

You Must Be Joking... Cares About The Endangered Species

Michalis Faloutsos
UC Riverside
Dpt. of CS
Riverside, CA
michalis@cs.ucr.edu

ABSTRACT

For those of us in academia, tenure is great. Unless you don't have it, in which case it pretty much sucks. In fact, it goes beyond sucking: it kills. Mainly your personal life, sometimes your spirit. I claim that assistant professors are an endangered species. Are we heading towards an environmental disaster that is going to destabilize the academic ecosystem or did I just have too much to drink last night? Either way, I think I deserve the Noble prize for Peace.

1. INTRODUCTION

The main thesis of this article is that being an assistant professor is not fun. In fact, it can be pretty bad.

Let me start with a disclaimer. I don't want to give the wrong impression that I have personally suffered, or that this is a comment of bitterness from what happens at my school. First, I have personally felt cared for and appreciated from the beginning that I joined my department. Second, if anything, our department is very high on friendliness and collegiality. People genuinely want others to do well, and they often try to help. I know many other CS departments that are also like that, but also know some that are not. I also know people that overly strict.

In case you are wondering why I am writing this, I really feel bad for assistant professors. This has reached the point, that if a PhD student would ask me whether they should go to academia, I would not recommend it without some reservation and some warning.

2. THE GAME HAS CHANGED

To paraphrase a little known early song by Bruce Springsteen: *"It is so hard to be an assistant in the city"*.

The shifting tides of research. Is it just me or are research topics moving fast? From, hot, to cold, going through reincarnations, to hot again under different jargon. As proof, let me say that we see a proliferation of *hot*-something conferences¹. Next thing you know, we will have an ACM list of "Who's hot". This clearly reveals the urgent need to find ourselves on the forefront of hotness, before other people claim it first and jump on it, thus, turning it into lukewarm instantly.

Unofficial statistics from my aging memory cells: research topics have a two year time span. In fact in almost a year

¹For more on this topic, see Jon's article on Jan 2008 CCR issue

a hot area becomes crowded, by the third year it becomes a joke or is received by reviewers as "Yet another paper on Blah".

Remember sensor nets? Sensors everywhere, the revolution, the "it" topic. Two years later it is the applications and the army deployments that is sustaining the area. I recently spoke with some sensor people that are moving on to other areas. Don't get me wrong, I am not saying that the sensor excitement was for nothing: good technologies were developed and commercialized and all. Some open topics still remain. However, I am seeing this from the eyes of an assistant professor, who is trying to find a topic to work on for the next 5 years, so that they can claim an area of expertise and ultimately achieve name recognition. Two years is too short of a time.

Can anyone predict, what will be a topic that will last? There may be some visionaries² and clairvoyants, but I am not sure that we can demand this from young professors.

There are some areas that are niches, and can continue to be active and less crowded for a while. But then again, if an area is too narrow, that's a criticism: good work but narrow area, does the community at large know her/his name? One can't win.

You may say: "I also started as assistant, and I managed to make a name for myself". Maybe, but don't forget that, in the past, people would invent anything and it was a new thing.

— I was bored last night so I sat down and defined boolean logic. A statement must be either true or false. It does not get more basic than that. And I don't want to hear anything about Gödel.

— Hey look at me, I invented bubblesort. I thought, given a hundred numbers, heck, someone would want to sort them.

— Good for you. I invented Finite State Machines on Sunday. It is great. I love drawing boxes. I somehow always knew they would be good for something.

— Hey guys, keep it quiet, I am trying to invent the Internet, here. Having just invented packets, it shouldn't be that difficult.

— Seven layer dip anyone?

— Layers you said? Hmm... Seven also seems like a good number. I think I could use this idea...

And so it goes.

Publishing. Publishing papers has become more vicious with time. I still have bitemarks near my jugular from the

²"I was a peripheral visionary. I could see the future, but only way off to the side." Steven Wright. The man is a comedic genius. Tube him.

most recent reviews. I am not going to talk about reviewing, since I have bitched about this more than once. However, see this from the point of view of an assistant, who is trying to reconcile two conflicting goals: (a) publishing quickly, before someone else scoops the work, (b) publishing in highly selective places, and (c) having a large numbers of publications. (ok, fine the conflicting goals turn out to be three, sue me.) I personally find that selecting which conferences to send your work to has become an optimization problem with mutually exclusive deadlines, different weights of importance, different success ratios, different page numbers. I wonder how come none has thought of some kind of dynamic programming or a greedy approximation solution that will maximize the total academic points within five years.

It is clear that in the past things were easier. I have no actual proof, but I believe that in the early days of Sigcomm, Dave Clark would submit his CV by accident and the TPC would have a lively discussion on the pros and cons of the document. Mind you, that even now, Clark may send his lunch discussion with Karen Sollins on a napkin, and it may be published. What's scarier is that this may very well turn out to be the most thought provoking submission. Am I contradicting myself here? Maybe. One thing is certain: it is a crazy world out there.

Furthermore, the expectations from a publication was much lower in the past. Conferences where about ideas. Now, a conference paper has to practically have reached the point of commercially viable alpha-tested product. I have seen papers with two implementations on different platforms. I was part of a paper where the new protocol was tried on two different testbeds. Yes, for well established groups this is a major hassle, but it is doable. For young profs, it is a killer.

Funding and assistant professors: Correct me if I am wrong but since approximately 1999³ success ratio for NSF proposals went from a happy one out of four to an abysmal one out of ten. Even for senior people, this hurts badly. Imagine now that you don't yet have an established group, connections with other researchers, connections with industry, and you have the research breadth of a single PhD in your portfolio.

