

Lecture #02

January 06, 2010

Today's Objectives

- Review Workshop/Conference/Journal procedures
- What makes a good venue?
- Outline of a good paper?

Today's Objectives

- Once upon a time, only journals were in libraries
 - Conference proceedings weren't and certainly not workshop proceedings
 - Even today, some workshop proceedings are hard to find
 - Though, now lots of stuff is online, either through a professional society (like IEEE, ACM, USENIX, or IFIP)
- Slowly, over time, more and more conferences had proceedings that were available in libraries
- Today, just about any proceedings is available online

Venue History

- Once upon a time, only journals were in libraries
 - Conference proceedings weren't and certainly not workshop proceedings
 - Even today, some workshop proceedings are hard to find
 - Though, now lots of stuff is online, either through a professional society (like IEEE, ACM, USENIX, or IFIP)
- Slowly, over time, more and more conferences had proceedings that were available in libraries
- Today, just about any proceedings is available online

Venue History (con't)

- ACM and IEEE lead the way in online access
 - Both have excellent digital libraries
 - Can access them for free from UCSB
 - <http://portal.acm.org/>
 - <http://ieeexplore.ieee.org/>
 - Can access them for free outside of UCSB by using UCSB Library proxy
 - <https://login.proxy.library.ucsb.edu:9443/login>
- The Library's search options are pretty good too
- The end result is that a couple of the major differences among workshops, conferences, and journals have disappeared.
 - It also means there's a lot more crap to sift through

Workshops

- Designed to focus on the hot topics
 - Topics are typically narrowly focused
 - Workshop may last for only a few years, though there are some that have lasted a long time (www.nossdav.org)
- The current popular model is to have workshops attached to conferences
 - Reduces expenses for conference and allow more people to attend with less travel
 - Though there are still some workshops that are independent
- Typically single track and lasts 1-3 days
 - That means anywhere from 12-25 papers
 - Workshops attached to conferences are typically single day

Workshop Program

- Most programs are posted right before the workshop
 - Take a look at some of the workshops you've identified to get a sense of how they are organized
- Most have a keynote (or opening invited talk)
- Some have a panel
- Otherwise, the majority of the time is spent listening to talks about papers
 - Anywhere from 18-30 minutes per paper (plus questions)
 - Using this model as a guide for this class
 - Some workshops try to increase interactivity by having all authors come up at the end of a session and do a mini-panel

Workshop Organization

- A *Steering Committee* that includes about 6 people
 - Maintains continuity from year-to-year
 - Chairs from a given year might serve on the SC
- An *Organizing Committee* with two parts:
 - One that handles logistics (General Chair)
 - One that handles selection of papers (TPC Chair)
- For workshops, sometimes the two are merged and there is a single chair that does both parts
 - Current trend is to have two co-chairs
 - Divide up responsibilities and provide checks-and-balance
- People are selected for positions based on reputation
 - Almost always a researcher: academia or industry lab
 - are considered prestigious but are a lot of work
 - Selections are done by SC

Technical Program Committee

- The group that decides which submissions will be presented at the workshop
- TPCs vary widely in how they operate
 - Lots of ways of getting from set of submissions to set of selected papers
- Basic Steps
 - Paper submission
 - Paper assignment to TPC members
 - Review phase
 - Decision phase
 - Camera ready paper

Paper Submission

- These days, via a web site
 - EDAS is pretty widely used (<http://edas.info/>)
- Papers can be single blind or double blind
 - Single blind: author names are on the paper (reviewers are not identified)
 - Double blind: papers are submitted anonymously
 - An experiment was tried with IEEE Global Internet where the reviewer names were identified
- Papers must follow a specific length and format
 - Provides level playing field
 - Good insight into quality of venue: more pages means more contribution and generally better quality

