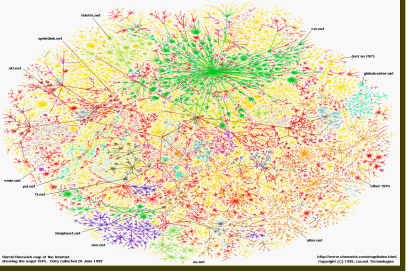

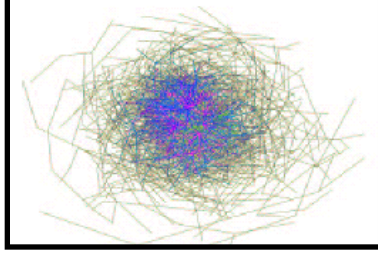

CS525

Advanced Distributed Systems

Spring 2016

Indranil Gupta (Indy)

Wrap-Up

January 19 – May 3, 2016

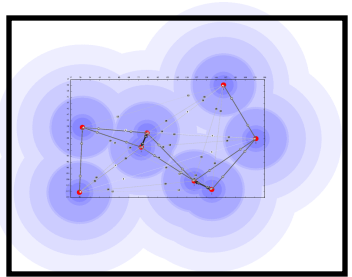

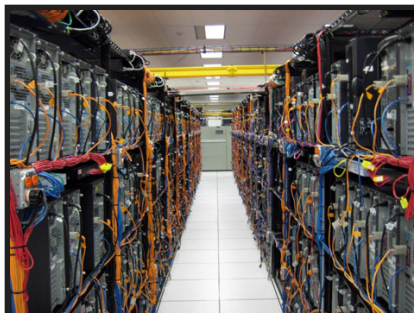

Agenda

- Wrap-Up of Discussion started at Course Beginning
- Articles

Can you name some examples of Operating Systems?

Can you name some examples of Operating Systems?

…

Linux WinXP Vista Unix FreeBSD Mac OSX

2K Aegis Scout Hydra Mach SPIN

OS/2 Express Flux Hope Spring

AntaresOS EOS LOS SQOS LittleOS TINOS

PalmOS WinCE TinyOS

…

What is an Operating System?

What is an Operating System?

- User interface to hardware (device driver)
- Provides abstractions (processes, file system)
- Resource manager (scheduler)
- Means of communication (networking)
- …

Can you name some examples of Distributed Systems?

Distributed Systems Examples (From Beginning of Semester)

- Client-server (e.g., NFS)
- The Internet
- The Web
- A sensor network
- DNS
- BitTorrent (peer to peer overlay)
- Datacenters
- Hadoop

Do these look familiar?

- HBase, HDFS, Cassandra, Paxos, Mapreduce
- Grid, Gnutella, BitTorrent, Napster, Chord, Pastry, PAST, Ivy
- AWS (EC2, S3, EBS), Google Cloud
- Storm, Naiad, Heron
- Yesquel, Tapir, Tango
- TinyOS, TAG
- Epidemics, Bimodal multicast
- BlinkDB, Succinct, E-store, Centiman
- FlashGraph, Giraph
- IronFleet, P2
- These are some of the systems you've seen during CS525 this semester
- **77% of the papers** were new in SP16 compared to SP15!

What is a Distributed System?

The definition we started with

A distributed system is a collection of entities, each of which is autonomous, programmable, asynchronous and failure-prone, and which communicate through an unreliable communication medium.

- Our interest in distributed systems involves
  - algorithmics, design and implementation, maintenance, study
- Entity=a process on a device (PC, PDA, mote)
- Communication Medium=Wired or wireless network

A range of interesting problems for Distributed System designers

-
- P2P systems [Gnutella, Kazaa, BitTorrent]
- Cloud Infrastructures [AWS, Azure, GCE]
- Cloud Storage [Key-value stores, NoSQL, BigTable]
- Cloud Programming [MapReduce, Pig, Hive, Storm, Pregel]
- Coordination [Paxos]
- Routing [Sensor Networks, Internet]
-

A range of challenges

-
- Failures: no longer the exception, but rather a norm
- Scalability: 1000s of machines, Terabytes of data
- Asynchrony: clock skew and clock drift
- Security: of data, users, computations, etc.
-

Some of the Topics We've Covered

- Clouds and their predecessors (e.g., Grids and timesharing)
- Overlays and DHTs
- Sensor motes and TinyOS
- Basics – Lamport timestamps, Consensus, Snapshots, Failure detectors
- Epidemics
- Paxos
- Mapreduce
- Structure of Networks
- P2P Apps
- A Touch of Sensor Nets
- Key-value stores
- Dataflow Programming
- Adaptive Stream Processing
- Is CAP Dead?
- Storage Systems
- Cluster Scheduling

Past

Present

Some of the Topics We've Covered (2)

