Advanced Topics in Computer Networking

WS2010/11

Introduction

Xiaoming Fu

Yang Chen

Tianyin Xu

Overview

- Basic concepts: Gain experience in **Reading**, **Writing** and **Presenting** research ideas in English
- Learn **knowledge** about new and emerging research topics and technologies in networking
- Organize **discussions**, exchange **ideas**, **experiences** between students and researchers

Research: Different Aspects

Problem selection

Research Methodology

Presentation

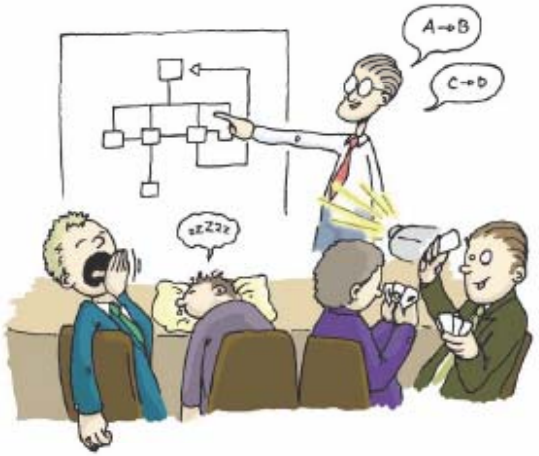

What we plan to do?

- It's **research-oriented**. Therefore, the main focus will be course project and paper reading.
- It's **Seminar**, so we will.....Discuss, Discuss and Discuss with prepared reading materials
- Every week (November – January), we discuss **one** paper
- Final "formal" discussions in January (formal presentations, Q&A)

How to evaluate?

- Part I: Session participation (40%)
  - Including questioning, attendance, paper presentation, paper reviews
  - Opportunity to practice everyone's presentation & reading skills!
- Part II: formal final presentation (30%)
- Part III: 12-15 pages essay (30%)

Main Topics

- P2P Networking
  - Internet Geometry
  - Distributed Hash Table
- Online Social Networking
  - Facebook, Twitter, LinkedIn…

Detailed Lecture Plans (1)

- 05.11.2010 [Session 0]
  - Course Introduction (how to read, write....)
  - Invited talk by Prof. Tilman Wolf
- 12.11.2010 [Session 1]: Peer-to-Peer Networking
  - Paper resding
- 19.11.2010 [Session 2]: Peer-to-Peer Networking
  - Paper reading
- 26.11.2010 [Session 3]: Peer-to-Peer Networking
  - Paper reading
- 03.12.2010 [Session 4]: Peer-to-Peer Networking
  - Paper reading

Lecture plans (2)

- 10.12.2010 [Session 5]: No class
- 17.12.2010 [Session 6]: Peer-to-Peer Networking
  - Paper reading
  - Decisions on assigning papers for the final presentation!
  - Introduction of online social networking
- 07.01.2011 [Session 7]: Online Social Networking
  - Paper reading
- 14.01.2011 [Session 8]: Online Social Networking
  - Paper reading
- 21.01.2011 [Session 9]: Online Social Networking
  - Paper reading
- 28.01.2011 [Final presentation]

Seminar Preparation and Process

- Before the seminar (Wednesday 23:59):
  - Read the selected paper
  - Prepare the review of the paper and send it to (E-mail: yang.chen@cs.uni-goettingen.de):
    - Summary of the paper
    - pros and cons of the paper (your conclusion)
    - With the title '[ATCN10] XXX'
- During the seminar:
  - one is chosen for giving a presentation of the paper.
  - And the list of pros and cons is discussed by all the participant
    - Debate between two groups

How to read

- Preparation: pencil, paper photocopy of article
- Decide what to read: read title, abstract
- Rough round: what're the major points?
  - Introduction, conclusion and first paragraph of the main work
  - React to the what you read: write down/put note
- Read in depth:
  - build a structure of the problem/system
  - challenge/verify the credibility of the idea: do the assumption, method, statistics & conclusion sound?
  - construct your own example
- Summarize what you read

Slides of the paper

- Would be really helpful for understanding of the paper …
  - Enter the name of the paper author (and/or Project name, Univ.) in Google
  - Access the webpage
  - You may find the slides (and/or Project Homepage)there!
  - Sometimes not available ☹

Paper review and discussion

- Major contributions of the paper
- Differences from related work
- Sparking ideas and techniques proposed in the paper
- Weaknesses
  - technical errors, unrealistic assumptions, unanswered questions, limited measurements
- Ways to improve the paper technically
  - To overcome above weaknesses
- Interesting topics discussed in its future work
- Ideas for extending the work

Presentation

Give clear presentation of your work

How to present

- Read the article beforehand, sketch your impressions about the author's idea
- Decide which is the best idea in the paper
  - Write it down and justify with a line sentence
- Figure out how to get your audience as quickly as possible to the point where they can understand this idea
- Elaborate the idea in details
  - Background, proof, benefits/difference over others
- Summary: thinking/justification of the paper in your own words

Essay

- Title
- Abstract: a broad overview of the report, end with a short statement of the major results of your investigation
- Introduction (or Motivation of the work): expands the abstract, get specific about your investigation
- Related work: unlikely invented sth completely new
- Methodology/Technical approach: what's important idea of the work? Bring your reader to the method
- Advantages and improvements of the work
- Results: testing of "implemented" method
- Future thoughts along the direction (has the problem fully resolved?)