Phases

- Paper assignment
 - Either the TPC Chairs assign papers
 - requires TPC Chairs to know TPC member expertise (and more work)
 - TPC Members can bid for papers based on interests
 - Sometimes leads to conflicts, but still pretty rare
- Paper review
 - Some period of time for TPC Members to review papers
 - Almost all workshops allow delegation of reviews, but many prefer the TPC Member to do the reviews him/her self
- SO: keep in mind that paper could be reviewed by expert in your area (maybe even someone whose work you are citing), a complete non-expert, or even a student

Phases (cont)

- Paper selection
 - Sometimes there is a “discussion phase”
 - All reviewers get to see the other reviews and comment on them
 - Sometimes one TPC Member is assigned to write a summary of the discussion and a consensus recommendation for the paper
 - There is almost always a “TPC meeting”
 - Can happen on a teleconference, but those are very hard to run
 - Trend is to have TPC meetings at other meetings/conferences
 - Typically last a day depending on number of papers to review
 - Papers are divided into “clear accept”, “discuss”, and “clear reject”
 - Each paper is discussed for between 5-10 minutes
 - Once accepted list of papers are selected, TPC Chairs organize them into sessions
 - Typically, the better the venue, the more thorough the review process will be

Phases (cont)

- Camera ready preparation
 - Authors have approx a month to prepare final version
 - Sometimes, if there are concerns about the paper, it is “shepherded”
 - A TPC member will be assigned to review paper and make sure concerns are addressed
 - Otherwise, authors are free to do as little or as much as they want to change the paper
 - There are few checks, so an author COULD make wholesale changes, but so far, that’s pretty uncommon

Conferences

- Today, conferences are pretty similar to workshops
- In addition to paper presentations, keynote, and panels, might also have:
 - Tutorials (half- or full-day), poster session, work in progress session, workshops, more panels, and more invited talks
 - Can be either single track or multiple track
 - Sometimes, tracks can have their own TPC Chair (ICDCS and ICCCN are two examples that come to mind)
- With more submissions, the TPC process becomes more complicated
 - More TPC co-chairs, maybe TPC vice chairs, TPC area chairs are becoming popular

Journals

- Basic difference is that there is no event to show up and present
 - Result is that journal articles MAY be less visible
 - Were once considered more prestigious, and the top journals still are, but for CS, the major venue is conferences
 - For tenure and promotion, journals are still important
- Submissions are either to special issue (fixed submission date) or open call (submit any time)
- Two options:
 - Assigned to Associate Editor who finds three reviews and recommends decision to Editor
 - Assigned to three Associate Editors who do the reviews and the Editor makes the decision
 - Both are about equally common

Journals (con't)

- Instead of accept or reject, journals may have a “resubmit after revision”
 - Authors make changes, describe in addendum what changes were made (response to reviewers), and then the paper goes through another round of reviews
 - This step is less common for special issues (since the schedule can't accommodate multiple review cycles)
- Acceptance rates for journals are generally higher than conferences
 - Obviously varies, but my sense is most are about 25%
 - Higher degree of self selection

Finding Places to Look

Why Look?

- For ideas...
- But more important if you are trying to do something new
- Need a systematic way of checking the literature
- The challenge is that the amount of work being done these days is H-U-G-E

Two Ways to Search

- Use a search engine
 - Google scholar, Citeseer, UCSB Library
 - In general, if you don't know what you are looking for, you won't find it
 - Probably okay if you are just looking for ideas, but certainly not if you are looking to do something new
- Know where to look
 - Know where relevant research is being published
 - Know what the good venues are
 - Know who is doing related work
 - A hard skill to learn, but necessary to do original research
 - Comes with practice and experience
 - HW1 was a first step

Good v. Bad Venues

- Generally only have to look in good venues
 - The idea might show up in a bad venue, but unlikely, and if it does, probably won't be detailed enough
- So, how to measure a good venue?
- Generally start with ACM and IEEE sponsored workshops, conferences, and journals
 - Depending on the area, might also see USENIX—especially for operating systems (OSDI) and networking (NSDI)
 - There are a few unsponsored events that are good, but be careful