- Distributed Machine Learning
- Elasticity
- Reliability
- Latency = $$$
- Graph Processing
- Approximation
- Aggregation
- Verification and Models
- Debugging and Performance
- Lots of Industrial Systems

Future

- H. G. Wells, G. Hardin, Levin-Redell

CS 525 and Distributed Systems

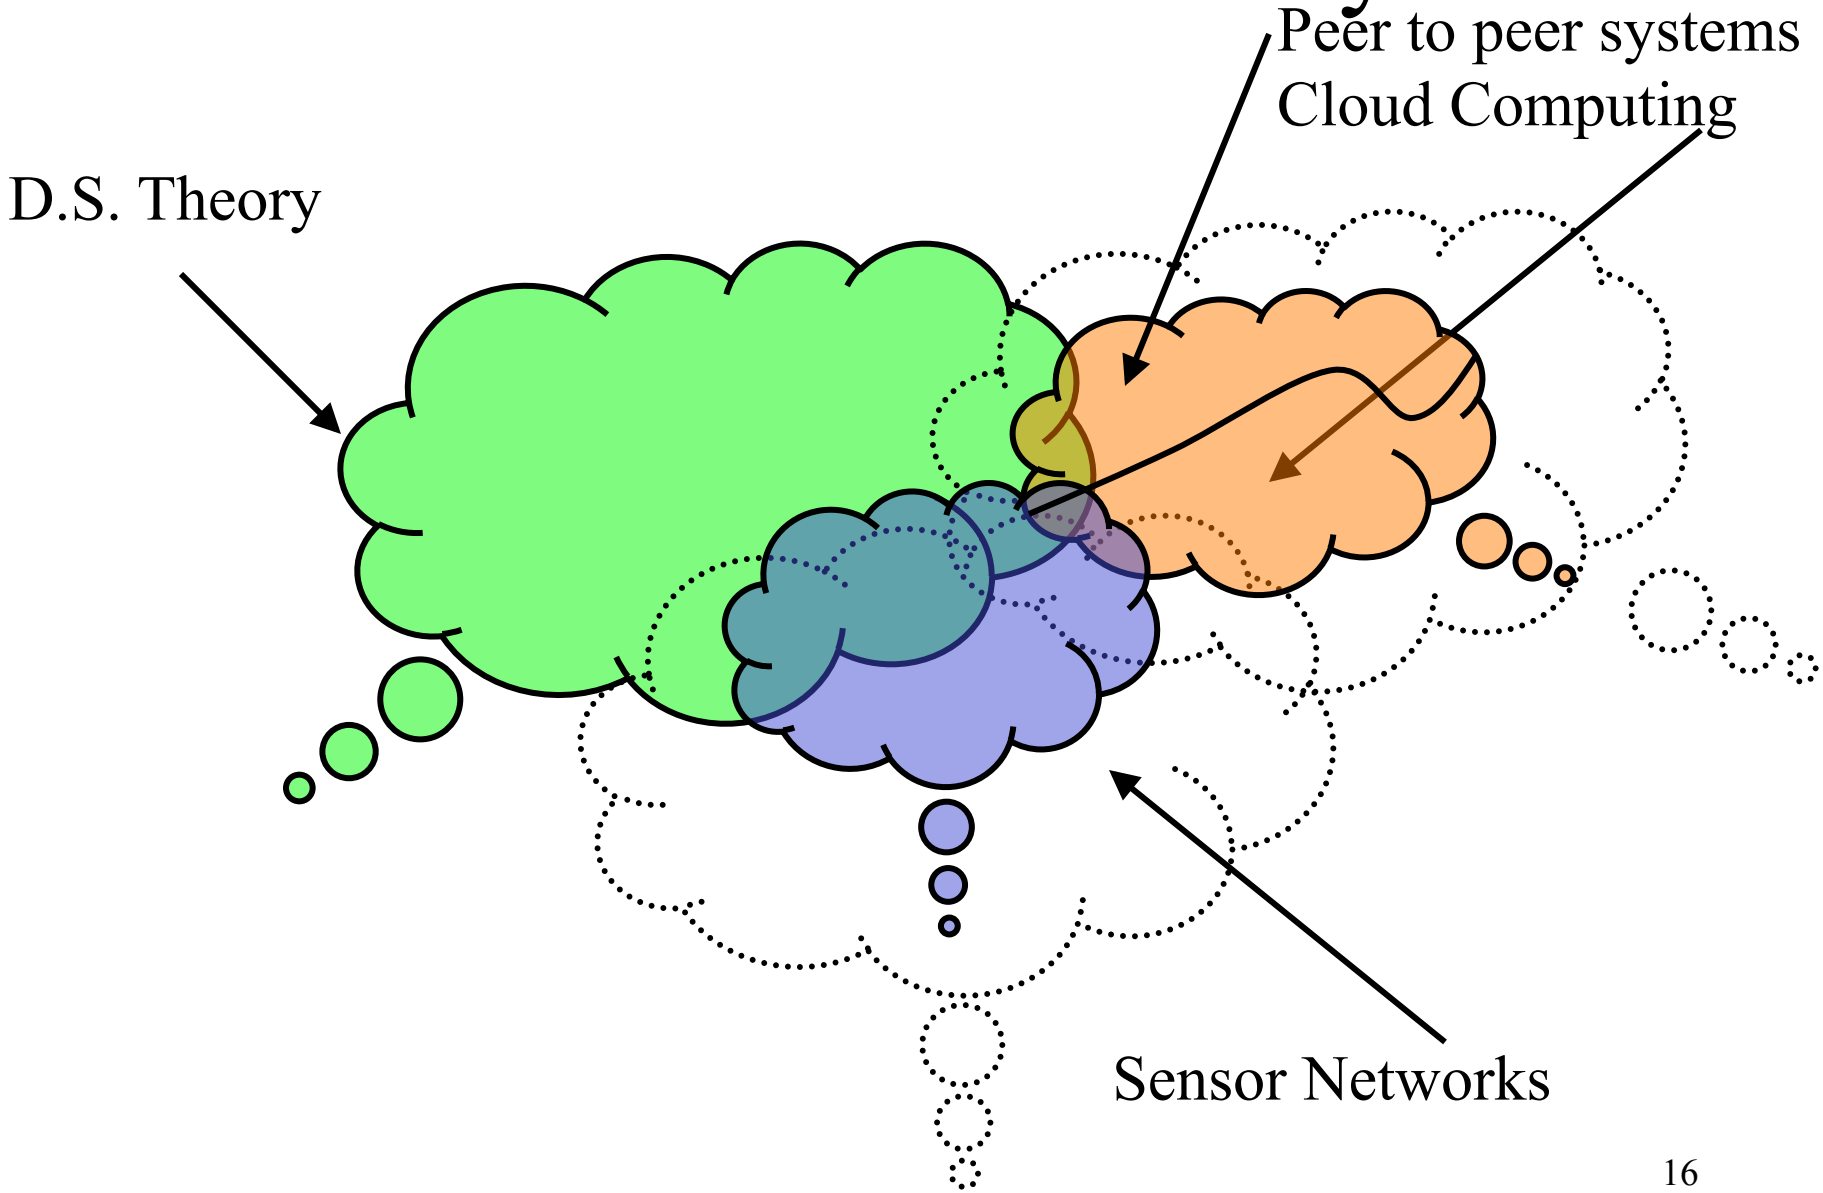

Interesting: Area Overlaps

Course Projects!

