

Doing a Systems PhD

Steve Hand

**EuroSys Doctoral Workshop
Salzburg, 10th April 2011**

**University of Cambridge
Computer Laboratory**

Systems Research is...

- Work in OS, file-systems, databases, networking, language run-times, system security, ...
 - Not a 'hard' science
 - No ground truth to be discovered
 - Get to create the universe!
 - Things can be "sort of" right
 - Absolutes are rare
 - Key skill: critical thinking
-

Critical Thinking - Reading

- Read a paper and consider:
 - Do I like it? Hate it? (opinion)
 - What problem is it trying to solve?
 - How does their approach differ from previous ones?
 - (how much previous work do I know about – read it! (reference chaining))
 - Does it work?
 - What could be improved?
-

Critical Thinking - Writing

- Consider a paper (or your thesis) as an *argument*
 - What is the problem?
 - If not well known, *why* is it a problem?
 - Why are all previous approaches *insufficient* (broken / wrong / stupid)?
 - What is your approach?
 - how does it work?
 - how well does it work?
 - how does it improve on previous attempts?
-

PhD Outcomes

- Drop out, do something else
 - Finish, do something else
 - Finish, join/start a start-up
 - Actual PhD work may be important
 - Finish -> Academic Research
 - Finish -> (good) Industrial Research
 - Last two: academic track record, references, publications important
-


Doing a PhD: What's involved?

- Choose an advisor and a research area
 - Write a research proposal, apply & get accepted
 - Then do some or all of:
 - Build some stuff (e.g. new multimedia FS)
 - Measure some stuff (e.g. power use on laptop)
 - Evaluate some stuff (e.g. performance of new web server scheduling algorithm)
 - Analyze some stuff (e.g. queuing theoretic model of the Internet)
 - Write a dissertation
 - Submit and get viva'd – scary!
-

Choosing an Advisor

- How much time do they have?
 - Is their technical background a good fit?
 - How do they like to work?
 - How do they like their students to work?
 - Ask their students!
 - Do they expect you to work on 'their' topic (RA for N years)?
 - Or do they expect you to come up with your own topic?
-

$B > \{ A \text{ or } C \} \gg D$



A=Prima donna

B=Savior

C=Pushover

D=Nightmare

Choosing a Topic

- Don't have to fix on this on day one!
 - Do some stuff, work with others, write some papers, and *then* choose
 - 6 – 18 months is fine
 - How to choose something?
 - Read a paper you hate – decide to fix it
 - Read a paper you love – join effort
 - Discover a problem – aim to solve it
 - “Scrabble” – invent something
-

Risks – Underestimation

□ Underestimate PhD

- “Bob got a PhD for writing a web server – I’ll write a web server too!”
- (Fail to realize Bob’s PhD was actually about techniques in zero-copy I/O)
- Read nothing but write lots of code
- Zero publications, zero justification

□ Result:

- Never submit (“don’t like writing”); or
 - Submit, sit back confidently, and get blown to shreds in PhD exam
-

Risks - Overestimation

- PhD ~ = Nobel prize in Physics
 - All existing problems are simple
 - All existing solutions are stupid / trivial
 - Read everything and dismiss it
 - Invent new words or language or logic or ontology just to *express* the problem
 - Cannot explain problem (or solution) to mere mortals – stupid them!
 - Zero publications (“Not quite finished yet”)
 - Result:
 - Never submit (“No one understands me”)
 - Submit, sit back smugly, get ripped apart in PhD exam, and blame the idiocy of the world.
-

Risks – Awaiting Orders

- Treat PhD like high school
 - Ask supervisor what to do
 - Don't understand it
 - Don't agree with it
 - But do it anyway
 - Finish task and wait for next task
 - Play WoW or do consultancy in meantime
 - Write nothing, read nothing
 - Repeat until PhD deadline and realize
 - You don't have a thesis; or
 - You have a thesis you don't understand
-

Risks – Isolation

- Come up with plausible topic and reasonable approach / solution
 - Start work, but don't make as much progress as you'd like
 - Oh no! Other people seem to be doing fine!
 - Stop coming into department
 - Avoid supervisor
 - Avoid peers
 - Stop showering
 - Heroin
 - Death
-


Better Strategies

- Work with others from day one
 - From 1 year in, aim to have a current “draft” of your PhD in your head
 - Do *something* every month
 - Read + critique a bunch of papers
 - Write some code
 - Do some measurements
 - Write down results, designs, ideas, ...
 - Dual-task if at all possible: left brain / right brain parallelism
-

Aside: Breadth is Important!