How to cite references

- If a new term appears in your context for the first time, try to explain with your own words, and better cite a reference which can explicitly explain it
- If an approach is new but presented by others, cite the corresponding paper
- If a figure or data is copied from other resource, please cite them

=> improper/missing references are depreciated and regarded as unethical

Networking Research: Resources (1)

- Academic Journals: well-established results
  - **IEEE/ACM ToN**: top journal in the field
  - IEEE JSAC: special issues
  - Elsevier ComNet
  - ACM/Kluwer WINET, IEEE ToW/TMC, Kluwer MONET, Wiley WCMC: wireless/mobile domain
  - ACM CCR, MC2R: fast reviewed paper
- Magazines: tutorials/surveys
  - IEEE ComMag, IEEE Network, IEEE Internet Computing

Networking Research: Resources (2)

- Good Conferences: fast, peer-reviewed papers
  - **USENIX SOSP, NSDI, OSDI**
  - ***ACM SIGCOMM***
  - ***ACM MOBICOM, ACM MOBIHOC, ACM MOBISYS, ACM CoNEXT***
  - **IEEE INFOCOM,**
  - **IEEE ICNP**, IEEE ICDCS
  - IEEE/IFIP IWQoS
  - ACM SIGMETRICS, IMC, PERFORMANCE
- Q: How to find good ones from tons of resources?
  - IEEE, ACM sponsored ones normally **may** be better than others
  - Papers from top labs/univs: **MIT, Berkeley, Stanford, CMU, Princeton, Cambridge,** Caltech, UIUC, GIT, Cornel, UCLA, Washington, UPenn, UMass, Columbia; ETHZ, Karlsruhe, TUM; Bell Labs, AT&T Labs, IBM, MSR, Intel research, etc.

Networking Research: Resources (cont.)

- Hmm… it's not easy to remember so many names before starting this seminar
- Is there some way easier? Yes!
  - Many researchers have provided something useful for everyone
  - Also, we will introduce one conference per week in this seminar