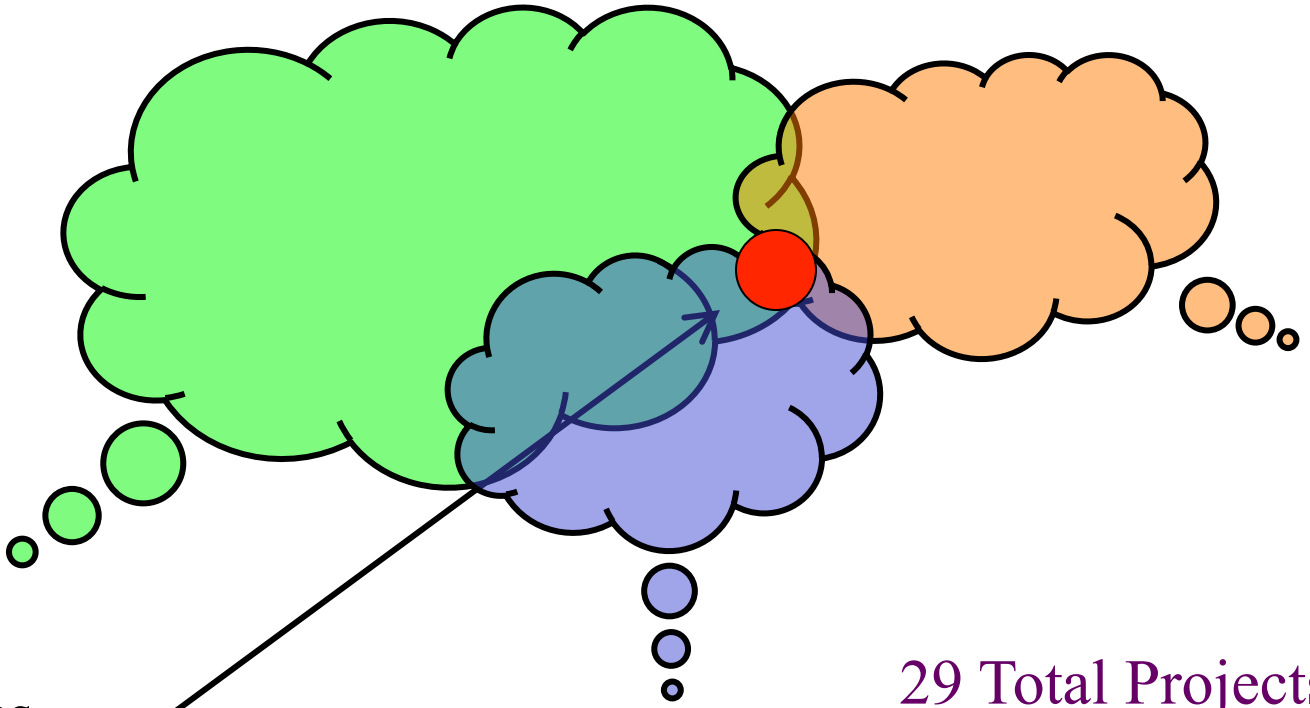

Epidemics
NNTP
Gossip-based ad-hoc routing

29 Total Projects
= 20 Research
+ 9 Entrepreneurial

Course Projects (Research)

1. Update Atomic Isolation: Highly Available Transactions without Lost Updates
2. Evaluating the consistency and latency of distributed datastores
3. High-throughput Pessimistic Locking for Cassandra
4. Monitoring Fine Grained Application Energy Consumption for Offloading in Mobile-Cloud Computing Frameworks
5. Web Service on Low-Power Embedded Micro Servers
6. Synchronization in a Heterogeneous Bluetooth Low Energy Sensor NW
7. Distributed Validation in Wireless Sensor Networks
8. Eagle: Accurate people counting in real time
9. BEEER: Distributed Record and Replay for Medical Devices in Hospital Emergency Operating Rooms
10. Automated and Transparent Checkpointing in Apache Spark
11. Meta-less PathSim–a PathSim based Similarity Search on MapReduce
12. Opportunistic Scheduling in Multiple-tenant Graph Processing Systems
13. Efficient Guaranteed Message Processing in Apache Storm
14. Visualization of Astrophysical Simulations using Distributed Stream Processing
15. Getafix: Workload-aware Data Management in Lookback Processing Systems

Course Projects (Research)

1. Update Atomic Isolation: Highly Available Transactions without Lost Updates
2. Evaluating the consistency and latency of distributed datastores
3. High-throughput Pessimistic Locking for Cassandra
4. Monitoring Fine Grained Application Energy Consumption for Offloading in Mobile-Cloud Computing Frameworks
5. Web Service on Low-Power Embedded Micro Servers
6. Synchronization in a Heterogeneous Bluetooth Low Energy Sensor NW
7. Distributed Validation in Wireless Sensor Networks
8. Eagle: Accurate people counting in real time
9. BEEER: Distributed Record and Replay for Medical Devices in Hospital Emergency Operating Rooms
10. Automated and Transparent Checkpointing in Apache Spark
11. Meta-less PathSim–a PathSim based Similarity Search on MapReduce
12. Opportunistic Scheduling in Multiple-tenant Graph Processing Systems
13. Efficient Guaranteed Message Processing in Apache Storm
14. Visualization of Astrophysical Simulations using Distributed Stream Processing
15. Getafix: Workload-aware Data Management in Lookback Processing Systems

Storage

Mobile and Sensors

Batch Processing

Stream Processing

Course Projects (Research + Entrepreneurial)

16. Profiling the Spot Market for Improved Efficiency
17. A Distributed Hypervisor with Automated Controller Deployment in Software-Defined Networks
18. Analysis of Large Scale Machine Learning Systems Based on Distributed Stochastic Gradient Descent
19. Analyzing the data-plane for anomaly and misconfiguration detection
20. OpenTimer 2.0: Distributed Timing Analysis at Scale

1. CollabOnCode
2. MyDrive – One stop distributed storage solution for all
3. Distributed Recommendation System Based on Graph Processing Engines
4. On-demand Games: Anyone, Anytime, Anywhere
5. FastNet: A P2P-Content Delivery Network
6. AcomoStreaming
7. GeoTweet: Real-time Geolocation Analyzer for Tweets
8. Synchronized Mobile P2P Video Sharing
9. Distributed System for 3D Video Conversion

Course Projects (Research + Entrepreneurial)

16. Profiling the Spot Market for Improved Efficiency
17. A Distributed Hypervisor with Automated Controller Deployment in Software-Defined Networks

Clouds&VMs

18. Analysis of Large Scale Machine Learning Systems Based on Distributed Stochastic Gradient Descent

Distbd. ML

19. Analyzing the data-plane for anomaly and misconfiguration detection
20. OpenTimer 2.0: Distributed Timing Analysis at Scale