Good v. Bad Venues (con't)

- Acceptance rate is one of the better metrics
 - <http://www.cs.ucsb.edu/~almeroth/conf/stats/>
 - Approx 10%: best in the field
 - 10%-25%: very good
 - 25%-35%: probably okay
 - >35%: iffy
 - If it isn't on this page, it should raise a yellow flag
- The other really good metric is the reputations of the TPC chairs and members
 - Okay, hard to judge without lots of experience
 - A gross approximation is the quality of the organization the TPC chairs and members are affiliated with
 - But even this is hard to judge since overall reputation doesn't mean they are good in networking and some specific sub-discipline

Good v. Bad Venues (con't)

- There are other metrics but they are less reliable
- Page limits of submissions
 - Another subtle clue: the more stringent the message to adhere to the formatting requirements, the more competitive: highly competitive venues try hard to keep playing field level
- Number of tracks and submission
 - Generally less visible the more tracks (which was the whole point of submitting to a conference in the first place)
 - IEEE Infocom is the major exception (but “number of tracks” is only a secondary consideration—Infocom’s acceptance rate is quite low, so everything works out)

Evaluating a Paper

High-Level Objectives

- Primary objectives
 - Novel area of research
 - Should be a new problem or a problem with minimal work in the field
 - Novel solution to a known problem
 - Fine if other solutions exist, but proposed solution has to be innovative, solve the problem in a novel way, more efficient, SOMETHING
 - If there are other solutions, a comparative evaluation needs to be done
 - Pretty common to apply a known technique to slightly different problem
 - Novel evaluation methodology
 - It is one thing to evaluate something using a different simulator or measured data set, but it is more useful to get a new result that provides more insight
- For top conferences you need the first, the second, and if not the third, a very, very thorough evaluation
 - Difference conferences have different “styles”: most noticeable is the kind of evaluation

Secondary Objectives

- Well written
 - Hard to judge how important this is—it can be used as either a strength or weakness
 - Reviewers sometimes argue that the primary contributions are so strong, it outweighs this factor; also, writing can be fixed
 - Also used as a reason to reject a paper that otherwise is just borderline
 - More common, leads to misunderstanding of the paper
- Proper citation of related work
 - Worst case is when reviewer is an author of something he/she considers relevant
 - Also used to reject papers when there aren't other good reasons

Evaluating a Paper

- Without some experience, may be hard to judge contribution
 - And how much better a proposed scheme is than what has been proposed before
 - In these situations, generally rely on the author to explain it to you
 - Mention this in the review (most reviews have a “Private Comments to TPC Chairs” that the authors won’t see)
- You may read a paper and are lost
 - That’s generally the authors fault
 - After all, you do have a BS degree!
 - There is quite a bit of discussion about how understandable papers should be and whether a paper should be readable by people outside the sub-sub-sub-sub-discipline

Writing a Paper

- As you read papers, consider what you look for in a paper and what works well for you are the kinds of techniques you should use when you write your own papers
- Another good judge is to try explaining an idea to someone (the reason we do practice presentations)
 - Whenever they ask questions, note the question
 - In a paper, there isn't an opportunity for questions, so once you lose the reader, that's it, you're done
 - You only can use another person once: once they hear the idea the first time, the second time, they'll have more background for the explanation than a person reading a paper for the first time will have

Paper Structure

- Translating the goals of a paper into parts
- MAJOR CAVEAT: this is one way to do it; it is not the only way—it may not be the best way
 - The explanation I'm about to give is just one way of translating the major components into an outline

High Level Outline

- Abstract
 - Write last... it is a complete description of everything else
- Introduction
 - Weird in CS: isn't really an introduction—is more of an extended abstract
- Background/Related Work
- Architecture/Idea/Protocol/Measured-System
- Evaluation/Results
- Conclusions/Future Work