http://grid.hust.edu.cn/call/

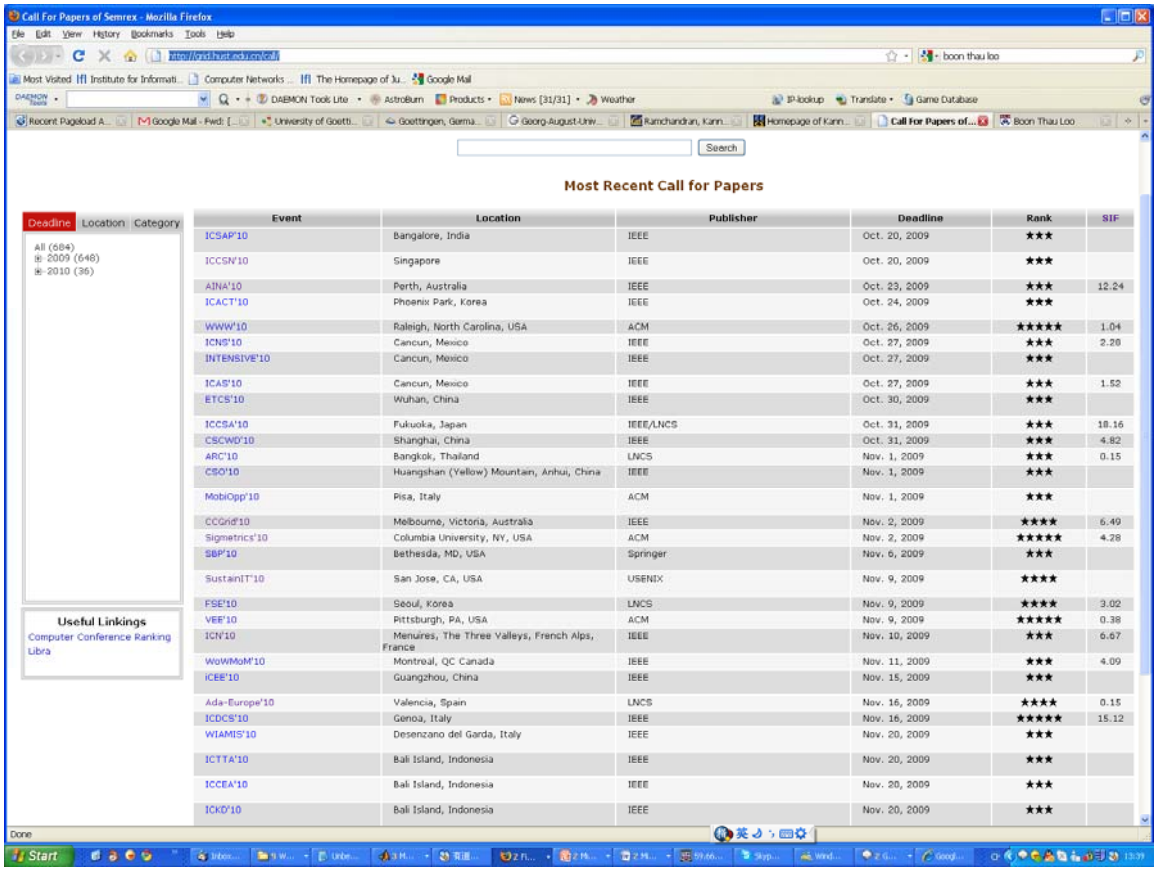

Stars

http://citeseerx.ist.psu.edu/stats/venues?y=2007

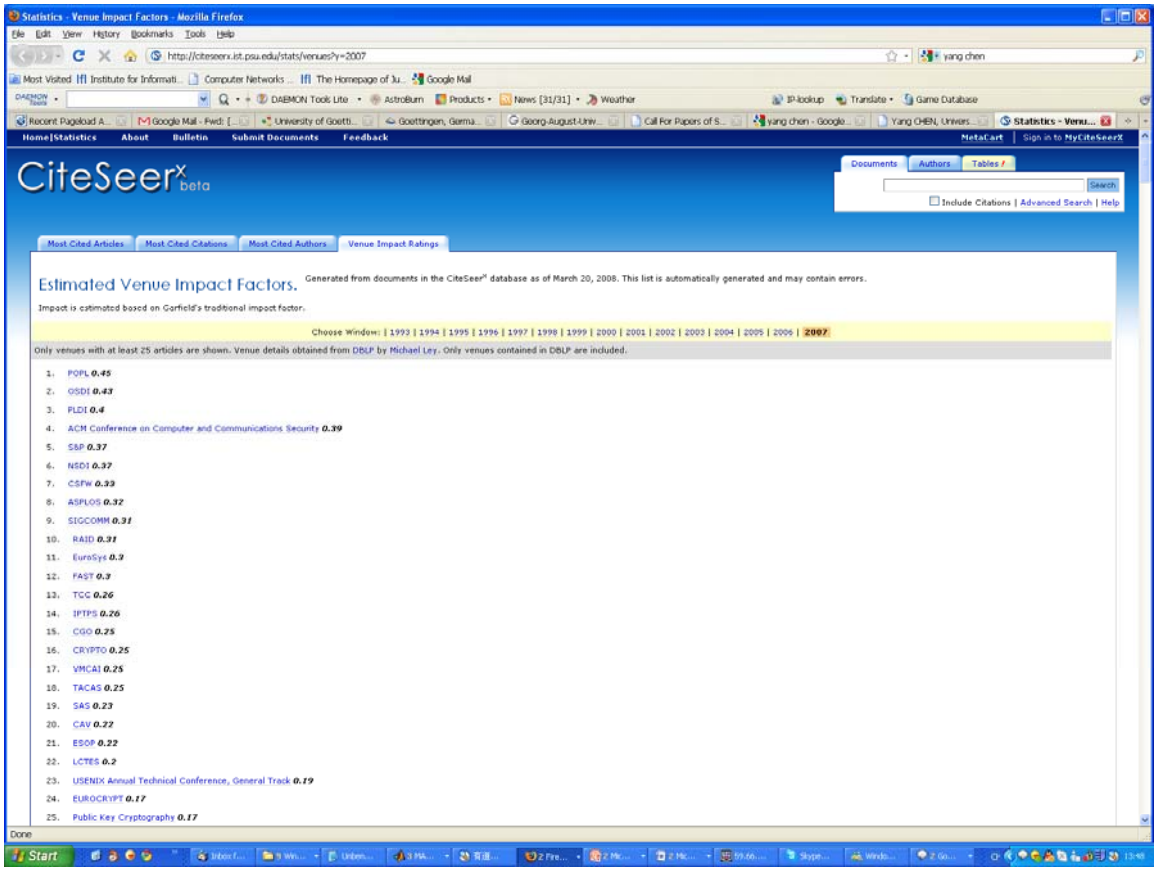

Impact Factor

http://www.cs.ucsb.edu/~almeroth/conf/stats/

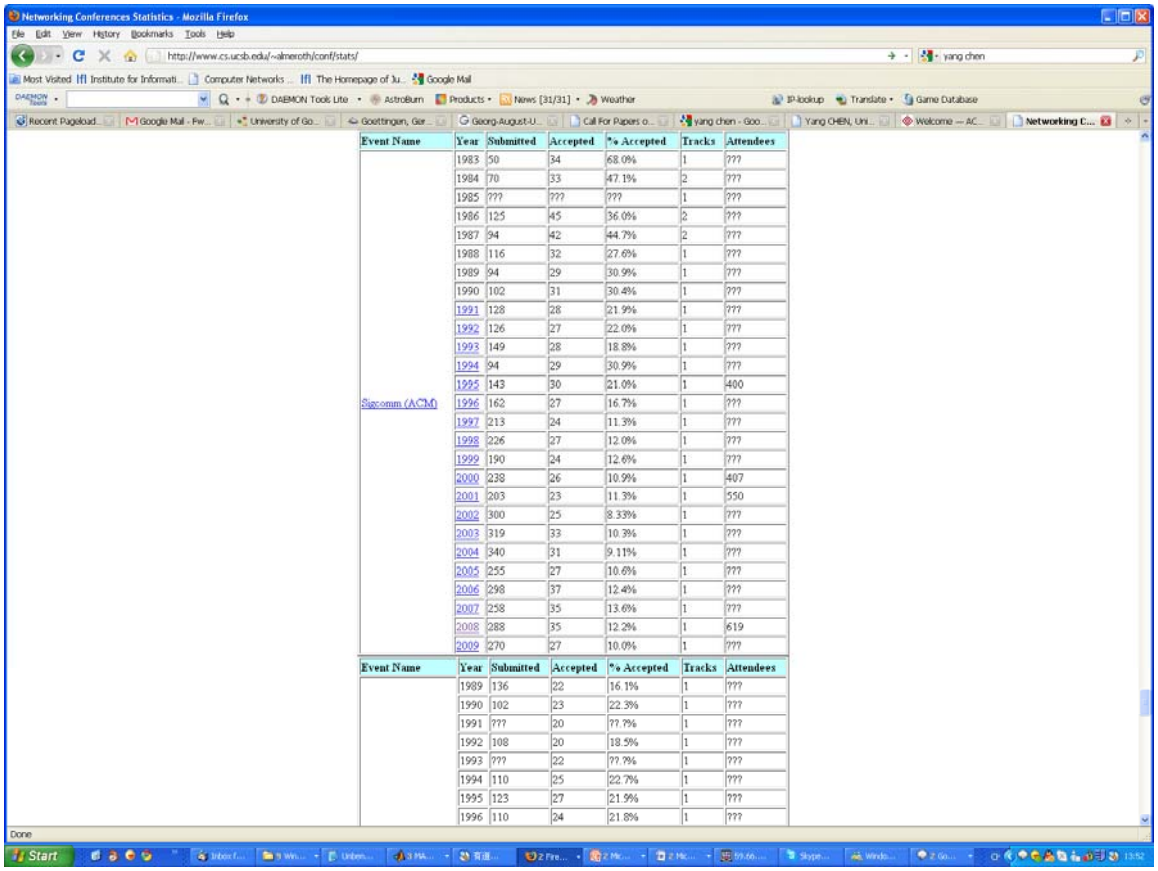

Event Name	Year	Submitted	Accepted	% Accepted	Tracks	Attendees
Sigcomm (ACM)	1983	50	34	68.0%	1	???
	1984	70	33	47.1%	2	???
	1985	???	???	???	1	???
	1986	125	45	36.0%	2	???
	1987	94	42	44.7%	2	???
	1988	116	32	27.6%	1	???
	1989	94	29	30.9%	1	???
	1990	102	31	30.4%	1	???
	1991	128	28	21.9%	1	???
	1992	126	27	22.0%	1	???
	1993	149	28	18.8%	1	???
	1994	94	29	30.9%	1	???
	1995	143	30	21.0%	1	400
	1996	162	27	16.7%	1	???
	1997	213	24	11.3%	1	???
	1998	226	27	12.0%	1	???
	1999	190	24	12.6%	1	???
	2000	238	26	10.9%	1	407
	2001	203	23	11.3%	1	550
	2002	300	25	8.33%	1	???
	2003	319	33	10.3%	1	???
	2004	340	31	9.11%	1	???
	2005	255	27	10.6%	1	???
	2006	298	37	12.4%	1	???
	2007	258	35	13.6%	1	???