The CAREER grant is a good effort to help assistants, but I am afraid that it has side-effects. I claim that it increases their stress level tremendously. CAREER grant has somehow stopped being thought of as "easier" NSF, where only assistants compete. It has become a badge of honor. And from there, it has spun out of control to the wide belief: if you don't have it, there is obviously something wrong with you.

Let me prove the madness of this situation. Consider an assistant, let's call him R, that gets a regular NSF grant as a single principal investigator. How would he compare with assistant C who gets his CAREER? Assume that for both, the amount of money is the same and everything else in their funding record is the same. Shouldn't assistant R be considered to have a stronger funding record? R competed against everybody: senior people, groups of people, international collaborations, illegal immigrants, extra-terrestrials, everybody. Logically, it is clear: R won a harder battle.

But, ask your local chair of the Department. Assistant C will get more kudos and congratulations, and be highlighted in the departmental webpage more prominently. Does it

³Just before the turn of the millenium. Is it pure coincidence?

make sense? This unfortunately happens across the board, and it is not based on a single department.

The change from 1/4 to 1/10 success ratio in NSF proposals changes fundamentally the nature of the game. With 1/4, if you write 4 proposals a year, you expect to do ok. With one out of ten, you have to write 10 proposals a year to ensure the same level of funding. In other words, you should be writing one proposal almost a month. Think about it for a moment. You may say: but the quality proposals will surface anyway, yada yada yada. I hear you, although it pains me to do so. Quality definitely plays a role, but only to some extent. Ok, fine, let us say that only half of the proposals are good, in the case of 1/10 success ratio, that still leaves a choice of 1/5 among good proposals. And oh, wait a minute, you can only submit 2 proposals each year in your area in NSF. Do the math: what are the chances of a young prof getting NSF funding in the first two years?

You may say: assistants should pair up with senior people. Which will then create the comment, "This proposal does not count, his co-PI was Prof. Bigname". Plus getting access to prof. Bigname is not trivial, and one's department may not have Prof. Bigname, while the available Prof. Mediumname may be a jerk.

But so-n-so is doing great. Yes, this is the comment I was expecting. The old proof by example. Sure, computer scientists can be buddies with Bono and Al Gore, look at Bill Gates' last day at work video. Right. For every so-n-so that is doing great there is a large number of assistants that work just as hard but either they are not quite as brilliant, or lucky⁴ like so-n-so. If you think about it, it is a zero-sum game⁵: how many revolutionary high impact papers can you have in a year? If you think about it, it is a bit like saying how many top-10 most influential papers of the year can you have in a year?

Clearly, many of the above hardships apply to all of us in academia. The point is that it is that much harder for young professors who lack the network of connections, the momentum of an ongoing program, a breadth of research activities and ultimate experience.

3. WHAT CAN YOU DO?

If you are not part of the solution, you are part of the problem. I am not sure if this is true, but I always wanted to say this⁶.

Seriously, now, if you are tenured, you are definitely part of the problem, unless you are actively doing something to help. If you are assistant, you are the problem.

It boils down to this: we need to cut assistants some slack. Enough with pointless comparisons to what we did five, ten, twenty years ago. It is a new game. Plus, I don't understand

⁴If you think that luck has nothing to do with, you can stop reading this article, as you most likely live in a different world than this article describes. I am not saying imaginary, just different.

⁵Wikipedia: "In game theory and economic theory, zero-sum describes a situation in which a participant's gain or loss is exactly balanced by the losses or gains of the other participant(s)".

⁶Once in a conference, I got to go in the second cab of a large group, where someone in the first cab knew where we were going. When the cabdriver asked where we were going, I used the "Follow that cab" line, meaning it. It was great. One less thing in my to-do-before-I-die list.

it, if people in the industry hire⁷ someone, they try to make them succeed. It is partly the task of the hiring manager and rest of the team to integrate the new person. They don't say: "It is your honor to be here, show us what you can do".

In more detail, here are some things that I would recommend.

a. Be nice. Sometimes a good word of encouragement can go a long way. In contrast, if an assistant tells you of their difficulties, and you find it as a good opportunity to show off or patronize, you are not helping.

b. Be mindful⁸. Anticipate problems, be proactive. Try to see if they are lost somewhere, or if they are spending too much effort on something that is not as important.

c. Remember the three I's. I am not talking about the Internet Indirection Infrastructure, the other ones: introduce, include, invite. Include them in a proposal. Ask them co-supervise a good student of yours. Introduce them to an important colleague from another university or an industry contact.

d. Don't drown them in busy work. Eliminate their administrative duties, don't dump on them things that you would have to do, if they were not there. The approach "they are young, they should pay their dues" is a rationalization. Once Kevin Jeffay said in a workshop for new faculty, that every year he finds an administrative duty that was given to an assistant and takes it over.

Jennifer Rexford correctly pointed out after reading my first draft: "folks can do these things not only for assistant profs at their own institutions, but also those at other schools." I am glad Jen exists: I could not agree more. It is like sitting next to the new kid at school just because you don't want them to eat lunch alone.

Final disclaimer: I don't claim that I am doing all these things. I want to believe that I have done some of those things here and there, but it has not been enough by far. So, if you are asking, yes, I am most likely part of the problem too. We can all do better.

In conclusion, I would like to officially launch a campaign: "Save The Assistants."

⁷There the interviewing process is often ridiculous, see an earlier column.

⁸You have to admit I chose an awesome word here. Merriam-Webster defines it as: 1 : bearing in mind, 2 : inclined to be aware. However, in Buddhism it takes a different meaning: "Mindfulness is a technique in which a person becomes intentionally aware of their thoughts and actions in the present moment, non-judgmentally." Initially it looks like a good idea, but I don't know, it sounds like too much work.