed journals and magazines in the field (e.g., *Journal in Selected Areas in Communications* and *Communications Magazine*) also give useful hints for narrowing down the search for a latest problem. The community brainstorms before finalizing these calls, which make them quite useful.
- Satellite workshops in reputable conferences sometimes represent upcoming areas of interest. The previous and current year's Web pages can be compared to determine the newly added workshops in the conference. One can also look for the newly added fields in the "scope" of the conference. Calls for papers and the scopes listed in them may also provide useful information.
- The main method, however, is to do a thorough literature survey of the area. It should begin with review articles (not always available for newer fields) and magazine articles because they are less technical and written for nonexperts. They are useful in getting a bird's eye view of the area and relevant research issues. Conference and journal articles reveal the state of the art, which helps in indentifying open problems.

Before your exact problem has been finalized, it is often not necessary to concern yourself with every tiny detail of the paper. The aim should be to identify the following: the problem(s) that the paper is addressing, its novelty with respect to the literature before it (answering the question about how novel a problem should be before it is considered significant), the authors' approach to solving the problem (however, minor details may be skipped for now), the mathematical and software tools used, and the results and conclusions. Also note your likes and dislikes about the presentation and organization of the paper; it may help improve your manuscript in the future.

After gaining some experience in the subject, you may try and "guess" what the paper is about (the problem and solution approach) after reading the abstract (and perhaps the introduction and conclusion). It may lead to new directions worth exploring, without biasing your mind with the contents of the paper. For example, the paper might have used heuristics to tackle a problem, but you thought of something better.

The future work/conclusion sections of the papers (especially in review articles) provide invaluable hints about open problems (however, if the paper is new, the authors may be pursuing some of them). Make sure that the survey you are conducting is not shallow, but a "perfect" literature survey is also not possible.

Once an interesting problem has been identified, it should be discussed with your advisor to make sure it is significant and publishable.

Once an interesting problem has been identified, it should be discussed with your advisor to make sure it is significant and publishable. If not, later efforts in solving the problem and writing the results will be wasted. The problem should also be solvable in the given timeframe. This is especially important for new students. Gradually students are expected to improve in assessing the quality and feasibility of problems.

Some papers from the literature survey may require a closer look at this stage to fine tune the problem statement.

- The main tools to carry out the literature survey are search

engines such as Google Scholar and scientific databases like IEEE *Xplore*. Search engines can query many scientific databases simultaneously and often rank articles in order of relevance. The advanced search options allows you to query specific journal/conference database, find articles between specified dates, and search by author names. Links to articles that are freely available on the Web can also be found using search engines. If you are using your university's network, then you will have access to the scholarly databases to which your university has subscribed. The search tool in those databases can be directly used to find relevant articles. You may be able to refine search results recursively based on various criteria to narrow them down to the most relevant articles.

- Once an article is read, there is often a need to go through one or more articles in its reference list to better understand the material. This may potentially lead to a never-ending loop. Keep in mind that the literature survey need not be perfect. The primary goal is to find a gap in the literature, and, once found, you should move on to the next step. If the article is too cryptic, try to find a similar article by the same or different research group, maybe one that discusses preliminary results (a conference version of a journal article).
- The critique and feedback provided by the audience during conference presentations and other opportunities such as comprehensive exams may be helpful in deciding future directions. After the problem has been identified and is being solved, Ph.D. students are often required to present their preliminary findings in a comprehensive exam. The examination committee consists of faculty members who are close to the research area of the student. This is a good opportunity to get some feedback on the chosen problem from experts other than the thesis advisor. The discussion and the faculty members' perceptions about the future of the area may be used to decide the future course of the thesis. The purpose of such an exam is to make sure that the student is on the right track in research. So if you feel (or are told) that something is wrong, then seek the help of the committee members and your research advisor.
- In many universities, Ph.D. students are required to complete a few courses as a part of their degree requirement. Make use of this opportunity to achieve a good background in your planned research area. In the first semester (when you are busy deciding your research area), take courses that teach mathematical tools that are useful in many areas (such as optimization and random processes, among others). In the subsequent semesters (when you know which area in which you will work, and perhaps have an idea of a specific problem), you can take courses that will help you in tackling your problem. This is an additional reason why you should try to narrow down to a problem as soon as possible. Get the approval of your dissertation advisor before crediting a course. Universities also allow graduate students to audit courses. It can be done in cases where only a small part of the course is useful to you.

Doing some work in a different academic setting provides an opportunity to learn from a different research philosophy, methodology, and culture.

Logistics issues

Identify the resources that will be needed and that are at your disposal for solving the problem as soon as possible. Examples include computational servers, software licenses, components and instruments required for experimentation (e.g., solar panels, certain integrated circuits, human volunteers for behavioral experiments, and the funds available for this purpose and the time required to procure the equipment. Specialty items may take longer to procure in certain regions of the world. Being proactive and planning ahead may save some time.