	2008	288	35	12.2%	1	619
	2009	270	27	10.0%	1	???

Event Name	Year	Submitted	Accepted	% Accepted	Tracks	Attendees
	1989	136	22	16.1%	1	???
	1990	102	23	22.3%	1	???
	1991	???	20	??.?%	1	???
	1992	108	20	18.5%	1	???
	1993	???	22	??.?%	1	???
	1994	110	25	22.7%	1	???
	1995	123	27	21.9%	1	???
	1996	110	24	21.8%	1	???

Acceptance Ratio

Other Course Materials

Main literatures:

- Selected recent papers (mostly from afore-mentioned resources, will be uploaded in studIP as required).

Background reading – so many textbooks:

- Kurose & Ross: Computer Networking – A Top-Down Approach Featuring the Internet, 2nd ed.
- Schiller: Mobile Communications, 2nd ed.
- Peterson & Davie: Computer Networks - A System Approach. 2nd ed.
- Hofmann & Beaumont: Content Networking: Architecture, Protocols and Practices.

Introduction to Internet Geometry

World Map (Location in the Earth)

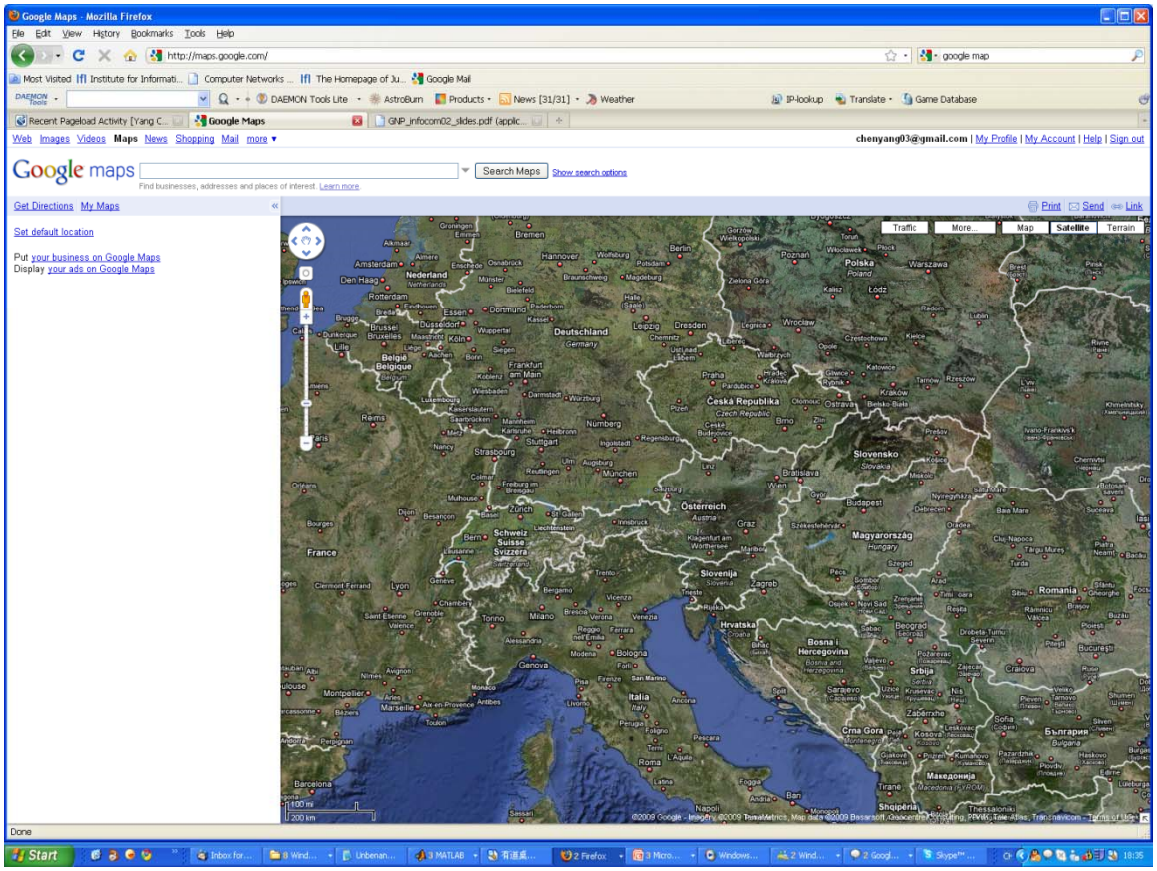

How to locate a place in the world?

