The Tragey of
The Computing-Research Common

Moshe Y. Vardi

Rice University
CACM EiC

"The Tragedy of the Commons"

Garrett Hardin, Science, 1968: A dilemma arising from the situation in which multiple individuals, acting independently, and solely and rationally consulting their own self-interest, will ultimately deplete a shared limited resource even when it is clear that it is not in anyone's long-term interest for this to happen.

Hardin's point: human population growth and the use of the Earth's natural resources

- Polution
- Fishing

The Parable

William Forster Lloyd, 1833:

Herders share a common pasture, on which they are entitled to let their cows graze. It is in each herder's interest to put the next (and succeeding) cows he acquires onto the pasture, even if the carrying capacity of the pasture is exceeded and it is damaged for all as a result.

The herder receives all of the benefits from an additional cow, while the damage to the common is shared by the entire group. If all herders make this individually rational economic decision, the common will be depleted to the detriment of all.

Solutions

- Privatization
- Usage fees
- Regulation
- Social norms

Example: Eliminating use of chlorofluorocarbons.

The Computing-Research Common

Observation: We conduct our professional life in an "professional pasture" – the computing-research common. We all contribute to and consume from this common:

- Journals
- Conferences
- Funding
- Reference letters
- ...

My Thesis: The computing-research common suffers from the tragedy of the commons.

Talk Outline

- Introduction
- Examples
- Discussion

Example: Journal of Logic Programming

In November 1999, the entire 50-person editorial board of the *Journal of Logic Programming* (Elsevier), led by Maurice Bruynooghe, resigned to protest Elsevier's exorbitant subscription rates and formed a new journal, *Theory and Practice of Logic Programming* (Cambridge U. Pr.).

Elsevier was left with the shell of the *Journal of Logic Programming*. It brought in new editors and renamed it the *Journal of Logic and Algebraic Programming* at an institutional rate of \$701/year, a slight increase over the price of the previous incarnation.

Question: Why did the new editors agrees to support Elsevier?

Answer: It was in their self-interest!

Open Access

Definition: Open-access publishing is "the publication of material in such a way that it is available to all readers without financial or other barriers."

October 2003 Berlin Declaration on Open Access to Knowledge in the Sciences and Humanities: unfettered access to knowledge!

- *Motto*: "Information wants to be free!"

Fact: As EiC of CACM, I am often asked: "Why don't you adopt the open-access model?"

- Very often the question is raised by non-ACM members – self-interest!

"Free" Is Not a Sound Business Model!

Facts:

- The annual cost of publishing CACM is $5M.
- Annual publishing costs for ACM total $12M.

Possible Sources of Funding:

- Author fee
- Member fee
- Advertising

Question: What is best for the Computing-Research Common?

- Will researchers get access to libraries' subscription budget?

Bigger Question: Who Decides?

Main Computing-Research Publishers:

- AAAI - non-profit association
- ACM - non-profit association
- IEEE Computer Society - non-profit association
- Usenix - non-profit association
- Elsevier - for-profit corporation
- Springer - for-profit corporation

They decide!

Who Decides?

Contrast:

- Non-profit associations:
  - Goal: promote computing
  - Means: Democratic associations – *We decide!*
- For-profit corporations:
  - Goal: Maximize profits
  - Means: Maximize revenue and minimize expenses – *They decide!*

Bottom Line: The associations are us.

- Churchill: "Democracy is the worst system of government, except for all the rest!"

The Tragedy

Fundamental Paradox: Why should for-profit corporations receive products and labor essentially for free and then charge us exorbitant rates?

- Ask your library how much it pays for ACM journals and DL versus analogous rates for Elsevier and Springer publications?

Question: Why do we continue to support for-profit journals?

Answer: Self-interest!

- Authors want to publish.
- Associate editors receive prestige.
- Editor-in-Chiefs receive monetary compensation.

Journals vs. Conferences

A Cliché: If *everyone but you* is driving on the wrong side of the road, then *you* are driving on the wrong side of the road!

Fact: We are the only technical discipline that considers conference publication as the primary means of publishing research results.

Why?:

- Once upon a time: was meant to offer fast dissemination to *complement* slow journal publication.
- Computing Research Association, 1999: Best Practices Memo *unintentionally* legitimized conference publication.

Journals vs Conferences– Merits

Journals

- *Pros*: fully fleshed articles, careful reviewing
- *Cons*: sloooow!

Conferences

- *Pros*: speed
- *Cons*: short papers, superficial reviews, balkanizatin of field

Why Practice Persists?: Self-interst!