- PhD itself typically in-depth:
 - You become the expert at something
 - But your time on the PhD program should cover more than this:
 - Work with others (& in other areas)
 - Internships particularly valuable
 - Post doctoral jobs typically favor a broader outlook (too narrow == bad!)
-

Managing your Supervisor...



Managing your Supervisor (1)

- Extreme #1: “The Gauleiter”
 - He or she has an idea and/or plan
 - Your job is to carry this out
 - Pros:
 - Should be clear what to do
 - Can make progress from day 1
 - Cons:
 - Little opportunity for creative thought
 - May feel lack of control or ownership
 - Potential outcomes: frustration, or apathy
 - Strategy: read, think & *argue*
-

Managing your Supervisor (2)

- Extreme #2: "The Don"
 - Vaguely interested in everything
 - Expects you to come up with an idea, and then go off and do something good (but may not mention this)
 - Pros:
 - Lots of flexibility and options
 - Lots of positive feedback from supervisor
 - Cons:
 - Easy to get stuck, or lost
 - Feedback may be vague or esoteric (low usefulness)
 - Strategy:
 - Attempt to engage him/her concretely in your work
 - Impress with your own erudition / intelligence
-

Managing your Supervisor (3)

- Most supervisors not at extremes, but will have different pros and cons
 - Key point: it is *your* responsibility to make your supervisor work for you
 - You're the one who wants to get the PhD
 - General Strategies:
 - Have relatively frequent meetings
 - Aim for concrete deliverables (e.g. whiteboard design, or draft paper)
 - Educate: be[come] the expert on your topic
 - Learn to argue/discuss/explain
-

Aim to Publish

- Writing a thesis with several publications under your belt is a lot easier than doing it from scratch!
 - Plus a good way to exercise your possibly atrophied 'writing muscles'
 - Work with others:
 - E.g. three 3-author papers a year for the same price as one single author paper
 - Get feedback:
 - Reviewers are often smart and dedicated
 - And even if they're not, they're representative of the research community (i.e. your examiners)
-

How to Publish (1)

- Start by writing down *something*
 - Hard to publish if don't have a paper :-)
 - Starting point usually either "stuff I've done" or "thing I believe"
 - "Stuff I've done" – first write a tech report which just describes it
 - Add `blank' related work section
 - Retro-fit argument of some sort
 - Give to peers / supervisor / others, get feedback, modify, repeat, ...
 - Submit to appropriate workshop / conference (with proximate deadline)
-

How to Publish (2)

- Or start with an idea / belief
 - Write out skeleton argument
 - Critique related work
 - Work out what you need to actually *do* to back up your argument, and then
 - Sketch out solution in paper, run past peers / supervisor / others, submit position paper
 - And/or start to do actual work
 - Add details / results etc as you go
-

In Both Cases: Use Others

- Come up with an outline argument
 - Run past peers / supervisor / others
 - Objectively consider feedback
 - Tweak / vastly rework argument
 - Repeat until fixpoint
 - Write a position paper or short paper
 - Run past peers / supervisor / others
 - Objectively consider feedback
 - Tweak / vastly rework paper
 - Repeat until deadline
-

Reasons for Paper Rejection

1. Paper not clearly written (at a word / sentence / paragraph level)
 2. Paper not clearly written (at a structural / argument level)
 3. Paper clearly written, but:
 1. argument is weak / false; or
 2. solution is obvious / incorrect; or
 3. experiments (or analysis) are poor
 4. PC are biased idiots
-

Writing up

- Need to write a *dissertation* which supports your PhD *thesis*
 - Typically 30-60K words:
 - Longest document you'll have ever written
 - Hard to ensure a single "story" throughout
 - Core ("meat") usually 1—3 chapters
 - E.g. design, implementation, eval
 - E.g. technique1, technique2, technique3
 - Produce drafts and get frequent feedback
 - Expect 6-12 months just for write-up!
 - Day after submission: best day of your life
-

Summary

- Systems Research is (or can be) fun
 - Can have “real world” impact, or make a fundamental contribution (or both?)
 - Keys to success are
 - Engage critical thinking – read a lot
 - Make everything an argument – use your colleagues and supervisor for feedback
 - Be proactive – educate your supervisor
 - Publish (or at least submit) papers
-

Questions?



- or why not just chat to me over a beer?
-