Networking

1. CollabOnCode
2. MyDrive – One stop distributed storage solution for all

Collaboration

3. Distributed Recommendation System Based on Graph Processing Engines

Recommendations

4. On-demand Games: Anyone, Anytime, Anywhere
5. FastNet: A P2P-Content Delivery Network
6. AcomoStreaming
7. GeoTweet: Real-time Geolocation Analyzer for Tweets
8. Synchronized Mobile P2P Video Sharing
9. Distributed System for 3D Video Conversion

Streaming

CS 525 Ongoing Projects

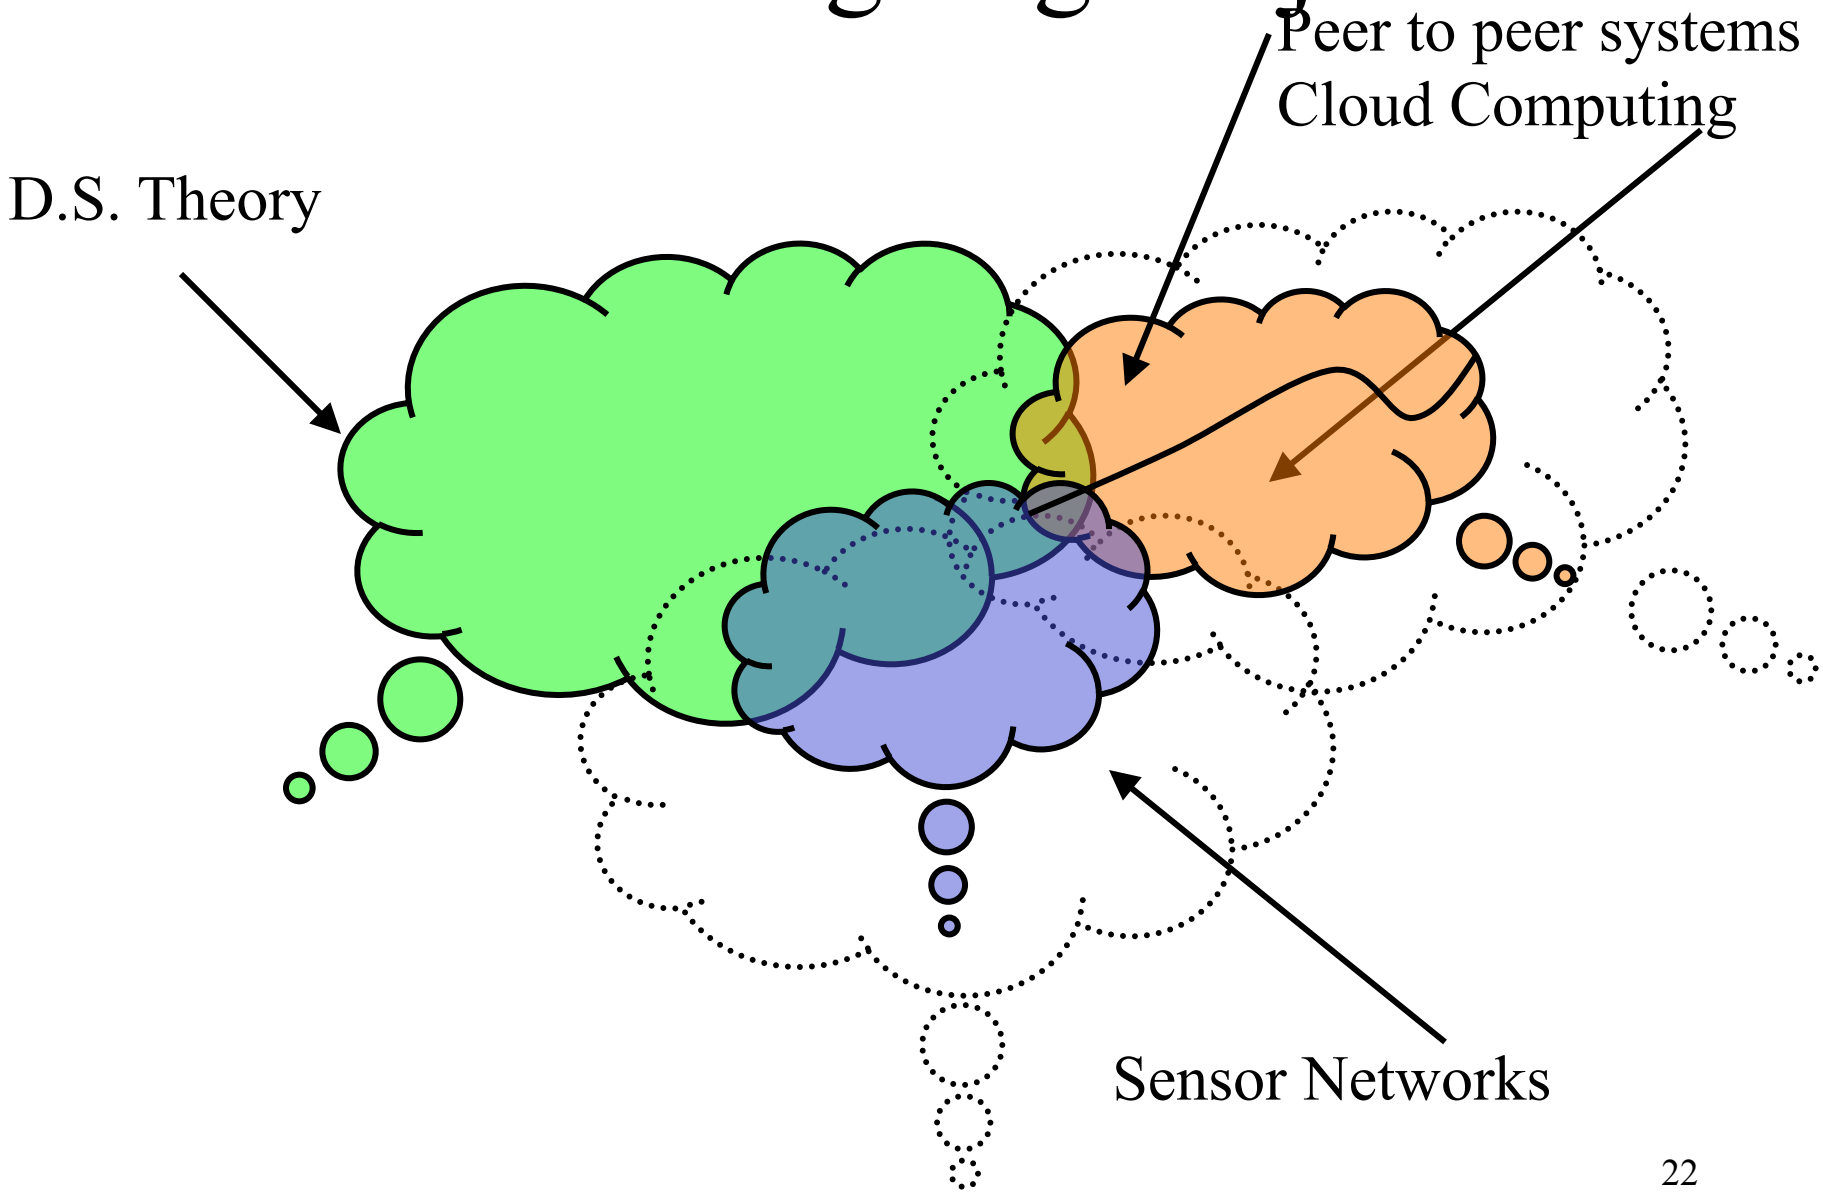

CS 525 Ongoing Projects

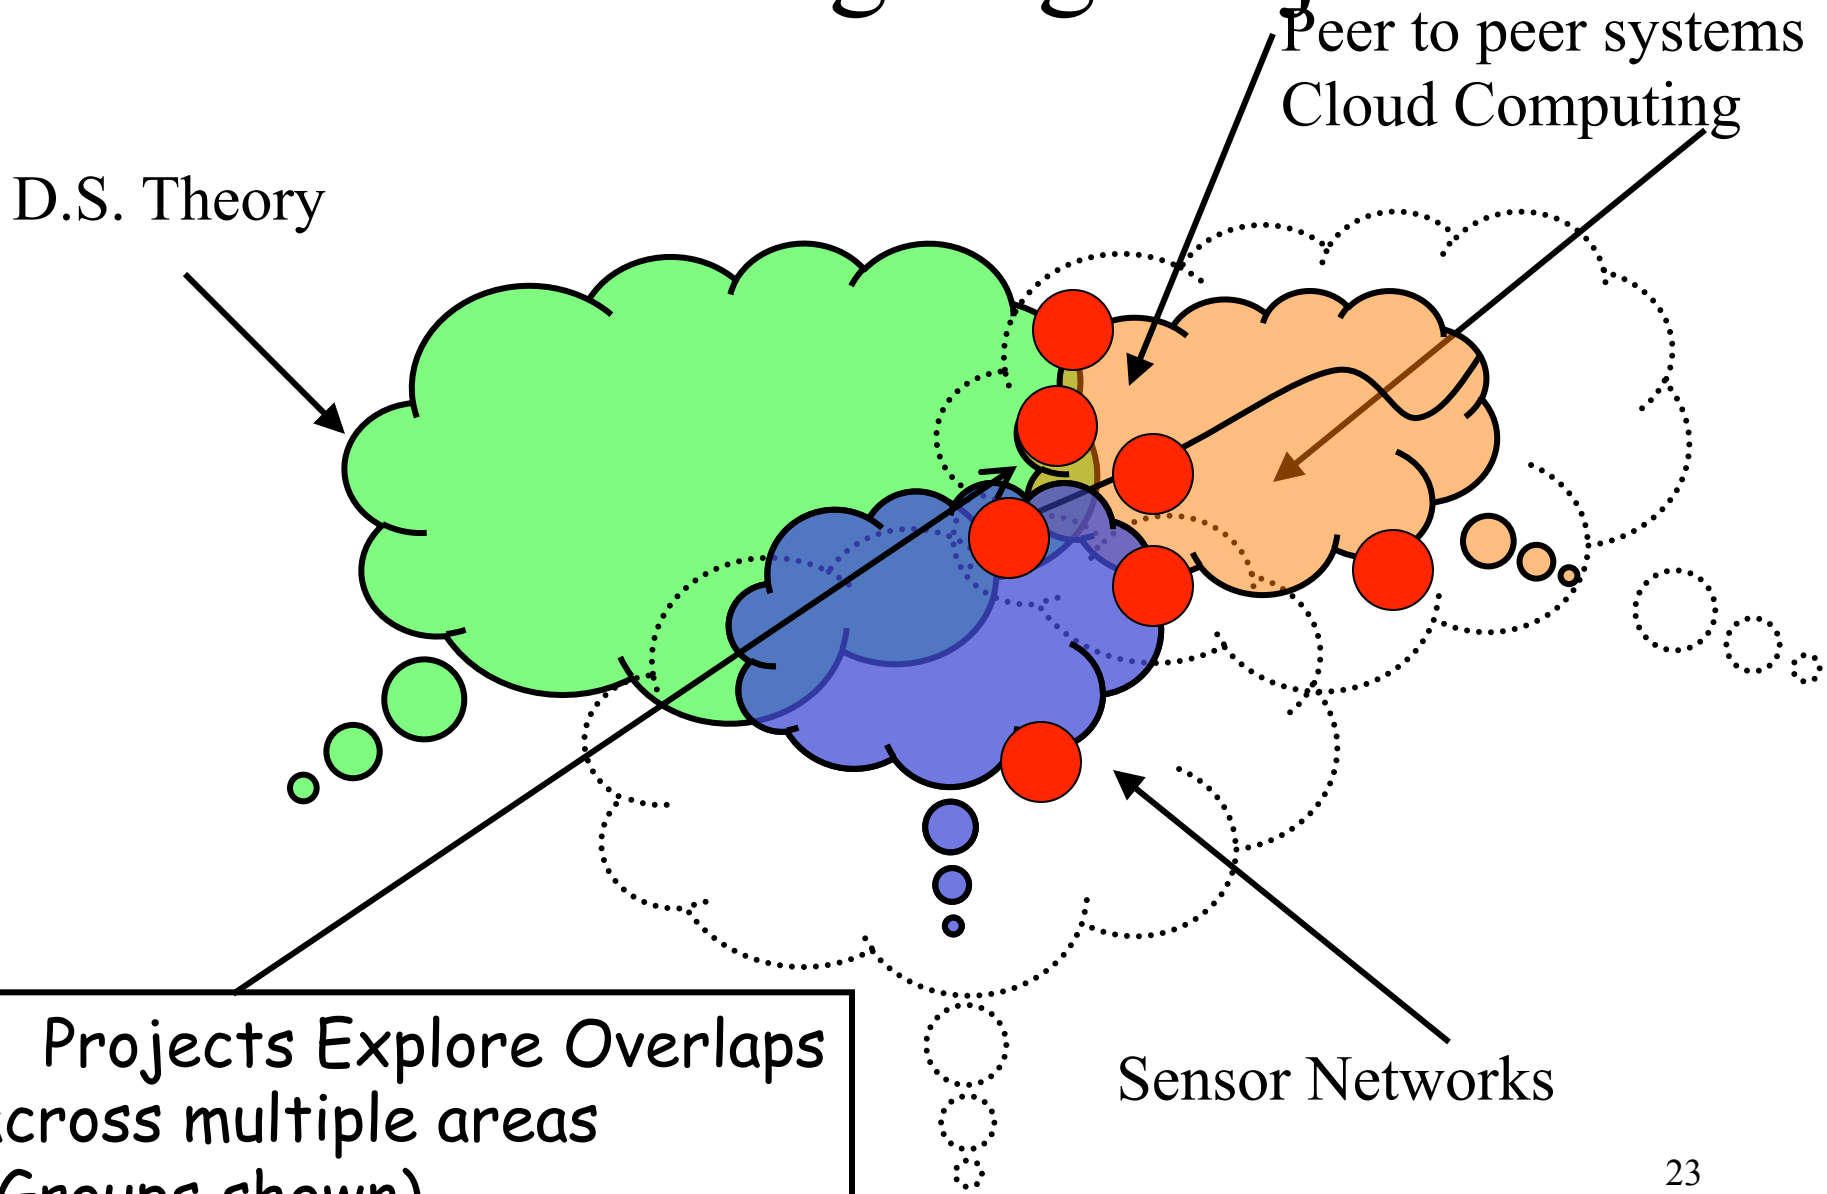

- Projects Explore Overlaps across multiple areas (Groups shown)

