CSCI 357: Algorithmic Game Theory

Lecture 19: Incentives in BitTorrent

Shikha Singh

Announcements and Logistics

- HW 8 graded feedback will be returned soon
    - Clarification about single deviation and automaton
- Project proposal due tomorrow 5 pm
    - Google form that identifies partner, topic
    - ~150 words description of your idea
- You are required to sign up for a meeting to discuss the proposal
    - Many slots next week: https://tinyurl.com/357projectmeet
- 1-Page Abstract (LaTeX) due via Github next Friday (April 29) 5 pm
- Midterm 2 will be available on Gradescope: Sun Apr 24 8am - Tues Apr 26 8 am
    - No lecture on Monday April 25

Questions?

Midterm 2 Topics

- All topics are fair game, but focus on topics in second half
    - Revenue maximization and VCG
    - Competitive equilibrium
    - One-sided and two-sided matching
    - Voting and social choice
    - Sequential and repeated games

Questions?

Last Time

- Defined finite and infinite repeated games
- Formalized file sharing game as a repeated prisoner's dilemma
- Analyzed trigger strategies for repeated prisoner's dilemma
- Takeways:
    - TfT never wins a head-to-head match with another strategy or is a SPE with itself, but does extremely well empirically

Today

- Discuss incentives in P2P file sharing
- BitTorrent reference client based on TfT strategy
- Other strategic agents that take advantage of BitTorrent
- Guideline on final project

Incentives in BitTorrent

BitTorrent

- Inspired by repeated prisoner's dilemma and tit-for-tat strategy, Bram Cohen introduced the BitTorrent protocol

Incentives Build Robustness in BitTorrent

Bram Cohen
bram@bitconjurer.org

May 22, 2003

Abstract

The BitTorrent file distribution system uses tit-for-tat as a method of seeking pareto efficiency. It achieves a higher level of robustness and resource utilization than any currently known cooperative technique. We explain what BitTorrent does, and how economic methods are used to achieve that goal.

Cohen, B. (2003, June). Incentives build robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer systems* (Vol. 6, pp. 68-72).

BitTorrent

- BitTorrent protocol is the dominant paradigm in peer-to-peer file distribution
  - Responsible for **80% of total peer-to-peer traffic** in 2013
- The protocol is designed to distribute a large file (a movie, an OS, a big software patch etc.) to many users
- For example, you may have used it to download Linux versions

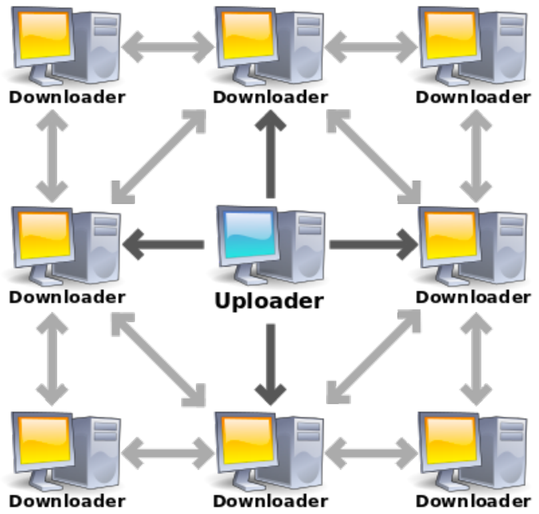

P2P File Sharing Setup

- A file is broken into **pieces** and (pieces may further be broken down into **blocks**)
  - E.g., 10 GB file broken into 10MB pieces
- Many reasons to break into smaller blocks and share at block level
  - More efficient use of network capacity
  - Changes it from a **one-shot PD** to a **repeated PD** with stages roughly corresponding to pieces of the file
- Downloads can often take a long time so there is room to track behavior of peers and respond accordingly
  - Room to strategize and have heuristics converge over time

