Identifying Research Problems

Peter Stuckey
IMDEA Software Institute and University of Melbourne

Workshop at UPM

Introductions

- Education
  - B.Sc and Ph.D from Monash University
- Postdoc
  - Postdoc at IBM T.J. Watson Research Center
- Currently
  - Professor of Computer Science at The University of Melbourne
- Research Interests
  - logic programming, constraint programming, discrete optimization
  - modelling languages, program analysis, bioinformatics (protein structure)

MiniZinc

- A high level modelling language for discrete optimization
- Open source (minizinc.org)
- Defacto standard for CP modelling
- Coursera course on MiniZinc
  - Modeling Discrete Optimization

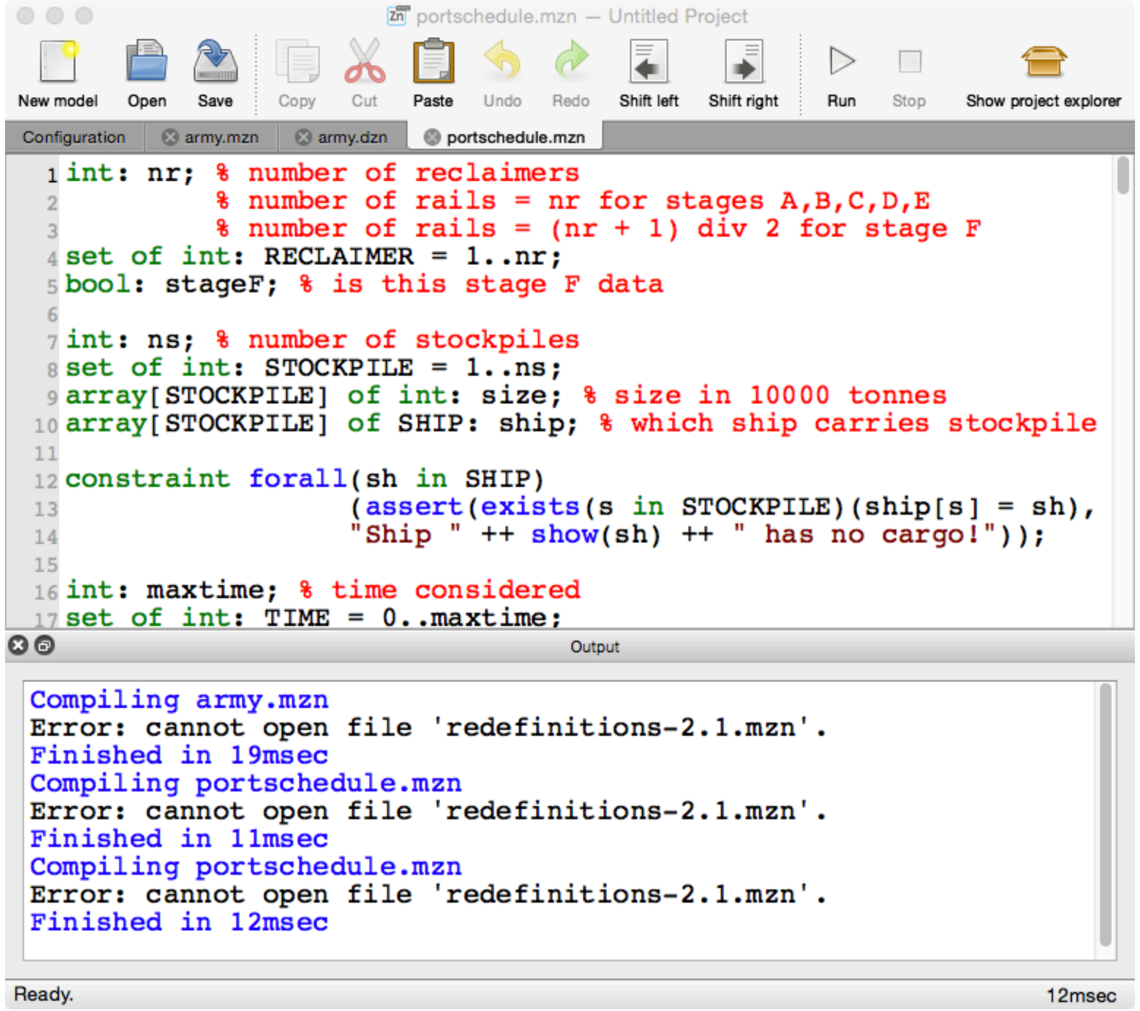

Melbourne

- The worlds most liveable city

Introduce Yourself (Who Are You?)