- Longitude
- Latitude

- With these two numbers (coordinate), the location of a certain place can be identified

- Can we assign the coordinate to Internet hosts?

Internet Coordinates

- Mapping Internet nodes into a Geometric space (e.g Euclidean space) while preserving the distances (delays)

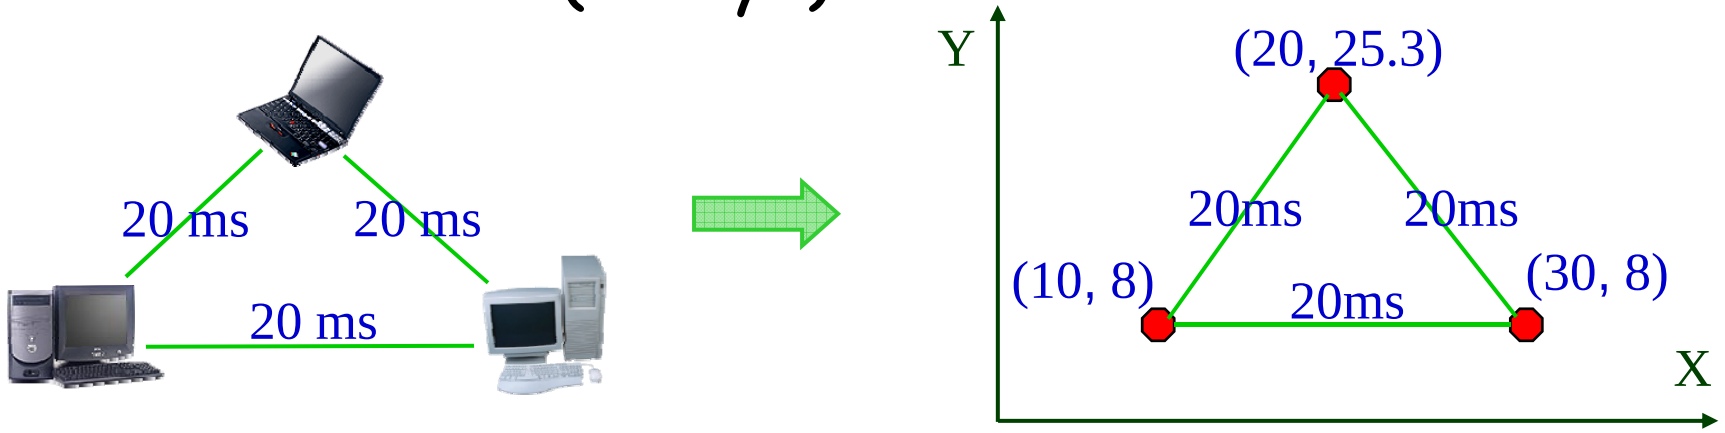

Internet Map

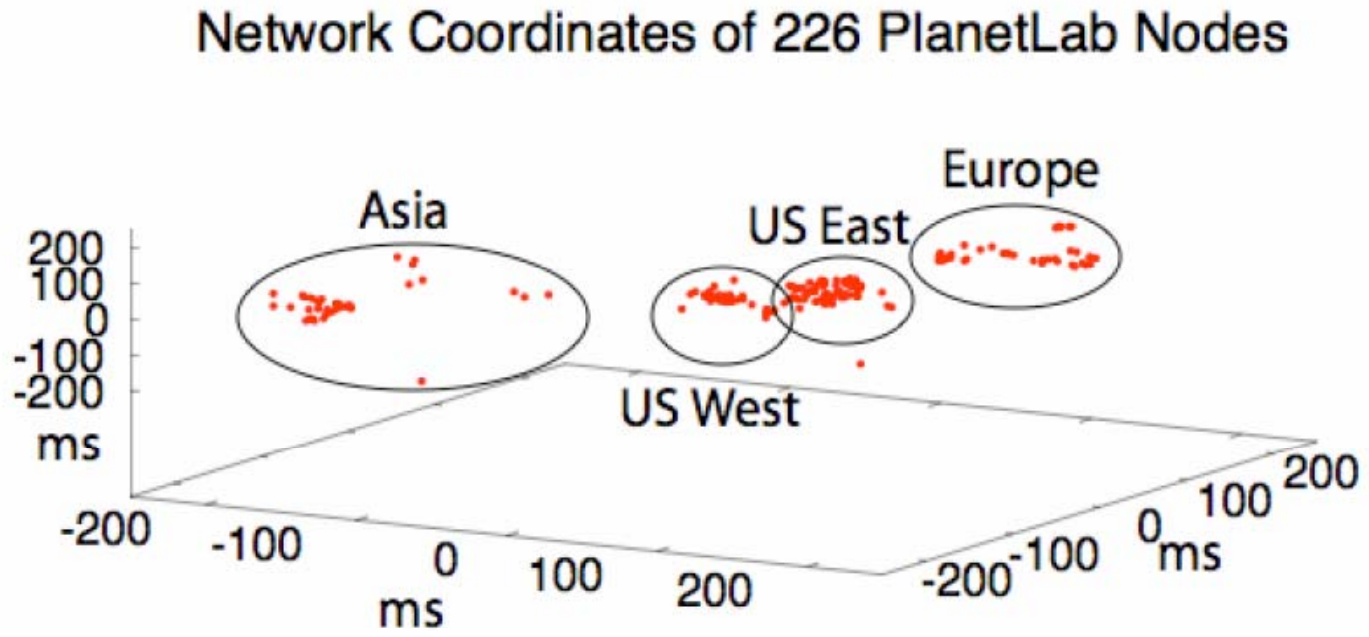