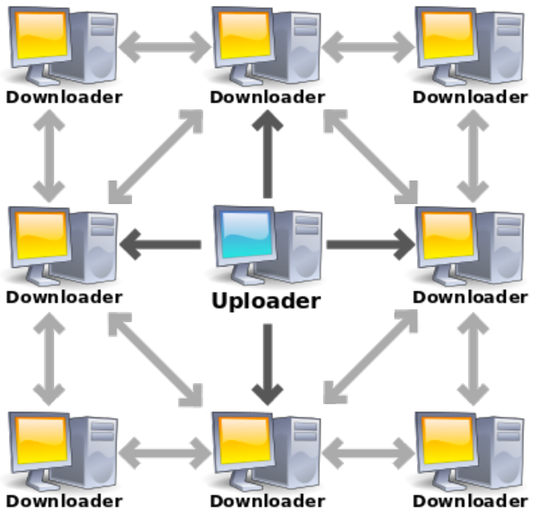

P2P File Sharing Setup

- There is at least one **Seed** (who has all the pieces of the file)
- Peers broadcast periodically & exchange information about who has which pieces
- Each peer sends **requests** to others for blocks of a particular piece
- Each peer also receives these requests from others and decides to **upload to** (**unchoke**) a subset of the requesting peers
    - May decided how much upload bandwidth to allocate to others
- After each round, download occur simultaneously from all peers who
    - you sent a request to and decided to unchoke you that round
- These exchanges happen until all peers have received all the pieces

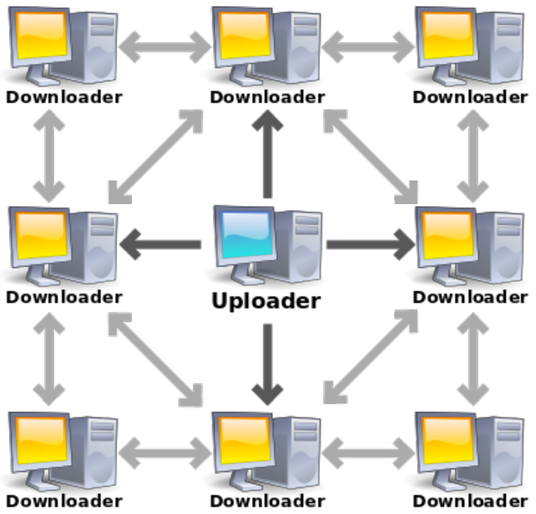

Request Strategy

- Determines which pieces to request from peers
- No harm in requesting from all peers who may have a piece you are interested in
- But some pieces may be more valuable than others
- Any ideas which one?
  - The pieces that are rare: not many peers have them yet
- BitTorrent client uses the **rarest-first request** strategy:
  - Request to download pieces with the **lowest availability**
- **Rarest first** strategy is designed to distribute rare pieces quickly and minimize bottlenecks

Upload Strategy

- Determines which requesting peers to unchoke, and how much upload bandwidth to assign to them
- Number of unchoke slots s is typically set to 4: means a user can upload to s of its peers in each round
  - Can choose to allocate bandwidth among them in any way, e.g., $1/s$ of its upload bw to each
- What would be a tit-for-tat strategy to determine who to upload to?
  - Peers who uploaded to you in the past
- What about the beginning when no one has pieces to share?
  - Need some way to get new peers "starter pieces" so they can share in the future

BitTorrent Upload Strategy

- BitTorrent reference client follows tit-for-tat **with optimistic unchoking**
- Given s unchoke slots, the reference client makes a new decision every round, unchoking $s - 1$ peers from whom it downloaded the most data in the past round
- Every three rounds, the reference client allocates an additional, **optimistic unchoking slot**, to a random peer (and leave this peer unchoked for three rounds)
- The optimistic unchoke is a “pity” upload to users who may not have much of the file yet
- The upload bandwidth is split equally among the s slots