- Name
- Year
- Topic
- Research Experience
- What do you hope to get out of this workshop

The hardest part of research: asking questions

- If you are smart enough to get into graduate school, you are probably good at answering questions
- That's what college trained you to learn
- Graduate school teaches you how to ask questions

- Once you ask the right question, the answer is often obvious
- The question you ask implies a framing of the problem

The characteristics of great research

- Recall a good research paper/talk/result
- What was remarkable about it?

The characteristics of great research

Great research needs:

- idea ⇦ today's focus
- execution
- presentation

- We will examine how to find a research idea

Types of Research Problems

- Foundational/Theoretical
  - New logic
  - Formalism
  - Better algorithm
- Experimental/Systems
  - Better system to do something (e.g., cloud computing)
  - New attack on a system (e.g, IoT)
  - …

Well-presented research convinces you:

about
the
problem

1. The problem is interesting
   - effect on human welfare, corporate profits, defense
   - *evidence:* other people have worked on it
2. The problem is hard
   - other smart people have tried it
   - obvious approaches do not work
   - problem persists over time
3. The authors have solved the problem
   - technical details
   - evaluation, such as proofs or experiments

Sources of research ideas

- Always be working on *something* (even if it won't be your thesis topic)
  - Think about problems, difficulties, or irritations
  - You might get sidetracked onto something more important
- Learn about other research: attend talks, read papers, take classes
  - future work they identified
  - connections they did not make
  - assumptions
  - other criticisms
- Your advisor
  - OK for starter projects, master's projects
  - by the end the PhD should be based on your ideas

Read, Attend, Discuss

- Read widely in your area
  - Your supervisor can start you off, but you should put in place mechanisms to alert you to new relevant results
- Attend conferences, departmental seminars, postgraduate seminars
  - Even if they seem to have nothing to do with your research
- Discuss with other postgraduates, supervisor, other academics
  - passageway conversations can lead to very interesting papers
- Even things that appear irrelevant may be useful
  - You don’t know what you need to know
- Referee papers
  - ask your supervisor for some refereeing jobs

Become an Active Reader/Listener

- When reading technical material/examining software/attending talk
- CANONICAL QUESTIONS
  - From where did the author seem to draw the ideas?
  - What exactly was accomplished by this piece of work?
  - How does it seem to relate to other work in the field?
  - What would be the reasonable next step to build upon this work?
  - What ideas from related fields might be brought to bear upon this subject?
- A written log of technical reading and listening can be useful

Your Advisor

- A direct source of research ideas
  - almost always the first problem you will work on will be suggested by your advisor
- An experienced filter of possible research ideas
  - discuss new research ideas/directions with your advisor
- Judging feasibility + importance of new ideas needs experience
  - this is where you advisor is helpful

- Learning to communicate well with your advisor is key

Collaboration

- Two brains on one problem

  \>>

  Two x (one brain on one problem)
- Strategies for collaboration
  - talk to other people,
  - read their drafts,
  - discuss their ideas without hoping to be an author
  - ask others for such help

Brainstorming

- An approach to generating ideas using a group
- Principles
  - Defer judgment
  - Reach for quantity
- 4 general rules
  - Go for quantity, “quantity leads to quality”
  - Withhold criticism, criticism inhibits idea creation, an accepting space
  - Welcome wild ideas, suspend assumptions, ignore the past
  - Combine and improve ideas, the power of the group

Brainstorming Exercise

- Think about problems that you encounter in day-to-day programming.
- Write a paragraph about the problem on a piece of paper
- Pass the paper left
  - write one sentence about the problem
    - things that might be related
    - some method that might be useful
    - some system/paper/person who has already solved it, solved something close
- Try to write something
- Continue until the paper reaches its original author

Think outside the box: A good research idea often "cheats"

- If you take the same approach as other people,
  you are likely to be no more successful than they are
- *Cut the Gordian knot:* take a different approach
- Example:
  - Artificial intelligence vs. brute force (chess, speech, etc.)

Bad ideas are essential

- Most ideas are bad ideas
- If you have lots of ideas, some will be good
- It's better to have bad ideas than no ideas
  - I can teach people how to filter
  - It's harder to teach courage and imagination
- Filtering is initially your advisor's role
  - You will get better with time,
    and spend less time evaluating and discarding bad ideas
  - Have faith in your own judgment, too
  - Your advisor is not infallible.
- Generate lots of ideas before filtering
  - Breadth-first, not depth-first, approach to the problem
- Don't look for a "home run"

Attack the Problem Now!