- Points represent locations of PL nodes in 3-d relative coordinate space

Internet Map

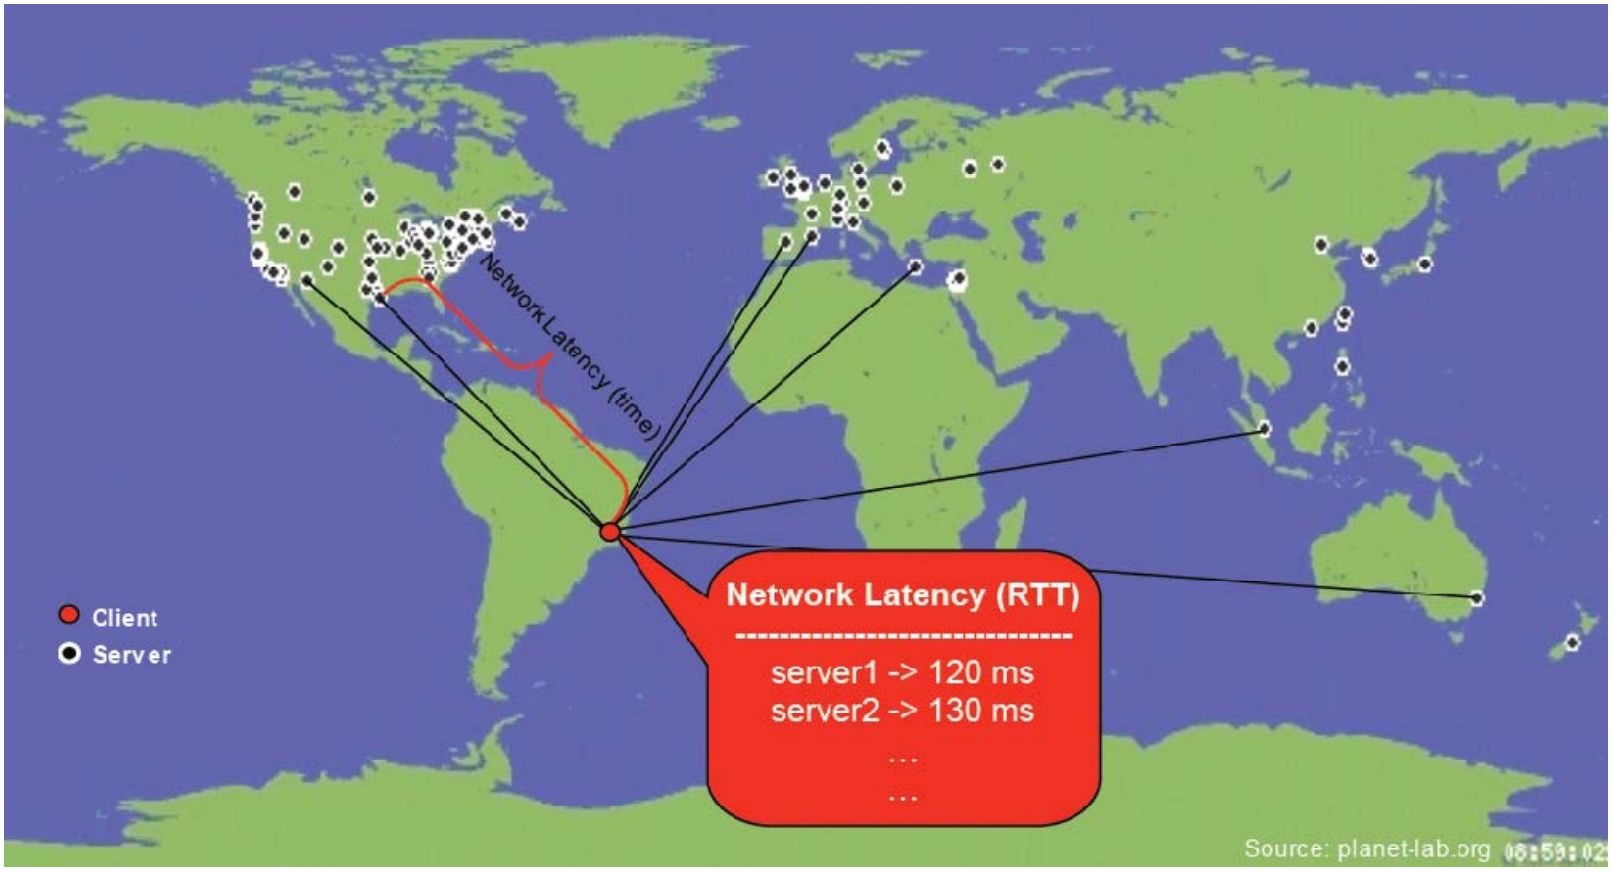

Distance Prediction

Can this Problem be Solved with a Peer-to-Peer Architecture?

- Flip the problem: Can end hosts maintain "coordinates" that describe their network locations?
- End hosts exchange coordinates to compute distance
  - High performance, high scalability