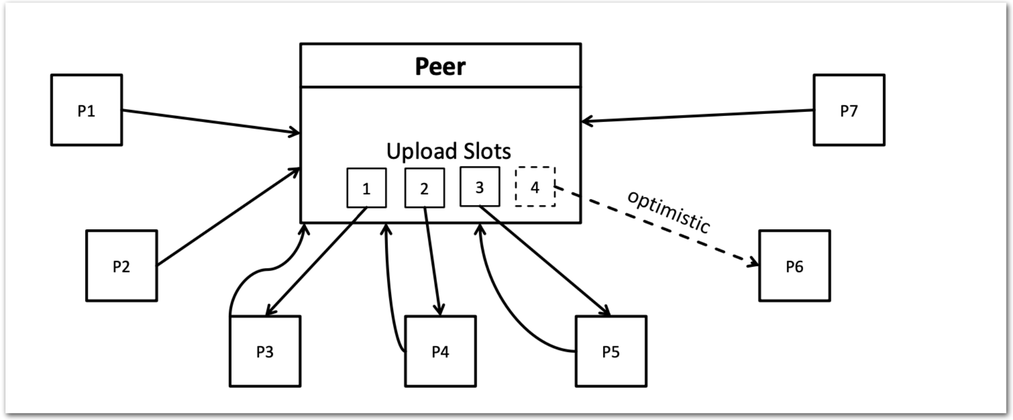

BitTorrent vs Repeated PD

- BitTorrent is inspired by repeated PD and Tit-for-Tat protocol
  - Unchoking strategy essentially reciprocates
  - Optimistic unhooking is essentially playing cooperate for the first time with a new opponent (providing a kind of conditional cooperation)
- Lot more going on than repeated prisoner's dilemma
  - Peers have more actions than just upload/not upload
  - E.g., they can vary "how much" to upload
  - More than one player so whether or not you get uploaded to depends on other competitors

Strategic Behavior in BitTorrent

Strategic Behavior on BitTorrent

- Many places for strategic behavior in P2P file-sharing
  - Piece-revelation: peers can be strategic about which pieces to reveal to others
  - Upload Bw: how to allocate upload bandwidth across peers by choosing number of upload slots/ how to distribute bandwidth
  - What pieces to allow an uncooked peer to download
  - What pieces to try to download, etc.
- Variations on how to handle each decision leads to different strategic clients, often designed to “game” BitTorrent