- Dont wait until you’ve read all the background material
- Develop a quick solution if you can
- Then check it hasnt been done before
- Get to know
  - CITESEER: http://citeseerx.ist.psu.edu/
  - DBLP: http://www.informatik.uni-trier.de/~ley/db/
  - Google Scholar: http://scholar.google.com.au/
- Use authors web sites

Heilmeier's Catechism
(questions to be asked about any research proposal)

1. What are you trying to do? Articulate your objectives using absolutely no jargon.
2. How is it done today, and what are the limits of current practice?
3. What's new in your approach and why do you think it will be successful?
4. Who cares?
5. If you're successful, what difference will it make?
6. What are the risks and the payoffs?
7. How much will it cost?
8. How long will it take?
9. What are the midterm and final "exams" to check for success?

Caveats

- More applicable to experimental research

- Foundational/basic research
  - Better algorithm to do something

  - Solve an open problem
    - Fully homomorphic encryption
    - P = NP (perhaps not)

  - Brand new concept or paradigm shift
    - Internet, ….

Automation

- Find a task that is currently done manually, or that is time-consuming or difficult
- Thoughtfully examine the manual process
- What reasoning does the person do?
- Can you automate that process?

Contrarianism

- Turn the conventional wisdom on its head
  - Do the opposite of what everyone else expects
  - You must do this thoughtfully: you have to justify why, and why it matters
- Example: sequentializing parallel programs
  - Solves the same problem, and is better in some ways

Notice complementary approaches

- Given two approaches to a problem:
  wherever one is used,
  try the other one

- Example: static and dynamic analysis
  - Static: sound, conservative
  - Dynamic: precise, does not generalized

⇒ combine ideas from both

- Example: semantic vs. syntactic analysis

Combine existing techniques

- Apply one after the other
- Create a hybrid analysis

Cross-fertilization

- Take an approach from one domain, and apply it to another
- Examples: Machine learning, statistics, biology, …

Fill in gaps

- Review your previous research
- Are there themes or commonalities?
- Find ways to fill in the gaps

- This is particularly useful for creating a PhD thesis out of a sequence of successful projects.

Generalizing

- How does this generalize? What are the broader implications?

Problem vs. technique

- A research project requires:
  - A goal or problem
  - A technical approach
- Advantages of starting with a goal
  - More likely to find a technical approach
  - More likely to be novel and relevant
- Advantages of starting with a technique
  - Quick to complete: you probably already have it implemented
  - Fits well with your other work

What is your secret weapon?

Use your unique set of skills and experiences

Why are you the right person to work on this?

- You have taken a particular combination of classes
  - *Example:* Google founders (algorithms/machine learning)
- You have a friend with a tool/technique/knowledge
  - Don’t be afraid to collaborate
- You thought of it first (new approach)
- You are willing to work harder

Passion for your work

- Your intellectual curiosity about a topic
  is necessary but not sufficient.
- Necessary:
  - Your love for the topic will affect your motivation (and possibly ability)
  - It helps if the topic was your idea
- Not sufficient:
  - It has to matter to other people as well

Creative Thinking (not specific to research)

- YOUR BRAIN DOES BETTER CREATIVE WORK WHEN YOU'RE TIRED
- EXERCISE CAN IMPROVE YOUR CREATIVITY
- AMBIENT NOISE LEVELS ARE BEST FOR CREATIVITY
- NOTHING IS ORIGINAL: CREATIVITY IS ALL ABOUT MAKING CONNECTIONS
- TRAVELING ABROAD MIGHT IMPROVE YOUR CREATIVE THINKING
- DIM LIGHTING MAKES US FEEL MORE FREE
- CONSTRAINTS CAN BE BENEFICIAL TO CREATIVE WORK
- A SEPARATE, MESSY DESK CAN IMPROVE YOUR CREATIVITY

Creative Thinking

- Brainstorming:
  - generating lots of ideas from a group of people
- Engaging your Right Brain
  - make something out of cardboard, balloons or Lego
  - mind mapping: a diagram of words interconnected by lines
- Go with the Flow
  - when you are being creative, push further, explore more ideas

Exercise

- Lets examine your research idea
- Does it convince us that

1. The problem is interesting
    - effect on human welfare, corporate profits, defense
    - *evidence:* other people have worked on it
2. The problem is hard
    - other smart people have tried it
    - obvious approaches do not work
    - problem persists over time

    - pstuckey@unimelb.edu.au

Exercise

- Lets examine your research idea
- Should it be
  - more defined, more focussed?
  - more general, less specific?
- Does it meet the Hellmeier criteria?