Risk assessment

Working in a new area (in which your advisor and other group members are also new) may seem a bit risky, but it is also rewarding. If the area is immature, there may be an opportunity to pluck some low-lying fruits. It is better to try and perhaps fail, than to not try at all. First, attempt to solve the simplest version of the problem, and gradually work your way toward the final goal. That being said, problems that are too risky should be avoided. It is important to make steady progress, and if that is not happening, you should seek help.

Some tips for other stages of research

1) You have to make a decision on when to stop working on a problem. Conferences and journals limit the number of pages of submissions. Once enough material is obtained for a reasonably convincing paper, it is a fair time to move on to the next problem. The papers read during the literature survey and the instructions to authors/reviewers on the conference/journal Web site will help in deciding how much material is enough.
2) There is always room for improving your writing and presentation skills, especially for nonnative English speakers. Spend some

time on organizing the graphs and tables. Show it to someone less familiar with your work to test the clarity of expression.

3) If you are working in an area that requires numerical computation or simulation that is resource intensive, try to optimize your code for efficient memory use and execution speed (whichever is the bottleneck). Doing so will allow it to handle larger systems. There are many optimized libraries available for modern programming languages that may be helpful. Freezing the code the first time it ran correctly is often not the best strategy.
4) Interaction with your advisor is important for a successful dissertation. Depending on the time your advisor allocates to meet students, the frequency of meetings, and his/her interaction style (domineering/participative/micromanager, etc.), you may have to devise a plan to best use your advisor's time. In the event that you disagree on a point with your advisor (which is not uncommon), identify the reasons that convinced you and present it to discuss his/her point of view. If both of you grew up in different cultures (pertinent to international students), make sure you understand each other correctly. Prepare for the meeting with your advisor to best utilize the time. Inform him/her of the important developments. If you are deviating from the prior plan (*e.g.*, adding additional assumption to a mathematical model for a physical phenomenon) bring it to his/her notice. It may keep you from wasting your effort on something that is obviously incorrect.
5) Conferences and journals accept significantly fewer manuscripts than those submitted. The acceptance rate is unlikely to be more than 50% (sometimes it is as low as 10%). So it won't hurt to be prepared for a rejection. Make sure to heed the comments made by the reviewers and the editor before resubmitting the article to the next venue.
6) It is useful to experience an internship in an industrial research lab or a foreign academic lab. The former may help secure an interview at a later time and exposes you to the research in an industrial setting. A research methodology driven by product development leads to unique challenges, requirements, and expectations. Doing some work in a different academic setting provides an opportunity to learn from a different research philosophy, methodology, and culture. Internships provide an opportunity for personal networking that are different from those gained during conferences and workshops.

Internships provide an opportunity for personal networking that are different from those gained during conferences and workshops.

Conclusion

It is worth spending some time in finalizing your initial problem because it lays down the foundation of the entire dissertation. Time spent on it may depend on the maturity of the area, effort spent, prior experience in research (on the same or other areas), and a bit of luck. Starting early is always useful. Current work should always be finished to the best of one's ability before moving on to the next problem.

Acknowledgment

I would like to thank my Ph.D. advisor Prof. Joy Kuri and the anonymous reviewers for their comments on this article.

Read more about it

• J. Eisner. (2014, Mar. 19). How to find research problems. [Online]. Available: http://www.cs.jhu.edu/~jason/advice/how-to-find-research-problems.html

• D. J. Boudah. (2011). Identifying a research problem and question, and searching relevant literature. *Conducting Educational Research: Guide to Completing a Major Project*. Thousand Oaks, CA: SAGE Publications. [Online]. Available: http://www.sagepub.in/upm-data/36330_Chapter2.pdf

• N. Walliman. (2001). Research and research problem. *Your Research Project: A Step-by-Step Guide for the First-Time Researcher*. Thousand Oaks, CA: SAGE Publications. [Online]. Available: http://www.sagepub.com/upm-data/40600_9781849204620.pdf

• S. Roger and M. Cobos, "Developing your electrical engineering degree thesis," *IEEE Potentials*, vol. 28, no. 4, pp. 12–16, July–Aug. 2009.

• S. Keshav. (2007, July). How to read a paper. *ACM Comput. Commun. Rev.* [Online]. Available: http://blizzard.cs.uwaterloo.ca/keshav/home/Papers/data/07/paper-reading.pdf

• V. Grover and M. K. Malhotra. (2003, Dec.–Jan.). Interaction between a doctoral student and advisor: Making it work. *Decision Line*. [Online]. Available: http://www.decisionsciences.org/DecisionLine/Vol34/34_1/34_1phd.pdf

About the author

Kundan Kandhway (kundan@dese.iisc.ernet.in) is a Ph.D. candidate at the Department of Electronic Systems Engineering at the Indian Institute of Science. He earned an M.A.Sc. degree from the University of British Columbia and a B.Tech. degree from the Indian Institute of Technology Guwahati. He is currently working on social networks.