Strategic Behavior on BitTorrent

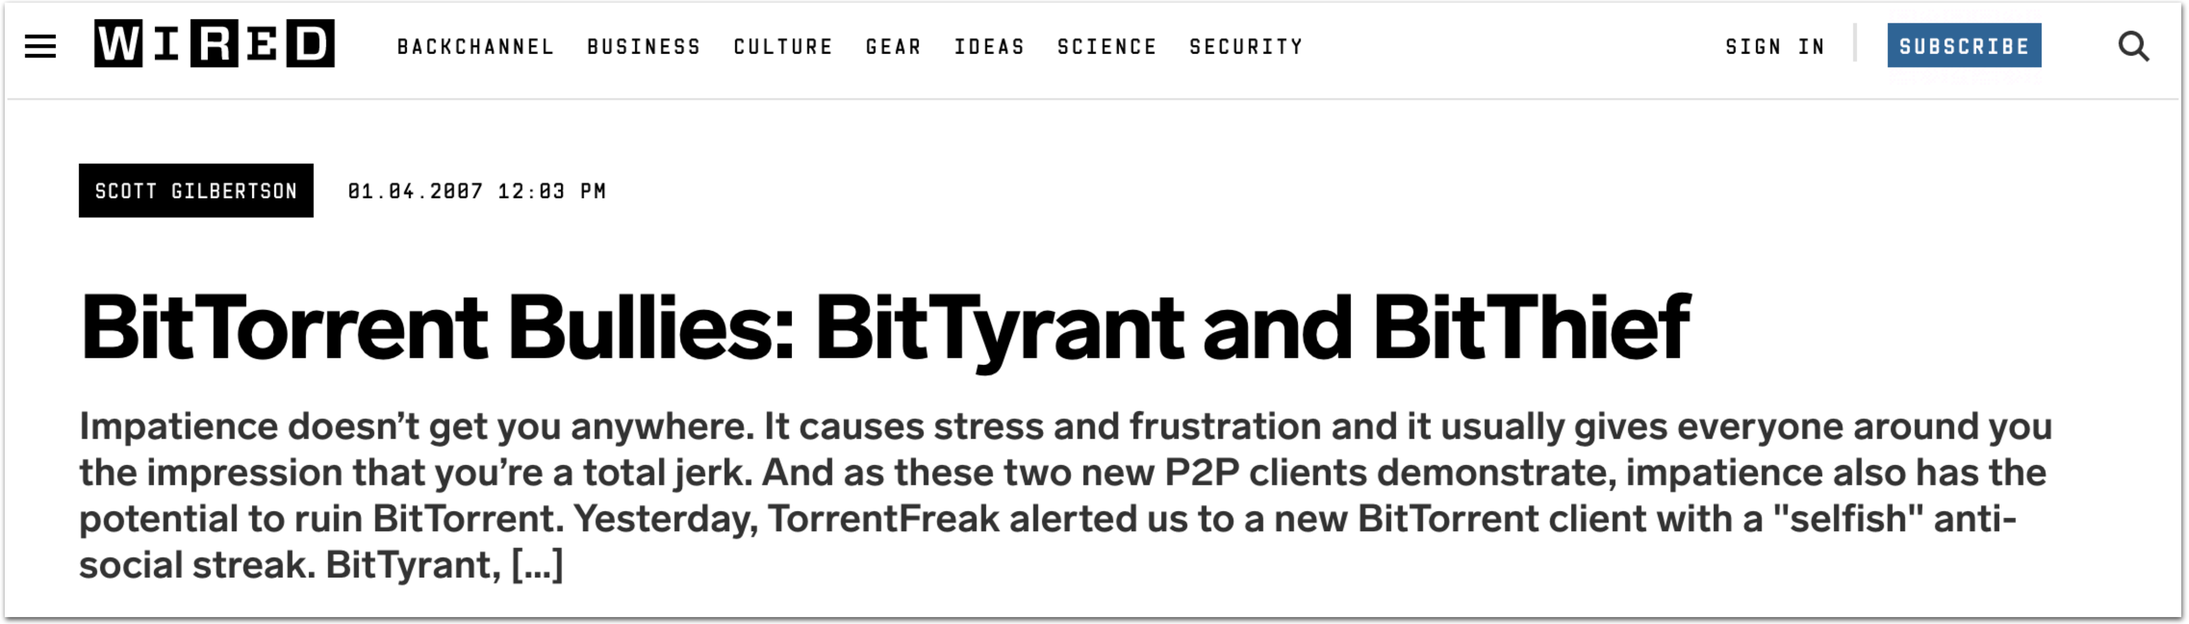

https://www.wired.com/2007/01/bittorrent-bullies-bittyrant-and-bitthief/

Bit Thief

- Unapologetically named BitThief client is to **leech off** a BitTorrent reference client
  - Download without ever uploading anything just like free riders in Gnutella
- The goal is to exploit the “optimistic unchokes”
- BitThief does this by pestering the tracker incessantly, asking for more peers to grow its neighborhood
- Downloads are slower than the reference client (because you don’t get the reciprocation advantage) but around 5x
- How to mitigate against such an incentive attack?
  - Have tracker ignore repeated requests in some window

Free Riding in BitTorrent is Cheap

Thomas Locher[1], Patrick Moor[2], Stefan Schmid[1], Roger Wattenhofer[1]
[1] Computer Engineering and Networks Laboratory (TIK), ETH Zurich, 8092 Zurich, Switzerland
{lochert, schmiste, wattenhofer}@tik.ee.ethz.ch
[2] Google Inc., Mountain View, CA 94043, USA
pmoor@google.com