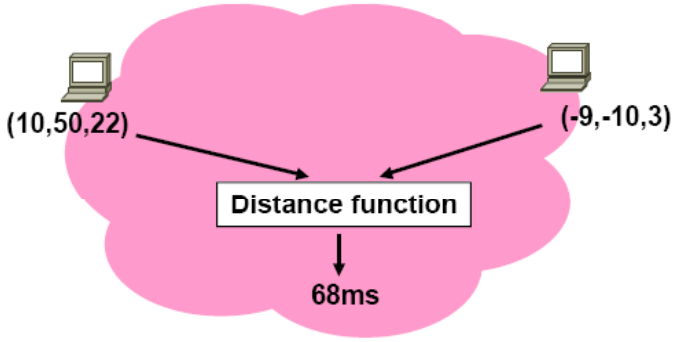

Application 1

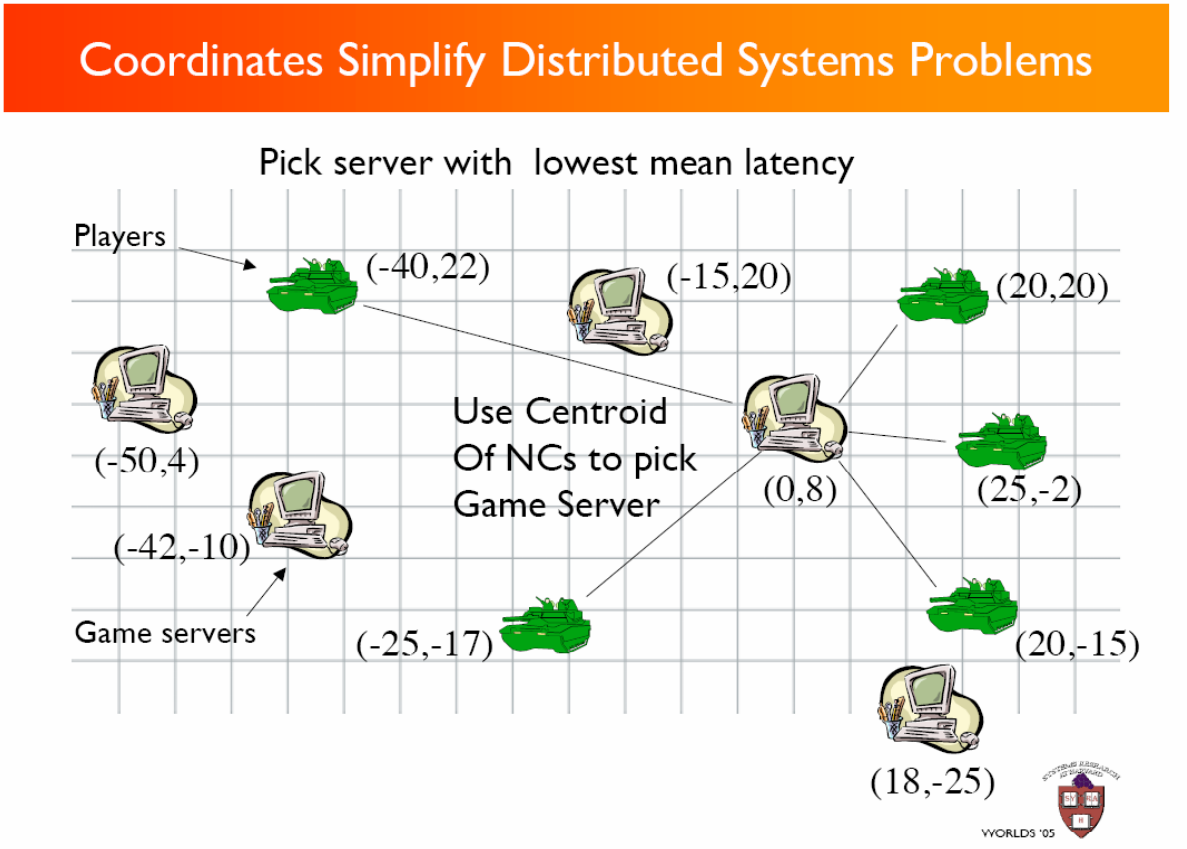

Application 2

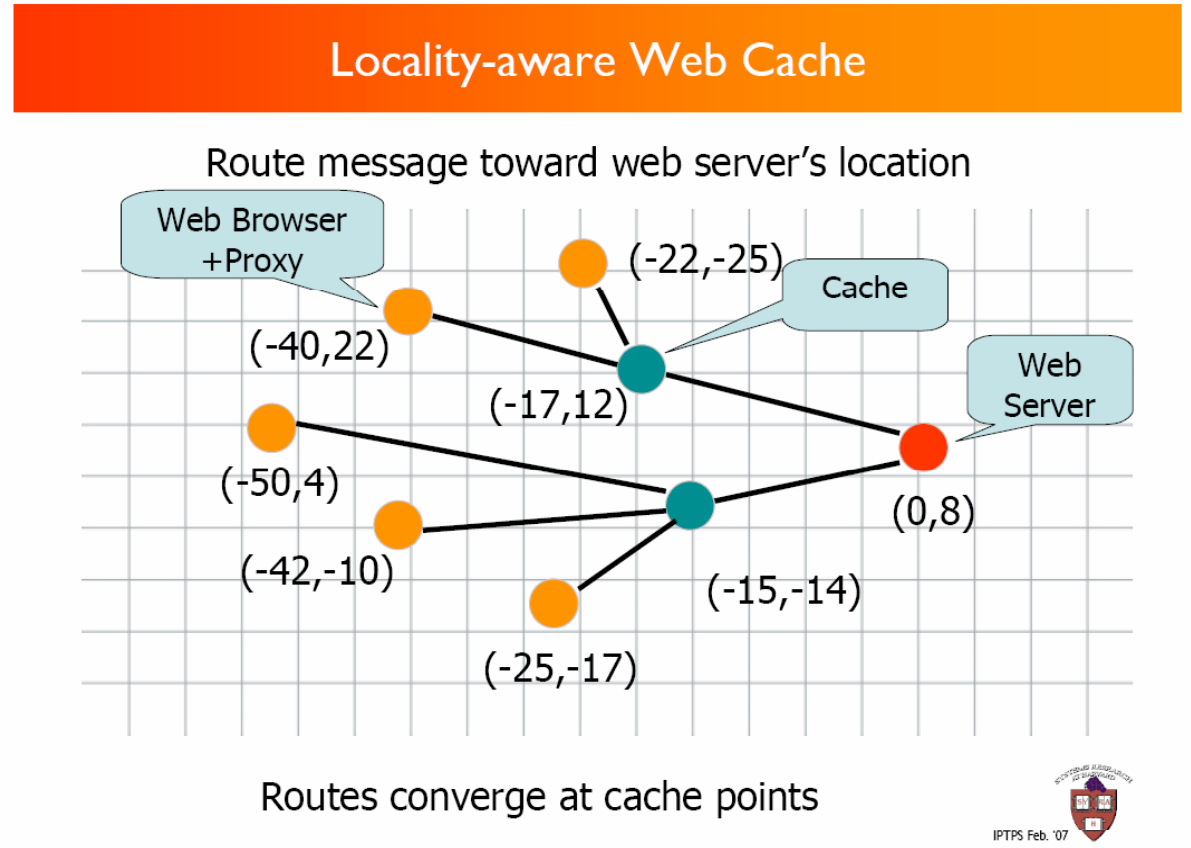

Paper Reading (This week)