- Easier to write a conference-length article than a full-length article.
- One accumulates "brownie points" by publishing in main conferences.
- Motivating conference deadlines

Journals vs Conferences - Reactions

- Lance Fortnow: "It is time for computer science to grow up and publish in a way that represents the major discipline it has become.
- Jeannette Wing: "How can we break the cycle of deadline-driven research?"
- Felipo Menczer: "I propose the abolition of conference proceedings alltogether."

Jano van Hemert: "For CS to grow up, CS journals must grow up first."

- "slow turnaround time, with most taking at least a year to make a publish reject decision and some taking much longer before publishing."

Why CS Journals So Slow?

Reasons for Slowness

- *Time to publish*: pipeline, page limit
  - *Solutions*: Utilize digital publishing!
- *Time to decide*: slow editors and referees
  - *Solution*: We are the problem!

Bottom Line: Making journals work better is up to us!

Hypercriticality

Feedback on CACM mostly positive, but:

"Although I have looked at every issue and at least glanced at every article, I have not yet found one good one."

and

"The level is unbelievably poor. It reads sometimes like a PR article for big companies. Donation to the ACM seems to be the main reviewing criterion. I would call the policy of ACM scientific prostitution, and I don't want to pay for a prostitute."

Are We Really Nasty?

- Ed Lazowska: "circling the wagons and shooting inwards"
- John L. King: "fratricide"
- Jeff Naughton: "bad reviewing is sucking the air out of our community"

Fact: Proposals submitted to the Computer and Information Science and Engineering Directorate of the U.S. National Science Foundation are rated, on the average, close to 0.4 lower (on a 1-to-5 scale) than the average proposal.

Why Are We So Nasty?

Two Theories:

- Computing systems are notoriously brittle. Mistyping one variable name can lead to a catastrophic failure. In our eternal hunt for flaws, we often focus on the negative and lose perspective of the positive.

- We typically publish in conferences where acceptance rates are 1/3, 1/4, or even lower. Reviewers read papers with "reject" as the default mode. They pounce on every weakness, finding justification for a decision that, in some sense, has already been made.

Another explanation: self-interest!

- It is rational to be hypercritical – classical Prisoner's Dilemma

The Golden Rule of Reviewing

Hillel The Elder: "What is hateful to you, do not do to your fellow!"

- *Silver Rule* of moral philosophy

Golden Rule: "Do unto others as you would have them do to you!"

Golden Rule of Reviewing: "write a review as if you are writing it to yourself!"

- Demand high quality
- But be fair, constructive, and respectful!

What Happened to Conference Reviewing?

Question: Did conference reviewing use to be better in the good old days?

Answer: Not really. It used to be *different*!

- 1975-1985: no reviews, only decisions!
  - 100 papers to read, no subreferees!
- 1985-1995: the emergence of reviews
  - "*selective conferences*" $\mapsto$ "*refereed conferences*"
- 1995-2000: from F2F to Web meetings

The Consequences of Web Meetings

A True Story:

A conference PC member receives a paper for review. He distributes the paper to his research group to "solicit their opinions of the paper." The group then embarks on improving the results of the paper under review. They submit their paper to another conference, three months before the first paper is presented in a conference; their paper is accepted. When eventually confronted (the four-months gap between the appearances of the two papers triggered questions), the PC member responded with "Was that wrong? Should I have not done that?"

Technology has Social Consequence

During the early 1970, running water was installed in the houses of Ibieca, a small village in northeast Spain. With pipes running directly to their homes, Ibiecans no longer had to fetch water from the village fountain. Families gradually purchased washing machines, and women stopped to scrub laundry by hand at the village washbasin.

Arduous tasks were rendered technologically superfluously, but village social life unexpectedly changed. The public fountain and washbasin, once scenes of vigorous social interaction, became nearly deserted. Men began losing their sense of familiarity with the children and donkeys that once helped them haul water. Women stopped congregating at the washbasin to intermix scrubbing with politically empowering gossip about men and village life. In hindsight, the installation of running water helped break down the Ibiecan's strong bonds–with one another, and with the land–that had knit them together as a community.

Face-to-Face vs. Web Meetings

F2F Meeting: high cost to PC member

- Cost of travel
- High reading load
- Social pressure to demonstrate competence

Consequence: At most one PC per year!

Web Meeting: low cost to PC member

- No travel cost
- Reduced reading load
- Reduced social pressure to demonstrate competence
- Ease of review delegation

Consequence: Low-cost brownie points $\mapsto$ several PCs per year!

Technology Has Social Consequence

Consequence I: Low-quality reviews!

- Review delegation (to "subreferees") evolved from a means to obtain additional expertise to a means to "train students" and reduce load – The PC member as "mini-editor"

Consequence II: Absence of socialization

- *Socialization*= process of inheriting norms and customs
- *Example*: Confidentiality as a lost norm

Praise Inflation

Maturing Field: specialization!