Bit Tyrant

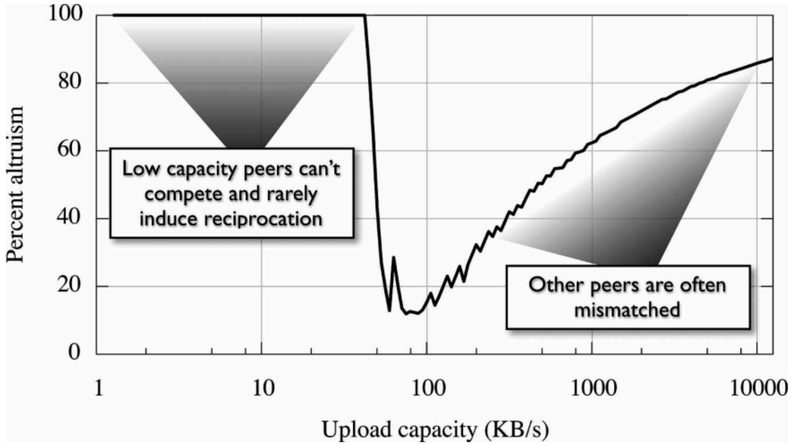

FIGURE 2: PERCENTAGE OF ALTRUISTIC UPLOAD CONTRIBUTION AS A FUNCTION OF CAPACITY

MICHAEL PIATEK, TOMAS ISDAL, TOM ANDERSON, ARVIND KRISHNAMURTHY, AND ARUN VENKATARAMANI

building BitTyrant, a (more) strategic BitTorrent client

Michael Piatek is a graduate student at the University of Washington. After spending his undergraduate years working on differential geometry, his research interests now include incentive design in distributed systems, network measurements, and large-scale systems building.

piatek@cs.washington.edu

Tomas Isdal graduated with a MSc in Computer Science and Engineering from the Royal Institute of Technology, Stockholm, Sweden, and is currently a graduate student in the Department of Computer Science and Engineering at the University of Washington. His interests include peer-to-peer and distributed systems, Internet measurements, and network security.

isdal@cs.washington.edu

Tom Anderson is a Professor in the Department of Computer Science and Engineering at the University of Washington. He is an ACM Fellow and a winner of the ACM SIGOPS Mark Weiser Award, but he is perhaps best known as the author of the Nachos operating system.

tom@cs.washington.edu

Arvind Krishnamurthy is an Assistant Research Professor at the University of Washington. His research interests are primarily at the boundary between the theory and practice of distributed systems. He has worked on automated mechanisms for managing overlay networks and distributed hash tables, network measurements, parallel computing, techniques to make low-latency RAID devices, and distributed storage systems that integrate the numerous ad hoc devices around the home.

arvind@cs.washington.edu

Arun Venkataramani has been an Assistant Professor at the University of Massachusetts Amherst since 2005, after receiving his Ph.D. from the University of Texas at Austin by way of the University of Washington. His research interests are in the practice and theory of networking and distributed systems.

arun@cs.umass.edu

PEER-TO-PEER SYSTEMS OFTEN APPEAL to scalability as a motivating feature. As more users request data, more users contribute resources. Scaling a service by relying on user contributions—the P2P approach—depends on providing incentives for users to make those contributions. Recently, the popular BitTorrent file distribution tool has emerged as the canonical example of an incentive-aware P2P design. Although BitTorrent has been in widespread use for years and has been studied extensively, we find that its incentive strategy is not foolproof. This article describes BitTyrant, a new, strategic BitTorrent client. For users interested in faster downloads, BitTyrant provides a median 70% performance improvement on live Internet swarms. However, BitTyrant also demonstrates that selfish users can improve performance even while reducing upload contribution, circumventing intended incentives.

Piatek, M., Isdal, T., Anderson, T., Krishnamurthy, A., & Venkataramani, A. (2007, April). Do incentives build robustness in BitTorrent. In *Proc. of NSDI* (Vol. 7).