Leftover Work

- Final Project Report Submissions – 11.59 pm, Sunday May 8th, 2015 (email softcopy to indy@illinois.edu and lmlesli2@illinois.edu, turn hardcopy in to 3112 SC).
  - At most 12 pages (+ any extra pages for refs), at least 12 pt font + 1 page for Business Plan
  - 30% of course grade
- **Final extension, Hard deadline**
  - (should contain hard and comprehensive data)
- We will work on all projects after the semester, in order to submit them to conferences/workshops!
  - Past CS525 projects (since Fall 2003) have produced a total of more than 10 journal papers, 20 conference papers, and 10 workshop papers, and multiple best paper award winners (ICAC15, IC2E16, BigMine12)!

Winners

- Final Project – There will be Three Best Projects announced on website soon after the May 10th
- After the semester, all projects are eligible for continued work.

Office Hours for this week

- Today: only from now – 4 pm (in my office)
- Thursday 3 pm – 4 pm (in my office)

Presentations

I hope you liked the selection of papers.

Special mention presentations

- Lots of good ones! (difficult to pick “best ones”)

- General comments to all for future presentations:
  - Use figures and animations – a picture is worth a 1000 thousand words
  - Keep an eye on the clock
  - Defer questions to end or offline if necessary
  - Plan for > 1 minute per slide

Reviews

Tough work, but

only way to ensure you remember

main ideas in paper

and your initial thoughts on reading it

Please preserve your reviews!

I hope you enjoyed writing them.

If your complaint is about the large number of papers you had to review….

Reviews

Tough work, but

only way to ensure you remember

main ideas in paper

and your initial thoughts on reading it

Please preserve your reviews!

I hope you enjoyed writing them.

If your complaint is about the large number of papers you had to review….you're right

Articles

Articles for this Class

- H. G. Wells, “World Brain”
- G. Hardin, “The tragedy of the commons”
- Levin and Redell, “How (and how not to) write a good SOSP paper”

H. G. Wells

- H. G. Wells, "World Brain" (1938)
  - Encyclopedias in those days written "for gentlemen by gentlemen"
  - H. G. Wells envisioned a University that is world-wide, and a base of knowledge that is global
  - He sought a "Permanent World Encyclopedia"
    - That can be read by anyone anywhere
    - That can be updated by anyone and from anywhere
    - That will be an archive of humanity and its actions
    - That will be an extension of humanity's memory
  - And he wrote this before the Internet was invented!
  - Has this been realized?

  (article taken from book "World Brain," book published circa 1938)

G. Hardin

- G. Hardin, "The tragedy of the commons" (1968)
  - Adam Smith in 1776 in "*The Wealth of Nations*" popularized the "invisible hand," the idea that an individual who "intends only his own gain," is, as it were, "led by an invisible hand to promote…the public interest"
    - Basis for stock markets and much of today's economics!
  - However, if there is a *commons* (think: open pastures, stock market, Internet, p2p, clouds, national parks, etc.), then the tragedy is that everything will be depleted so much that nothing will stay common anymore
  - Example of free pastures for farmers with herds of sheep: "Each man is locked into a system that compels him to increase his herd without limit -- in a world that is limited."
  - Hardin concludes: "It is our considered professional judgment that this dilemma has no technical solution."
  - This essay motivated the development of game theory, and selfish models
  - The tragedy of the commons is very visible in p2p systems (freeloading). Does it also reflect in Wikipedia?
    - The argument says there are no technical solutions, which means you need to incentivize (or de-incentivize) humans to solve the problem
    - Oil spills, other environmental disasters (oceans and wild lands are "commons")
    - Do clouds like AWS suffer from this? They're so cheap they're practically free.

Levin-Redell

- Levin and Redell, "How (and how not to) write a good SOSP paper" (PC co-chairs of 1983 ACM SOSP symposium)
  - original **idea** to a **real problem**
  - comprehensive and mature **evaluation**
  - chronological and logical **presentation**

Questions?

CS525 Course Evaluations

- Main purpose: to evaluate how useful this course was to you (and to get your feedback that will help improve future versions of the course)
- I won't see these evaluations until after you see your grades
- For Question 5 (Gender), leave blank
- Use **pencil only** (box of pencils being passed around – please return after use)
- I need a volunteer to:
  1. Collect forms and mail them (via campus mail only!) to ICES
  2. Return pencils to Donna Coleman (2106 SC)
  3. Return uncollected reviews to me

All the Best for Your Project!
Have a good summer.