- *Symptom*: Little ability to understand colleagues' work even in a small department

Consequence: total reliance on external evaluation letters

- We write evaluation letters for our students and close colleagues.

Growing Phenomenon: praise inflation

- *Example*: "He is my best student in the past five years!"

Outcome: "Red Queen Evolution"

Overselling and Overpromises

Reality: We compete for scarce resources!

- Paper slots
- Jobs
- Funding

Natural Tendency: oversell and overpromise
Compare:

- "We plan to study an alternative approach to processor microarchitecture."
- "We will develop a revolutionary approach to processor microarchitecture."

Fact: Most research is incremental, it is hard to identify revolutions

What Is The Solution to The Tragedy?

- Debate the issues
- Turn agreements into norms

More Debate, Please!

Vigorous debate is a form of inquiry, a way to expose all sides of the issue.

- Benjamin Franklin: "When Truth and Error have fair play, the former is always an overmatch for the latter. When Men differ in Opinion, both Sides ought equally to have the Advantage of being heard by the Publick."

Point/Counterpoint Feature of CACM:

- E-voting
- Net neutrality
- MapReduce vs Relational Databases

But: people recoil from debates!

Example

CACM, May 1979: "Social Processes and Proofs of Theorems and Programs" by De Millo, Lipton and Perlis

- Formal verification of programs is difficult to justify and manage.

30 years later:

- Formal verification is an industrial reality
- Numerous awards for research in formal verificaton

V., CACM, January 2010: The editors of Communications in 1979 erred in not accompanying a tendentious article with a counterpoint article.

Ensuing Debate - March 2010

De Millo & Lipton: "Though we support the sentiment as well, we question Vardi's judgment in using his editorial position to mount an attack on a 30-year-old article whose authors were neither forewarned nor given the opportunity to respond. It is an extraordinary event when the Editor-in-Chief of a professional journal uses his position to declare ex cathedra that a published article is "misguided," its arguments "off the mark," and prior editors "did err in publishing [the] article... without publishing a counterpoint article..."

V.: "The article in question is more than 30 years old. History, it is said, 'judges and rejudges.' I hardly view my offering of some comments, even if critical, on such a historically important article as 'mounting an attack.' Personally, if someone saw the need to disagree with an article of mine 30 years after its publication, I'd feel complimented. Most articles are long forgotten after 30 years."

How Can We Form Norms?

Malcolm Gladwell: "Why the revolution will not be tweeted"

At four-thirty in the afternoon on Monday, February 1, 1960, four college students sat down at the lunch counter at the Woolworth's in downtown Greensboro, North Carolina. They were freshman at North Carolina A.&T. black college a mile or so away.

"I'd like a cup of coffee, please," one of the four said to the waitress. "We don't serve Negroes here," she replied.

Within a month, thousands of students participated in these sit-in protests, and within a year, tens of thousands.

Gladwell: What makes people capable of this kind of activism?

Answer: Strong personal connections!

Why The Revolution Will Not Be Tweeted

Gladwell:

- Social media are tools for building networks.
- Networks are not controlled by a central authority.
- Ties are loose.
- Participants are not asked much.
- Networks have difficulty reaching consensus.
- They cannot make difficulty choices.

Observation: The Computing-Research Common is a network.

The Importance of Associations

Question: Why associate?

Answer: Associations enable us to go beyond networks

- Undertake coordinated efforts
- Chose philosophical directions
- Establish professional norms

A brief history of professional associations:

- *Then then*: Professional guilds - Roman Collegium (300 BC)
- *Then*: Royal Society (1660 AD)
- *Now*: Bar associations (the bar separate court from public)

Weak Associations in Computing Research

- Several "international" associations: AAAI, ACM, IEEE CS, IFIP, Usenix
- Several regional associations: BCS, CRA, EATCS, EASST, Informatics Europe...
- Heavy North American bias in ACM: 60% North America, 30% Europe, 10% other

Bottom Line: several weak associations = network!

First Step: Don't complain about ACM being American, join and make it more international!

- *Recent initiative*: ACM Europe

Summary

Main Points:

- The Computing-Research common suffers from the tragedy of the commons – we all behave rationally but the outcome is suboptimal to all of us.
- We cannot eliminate this completely, but we can ameliorate it by developing professional norms.
- We must debate the issues openly.
- Only a strong association can lift us from the local minimum in which we seem to be stuck.

Two Concrete Proposals

Please consider adopting these in your area!

- Return to F2F PC meetings.
  - Raise quality of PC decisions
  - Raise cost to PC members
  - Practiced by several leading conferences
- Demand conference submissions to contain all details in appendices
  - Make it harder to write conference papers
  - Make it easier to submit to journals