Bit Tyrant

- Goal: use upload capacity strategically
- Each user i maintains two estimates about other peers:
  - d_{ij} : download capacity expected to get from peer j if i uploaded to j
  - u_{ij} : amount of upload i would need to get j to reciprocate back
    - If j is getting uploads from a lot of people, it may not be worth it
- Based on these estimates, BitTyrant uses a simple greedy strategy
  - Get most bang-for-buck (like in Knapsack)
  - Sort everyone based on ratio d_{ij}/u_{ij} and upload in this order until capacity runs out
- Unlike reference client, does not split bandwidth equally

Strategic Piece Revelation

- Goal of keeping other peers interested in exchanging pieces for as long as possible
- If peer i has a piece that j wants, then we say that j is interested in i. Once j has all of the pieces that i has, it loses interest in i
- A peer could try to maximize the time others are interested in it by strategically revealing pieces
- BitTorrent client always reveals all the pieces it has
- Does underreporting ever help?

Levin, Dave, et al. "Bittorrent is an auction: analyzing and improving bittorrent's incentives." *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. 2008.

BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives

Dave Levin
University of Maryland
dml@cs.umd.edu

Katrina LaCurts
University of Maryland
katrina@cs.umd.edu

Neil Spring
University of Maryland
nspring@cs.umd.edu

Bobby Bhattacharjee
University of Maryland
bobby@cs.umd.edu

ABSTRACT

Incentives play a crucial role in BitTorrent, motivating users to upload to others to achieve fast download times for all peers. Though long believed to be robust to strategic manipulation, recent work has empirically shown that BitTorrent does not provide its users incentive to follow the protocol. We propose an auction-based model to study and improve upon BitTorrent's incentives. The insight behind our model is that BitTorrent uses, not tit-for-tat as widely believed, but an auction to decide which peers to serve. Our model not only captures known, performance-improving strategies, it shapes our thinking toward new, effective strategies. For example, our analysis demonstrates, counter-intuitively, that BitTorrent peers have incentive to intelligently under-report what pieces of the file they have to their neighbors. We implement and evaluate a modification to BitTorrent in which peers reward one another with proportional shares of bandwidth. Within our game-theoretic model, we prove that a proportional-share client is strategy-proof. With experiments on PlanetLab, a local cluster, and live downloads, we show that a proportional-share unchoker yields faster downloads against BitTorrent and BitTyrant clients, and that under-reporting pieces yields prolonged neighbor interest.

P2P in Practice

- Despite various ways to attack and deviate from the reference client, BitTorrent has been very successful
- Private BitTorrent communities:
  - Many private BitTorrent communities exist with trackers only accessible to users within the community
  - Typically such communities enforce sharing ratios, a minimum ratio of upload to download traffic per user
  - Understanding the dynamics in private tracker communities is an active area of research: e.g., are sharing ratios a good metric or should credit-based system be used?
- There also exist altruistic BitTorrent client

Project Overview

CS357 Project

- Analyze incentives inherent in a system or class of algorithms by applying the frameworks developed in class
- Goals of the final project:
  - Dig deeper into a topic in AGT
  - Read and understand **research papers** in the field
  - Apply the framework developed in class to new domains
- Ideal: groups of 2
- Scope: 3 week project
  - Effort should match the timeframe: not too ambitious, not too light

Deliverables

- Proposal (google form, 150 word proposal): due Fri April 22, 5 pm
    - Identify topic and partner
    - High-level ideas and how you will accomplish them
- 1 page LaTeX abstract: due Fri April 29 via Github
    - Formalize your project, include references and concrete plan
- 2 pages LaTeX report including research background: due Thur May 5
- Project presentations: last week of classes: due May 12 ~5-10 slides
- Final project report (at least 6 pages) due May 19
- **Check ins:** Meet with me at least twice for 20 mins

Suggested Project

- Each project must explore the **role of incentives and strategic behavior** in a market/ system or a class of algorithms
- Suggested project: Understanding Strategic Behavior in BitTorrent
  - **BitTorrent simulator** will be provided
  - Implement BitTorrent reference client
  - Read related research papers and implement strategic clients
    - BitTyrant: Do incentives build robustness in BitTorrent? by Piatek et al.
    - BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives by Levin et al.
- Open ended: design your own client, perform advanced analysis to compare existing strategic behavior