Introduction

- Overview paragraph (define area)
 - "P2P systems allow the exchange of files..."
- Identify problem, show significance
 - One problem with P2P systems is that people often cheat..."
- Mention related work and why it fails
 - "Reputation systems have been proposed to help identifying cheaters, however..."
- Introduce new idea (and take ownership)
 - "We propose a new ACTIVE reputation creation system"
- Overview evaluation and describe what it shows
 - "We demonstrate ACTIVE's better performance through a series of sims"
- Identify contributions and impacts
 - "With better reputation systems, P2P systems will become more usable ..."
- Standard organizational review
 - "The remainder of this paper is organized as follows."

Related Work

- Typically start with re-motivating, usually with more detail
 - Enough detail to understand the problem
 - Enough detail to understand what related work does and doesn't do
- Organize the related work into coherent groups
 - a good way of representing (related work and) alternatives is with a spectrum (Ex: QoS spectrum)
- Summarize group, give references, and identify deficiencies
 - identifying deficiencies is CRITICAL (otherwise, what is the need for your work?)
- At the end summarize what is missing
 - This becomes the transition to your proposed idea
- Sometimes Related Work can also include more motivation and background—especially necessary if the area is quite new

Solution

- Pretty straightforward... just present your solution
 - Hopefully there are some alternatives
 - slight variations on the theme you are proposing
 - will make comparisons more interesting
 - If there aren't alternative you are proposing, and instead, you are comparing your system to an alternate solution (and showing yours is better), you can either present that solution as part of related work or you can present it here as a known alternative

Evaluation

- Critical that you explain to the reader what is going to happen in the evaluation section
- Describe goals/hypothesis
 - "In our evaluation, we hope to show..."
- Describe methodology
 - "Our evaluation consists of a series of..."
 - Methodology is a good place to find fault in a paper, so be careful
- Describe assumptions
 - List assumption and justify what it is reasonable to make
 - Best reason is that assumption simplifies complexity but doesn't change results
 - "Too hard" is typically not a good justification
 - Another good place to find fault!
- List parameters, factors, nominal values, metrics...
 - Parameters: things that affect performance but you don't change
 - Factors: things that you vary
 - Nominal value: value of factor when you aren't changing it
 - Metrics: basis for comparison

Present Results

- First, given a high-level "road map" for evaluation
- Start with simple result, and explain clearly
 - Explain axes
 - Identify points of interest
 - Identify implications of results
 - If server satisfies 100 requests, is 100 requests a lot?
 - Is it better than other systems?
 - Identify open question to be answered by next graph
- Tell a "sequential story"
 - Create a "story" that results will follow
 - Helps to identify what graphs are needed (axes, legends, etc.)
- Be prepared to change architecture as you learn
- Be prepared to do 100s of simulations

Paper Themes

- Paper themes should run throughout the paper:
 - Summarized in the abstract
 - Introduced in the introduction
 - Are the focus of the solution
 - Are the goal of the evaluation
 - Are the conclusion of the paper
 - Ex: “We introduce a protocol that has better performance”
 - Ex: “We develop a coding technique that is robust to different kinds of loss and results in higher user satisfaction”
- That’s why you write the abstract and polish the introduction last
 - You aren’t sure until you build the evaluation what you can really show
 - You aren’t sure what you can really show until you’ve got a graph in hand that has that result

Conclusion

- Most papers just do a summary
 - Look for it in the papers you review
- A good conclusion also includes
 - A summary (definitely necessary)
 - Ties up the themes/goals of the paper
 - Provides implications of the research
 - Includes some future work ideas

Final Thoughts

- Pay attention to section transitions
 - Absolutely critical to provide a road map for the reader
 - Need to explain why what you are saying is important
 - Need to explain why you are saying what you are saying
- Novice writers often overlook just how much support is needed
- As always, think about explaining this in a talk and what you'd have to explain to your audience