

Changing The World



Robert O'Callahan
Mozilla Corporation

About Me

- Academic career
 - PhD (CMU)
 - IBM Research
 - OOPSLA, POPL program committees
- Mozilla career
 - Contributor
 - Developer
 - Manager

Responsibility

- Who is talented?
- What is the original meaning of “talent”?

Matthew 25:14-30

“Again, it will be like a man going on a journey, who called his servants and entrusted his property to them. To one he gave five talents of money, to another two talents, and to another one talent, each according to his ability. Then he went on his journey. The man who had received the five talents went at once and put his money to work and gained five more. So also, the one with the two talents gained two more. But the man who had received the one talent went off, dug a hole in the ground and hid his master's money.

“After a long time the master of those servants returned and settled accounts with them. The man who had received the five talents brought the other five. 'Master,' he said, 'you entrusted me with five talents. See, I have gained five more.'

“His master replied, 'Well done, good and faithful servant! You have been faithful with a few things; I will put you in charge of many things. Come and share your master's happiness!'

“The man with the two talents also came. 'Master,' he said, 'you entrusted me with two talents; see, I have gained two more.'

“His master replied, 'Well done, good and faithful servant! You have been faithful with a few things; I will put you in charge of many things. Come and share your master's happiness!'

“Then the man who had received the one talent came. 'Master,' he said, 'I knew that you are a hard man, harvesting where you have not sown and gathering where you have not scattered seed. So I was afraid and went out and hid your talent in the ground. See, here is what belongs to you.'

“His master replied, 'You wicked, lazy servant! So you knew that I harvest where I have not sown and gather where I have not scattered seed? Well then, you should have put my money on deposit with the bankers, so that when I returned I would have received it back with interest.'

“Take the talent from him and give it to the one who has the ten talents. For everyone who has will be given more, and he will have an abundance. Whoever does not have, even what he has will be taken from him. And throw that worthless servant outside, into the darkness, where there will be weeping and gnashing of teeth.'”

- From the beginning, God-given talent has carried responsibility
- Not just to refrain from evil, but to make the best use of it

The Magic Of Computer Science

- Make things happen at zero marginal cost
- Save one second per user per day
 - One life per 84,600 users
- 300M users → 3,000 lives

Things That Make Me Cry

Search For Impact

- Who is going to benefit from your work?
- How, and by how much?

How To Have Impact

Two Ways To Have Impact

- Build something to be used in practice
- Create and transmit knowledge that will help someone else build something practical
- Let's call the former “development” and the latter “research”
- Not mutually exclusive!
- Not based on location!

Why Do Research?

- Development has all kinds of constraints
- By dropping constraints, you can see more clearly and make faster progress
- Characterizing which constraints are dropped is a good way to understand “development” vs “applied research” vs “pure research”
- Example: programming languages

Research Vs Development

- Development has a simpler path to impact
- Research offers more crisp intellectual problems
- If you have to publish or perish, research lets you publish more

Problem Selection

Research Problems

- Many people think solving problems is the hard part of research
- Most research I see won't have impact even if it's completely successful
- Goal: a crisp, intellectually satisfying problem which also has immediate impact
- The space of problems is huge
- Desirable problems do exist!

A Cautionary Tale

- “Program slicing”
 - Extract the subset of a program P that preserves the value of a variable v
- First identified by Mark Weiser in 1984
- Hundreds of papers followed since
- Almost completely pointless

Solution-Driven Approach

- Identify skill or technology then find problems that fit the solution --- “applications”
- Examples from my career
 - Apply Hindley-Milner type inference to C and Java programs
 - Gather run-time statistics from Java programs using bytecode instrumentation
- An uphill struggle to have impact!

Problem-Driven Approach

- Identify important problems and solve them
- Example: debugging
 - Look at what programmers do, and figure out how to make it faster/cheaper/go away
- No preconceived ideas about solutions
- Use Wizard-of-Oz techniques to prototype solutions before you build them
- You may have to learn something new

Finding Problems

- The world is full of them
- Talk to people about what they do
- For software problems, look to open-source projects

Publishing

Aim High

- Hard to have impact outside the top conferences --- readership issues
- Often papers at top conferences repeat work in less-known conferences
 - The top conference papers get the credit

Tips

- Publishing in top venues is not that hard if you've selected good problems
- Good problems with mediocre solutions are often accepted
 - And widely cited!
- Read lots of proceedings to understand what is accepted (and how bad it is)
- Play the game
 - POPL (ahem)

Tips

- Reviews are quite random, don't be discouraged
- Clearly identify the problem to be solved
- Choose problems amenable to compelling evaluation

Negative Results

- Positive results are unsurprising
- Negative results are often surprising
- Negative results are more useful, if the reason is interesting
- You can't publish negative results
- This is bad :-)

Research Environments

Corporate Research

- Most corporate research benefits society, but not the company
- Companies do basic research out of vanity in cash-rich times
- Think of it as philanthropy

Integrating Research With Development

- Long-lived corporate labs survive by doing some development
- Splitting development makes little sense
- Prefer Google model
 - “Tech transfer” from labs becomes non-issue
- Start with research and add constraints until you have a product
- Use the same team (example: Mozilla JS)

Academic Research In NZ

- How can we have more great academic research done in New Zealand?
 - Without spending more money?

Gravitational Pull

- Top people want to live in NZ
- But they want to work with top peers
- A concentration of top people will attract more top people (and students, and money)
- This is why top universities stay at the top (see CMU)

A Modest Proposal

- Put our best people and resources at one institution
- It will attract more top people
- This is not a zero-sum game

Education

- Most students seem to want vocational skills, not computer science
- Some employers need students with computer science
- An elite institution could collect the students keen on CS
 - And leave the vocational training to other places

Impact Via Development

Development Is Underrated

- Research problems are not harder
 - Dropping constraints often makes things easier
- Development is better funded
- More direct impact
- Non-boring development jobs exist!

Development Environments

- Megacorporations are horribly inefficient
 - Go for specific projects
- Otherwise go for small organizations
- Consider contributing to open-source projects
 - Disseminate your work
 - Immediate impact
 - Learn and guide research

Tips

- Choose projects lower on the software stack
- Instead of building systems out of parts, build the parts
 - The design of the parts constrains what others can build
 - Affect more users
- Choose projects with more users

Don't Be Evil

- Don't have negative impact
- Don't work somewhere you have to apologize for

Conclusions

- Strive to have positive impact on the world using computer science
- You can have impact in development or doing research
- We could have more great research in New Zealand by concentrating it in one place