Matching Markets

- Lots of readable research papers related to deferred acceptance and matching
    - Analyze the role of strategic behavior
- Fairness in matching:
    - School choice, or gale shapely
- What if we focused on a different concept than stability?
    - Papers look at "globally" notions called popular matchings
    - Popular matchings are bigger in size than stable matchings
    - However, the incentives in popular matching algorithms are not well understood

Voting and Social Choice

- Study how different voting rules are to manipulate "on average"
- If we are care about a particular metric, are some voting rules better than others?
- Well known ways to generate "simulated voting data"
    - Mallows model and other probability distributions
- What if we had made all voting rules compete in a game?
    - GT Optimal Voting paper by Rivest studies mixed-Nash of such a game
- Social choice in other domains:
    - Tournament design or social rankings
    - Dividing resources fairly etc

Guide to Choosing a Topic

- Think about algorithms that you care about that may not be well designed in incentives with in mind
- Pick a topic you will enjoy working on (something fun!)
- Pick something that you will learn a lot from (something useful!)
- Pick something that is relevant to the course (something relevant!)
- Keep the scope in mind, make sure you have something that you can make progress on in ~3 weeks
    - Something not too ambitious but not too trivial
- Try to find CS/Econ papers related to the topic you have in mind
    - Topic must be technically interested, related to concepts in class

Topic Ideas and Papers

- Will post a more detailed summary of suggested project topic
- Will post a bunch of **recent relevant research papers**
- Some ideas of topics you can explore
- If you have an idea, propose it and we can discuss if it is viable
    - I can suggest research papers related to topic
    - Or ask to change it if needed
- Projects can be mostly implementation or mostly theoretical or between
    - BitTorrent project is very implementation heavy
- **Example projects from F2**0 posted on GLOW: check them out for reference

Theory vs Implementation

- Project must involve reading and understanding **at least two research papers**
    - And citing other relevant literature
- Even implementing an algorithm, important to understand the theory and concepts
- Purely theory: plan on going above and beyond just surveying existing literature:
    - New insights/observations or conjectures
    - Filling in important details, examples or cases that may have been overlooked
    - Simulating strategic behavior can help analyze situations where theoretical analysis may be tricky
- Always try to identify what your contribution in the project is going to be

Project Rubric

- Linked on the webpage, and consists of:
    - Scholarship / Background (5 pts)
    - Contributions (5 pts)
    - Correctness (5 pts)
    - Creativity (5 pts)
    - Presentation (10 pts)

Paper Reading Advice

- Use the three-pass approach suggested by Keshav
- First pass:
    - Read the abstract, introduction, conclusion and boldface
    - What information you can glean from this:
        - What are the paper's main contributions?
        - Is the paper well written?
        - Main keywords that the paper is about
- At this point, you should be able to identify the "closest" related work to the current paper: which can often be a useful resource

Paper Reading Advice

- Second pass:
    - Read the paper more thoroughly (but ignore proofs)
    - Pay special attention to theorem/lemma statements or graphs
        - What assumptions are the authors making?
        - Do the definitions used make sense?
            - Create examples and identify the defined notions?
- Start writing down questions about what you don't understand
    - It may be the answers are hidden in the paper but not obvious
- Take lots of notes! Use your own examples to explain the work

Paper Reading Advice

- Third pass: Virtually reimplement the paper
    - For implementation project this is literal; for theory projects, this means understand the concepts as well as if you coded it up
- At this stage, you should have your own "perspective" on the work
    - How would you do things differently if you wrote the paper?
    - Is there a notion the authors do not explain well, that you can supplement through your project?
    - How would you teach this paper to the class?
    - What are the most interesting and challenging aspects?
    - What new directions can this work take if you had more time?

Literature Survey

- For your technical writeup, you'd need to summarize the closely related literature to the topic you picked
- How should you approach it?
    - First, pay close attention to what the authors identify as the most related work to theirs
    - Follow "the trail of citations"
    - Find what's the state of the art is on a topic
    - Use google scholar!