- T. S. E. Ng and H. Zhang. Predicting Internet Network Distance with Coordinates-based Approaches. In Proc. of IEEE INFOCOM, 2002.

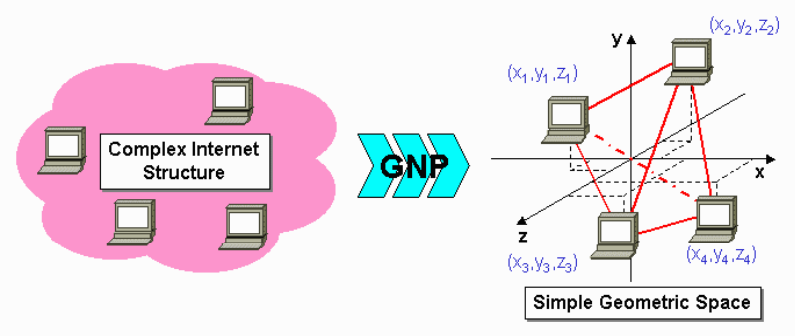

Project Homepage

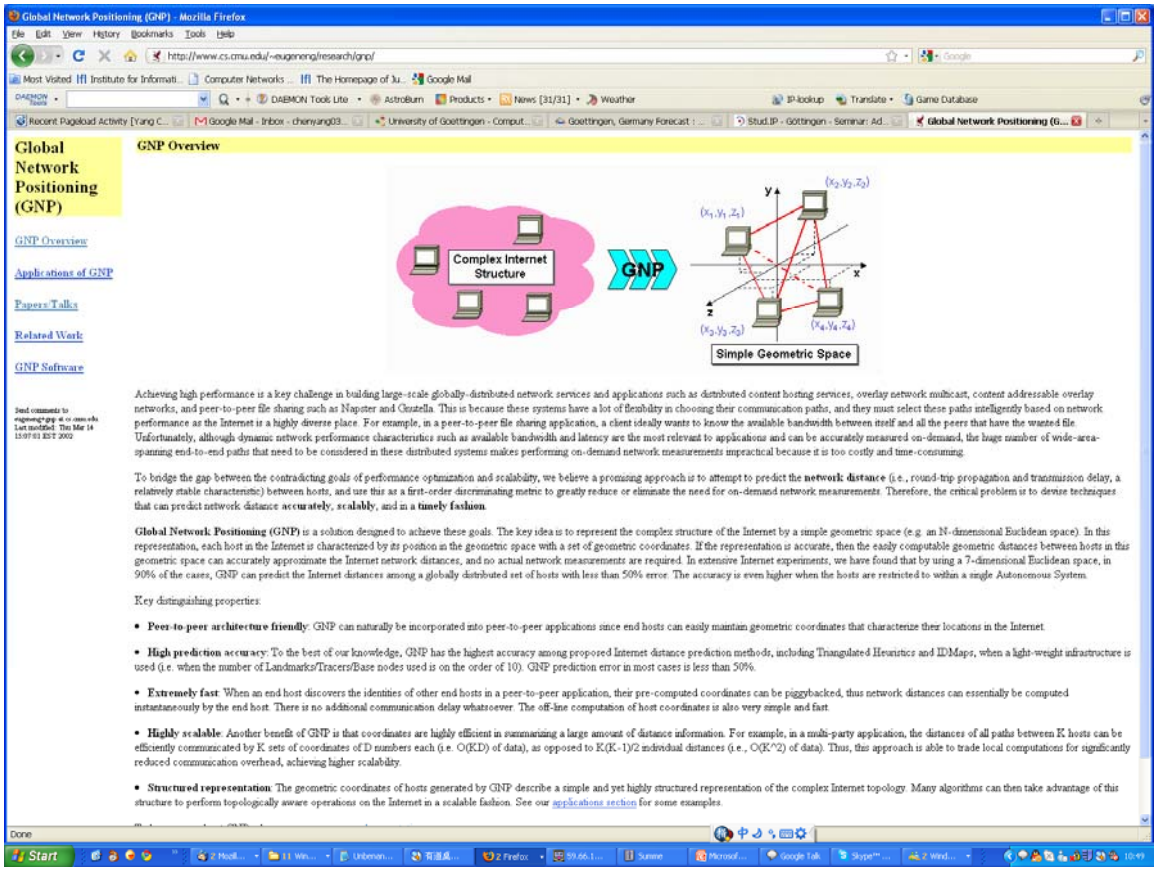

Any questions…

- Feel free to drop me e-mail

yang.chen@cs.uni-goettingen.de

Backup Slides

Choice of dissemination means

Forums

- Journals:
  - Well established NEW Results
- International conferences
  - New Results, Broad Audience
- Topical conferences
  - New Results, Selected Audience
- Workshops
  - Work in progress, New Ideas

Format

- Paper
  - Short (5-8),long(9-30pg)
- Abstract
  - 1pg, extended(2-3pg)
- Oral presentation
  - Short(10-