

Internet Economics

Networked Life

NETS 112

Fall 2013

Prof. Michael Kearns

The Internet is an Economic System (whether we like it or not)

- Highly decentralized and diverse
 - allocation of scarce resources; conflicting incentives
- Disparate network administrators operate by local incentives
 - network growth; peering agreements and SLAs
- Users may subvert/improvise for their own purposes
 - free-riding for shared resources (e.g. in peer-to-peer services)
 - spam and DDoS as economic problems
- Regulatory environments for networking technology
 - for privacy and security concerns in the Internet
 - need more "knobs" for society-technology interface

Can Economic Principles Provide Guidance?


- **Game theory and economics**, competitive and cooperative
 - strategic behavior and the management of competing incentives
- **Markets** for the exchange of standardized resources
 - goods & services
 - efficiency and equilibrium notions for performance measurement
- **Learning and adaptation** in economic systems
- Certain **nontraditional topics** in economic thought
 - behavioral and agent-based approaches
- Active research at the **CS-economics** boundary

The Internet: What is It?

- A massive network of connected but decentralized computers
- Began as an experimental research NW of the DoD (ARPAnet), 1970s
 - note: Web appeared considerably later
- All aspects evolved over many years
 - protocols, services, hardware, software
- Many individuals and organizations contributed
- Designed to be *open, flexible, and general* from the start
 - “layered” architecture with progressively strong guarantees/functionality
 - layers highly modular, promotes clean interfaces and progressive complexity
 - highly agnostic as to what services are provided
- Completely unlike prior centralized, managed NWs
 - e.g. the AT&T telephone switching network

Internet Basics

- Can divide all computers on the Internet into two types:
 - computers and devices at the "edge"
 - your desktop and laptop machines
 - big compute servers like Eniac
 - your web-browsing cell phone, your Internet-enabled toaster, etc.
 - computers in the "core"
 - these are called *routers*
 - they are very fast and highly specialized; basically are big switches
- *Every* machine has a unique Internet (IP) address
 - IP = Internet Protocol
 - like phone numbers and physical addresses, IP addresses of "nearby" computers are often very similar
 - your IP address may vary with your location, but it's still unique
- IP addresses are how everything finds everything else!
- Note: the Internet and the Web are *not* the same!
 - the Web is one of many services that run on the Internet



Full Internet map as of 18 Feb 1999
99664 edges, 88107 nodes (42443 leaves)

Hal Burch, Bill Cheswick
Printed 1999-03-23 15:25:11. <http://www.cs.bell-labs.com/~cheswick/index.html>

Internet Packet Routing

- At the lowest level, all data is transmitted as *packets*
 - *small* units of data with addressing and other important info
 - if you have large amounts of data to send (e.g. a web page with lots of graphics), it must be *broken* into many small packets
 - somebody/thing will have to reassemble them at the other end
- All routers do is *receive* and *forward* packets
 - forward packet to the "next" router on path to destination
 - they only forward to routers they are *physically* connected to
 - how do they know which neighboring router is "next"?
- Routing tables:
 - giant look-up tables
 - for each possible IP address, indicates which router is "next"
 - e.g. route addresses of form 128.8.*.* to neighbor router A
 - route 128.7.2.* to neighbor router B, etc.
 - need to make use of *subnet addressing* (similar to zip codes)
 - distributed maintenance of table consistency is complex
 - must avoid (e.g.) cycles in routing
 - requires distributed communication/coordination among routers
- Handy programs: *ipconfig*, *tracert*, *ping* and *nslookup*

The IP (Internet Protocol)

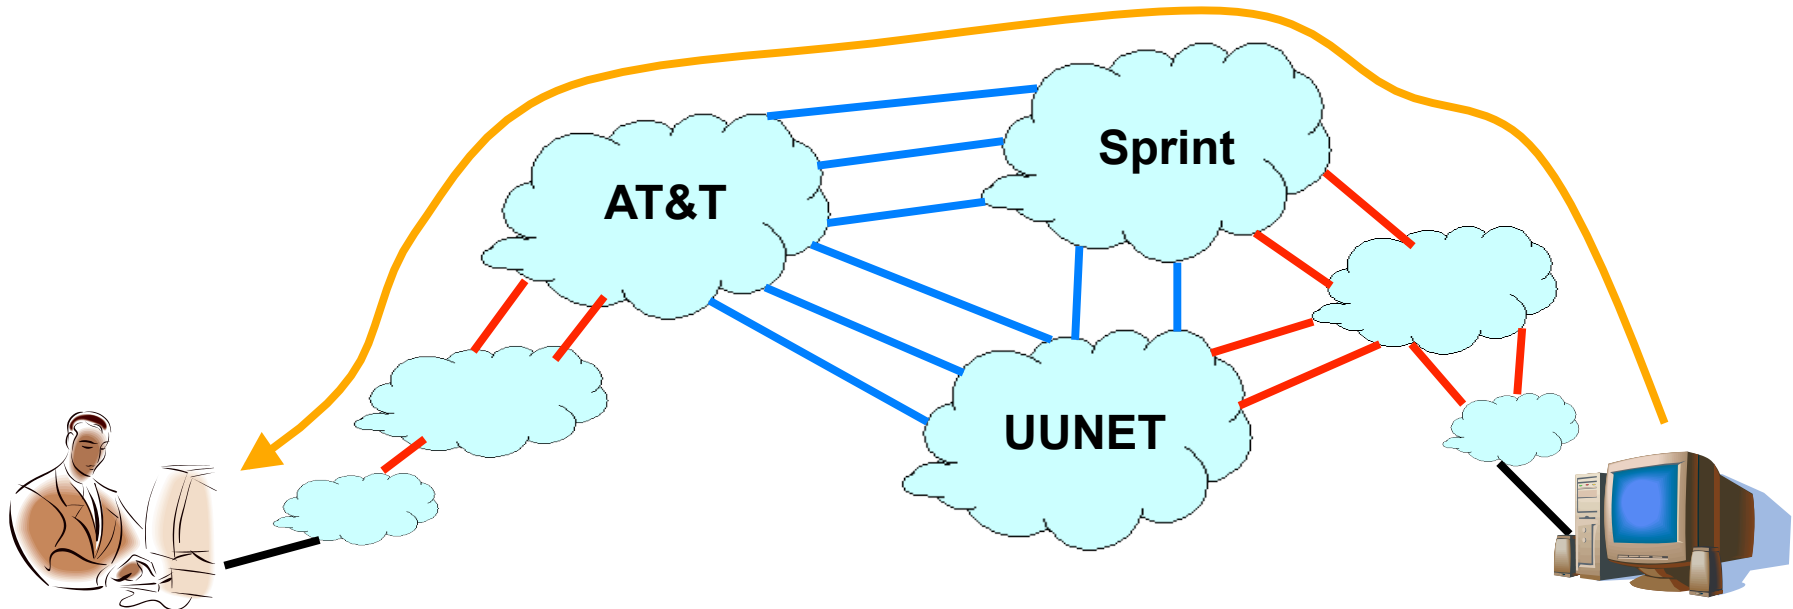
- There are many possible conventions or protocols routers could use to address issues such as:
 - what to do if a router is down?
 - who worries about lost packets?
 - what if someone wants their packets to move faster?
- However, they all use a single, simple protocol: IP
- IP offers only one service: "best effort" packet delivery
 - with *no* guarantee of delivery
 - with *no* levels of service
 - with *no* notification of lost or delayed packets
 - knows *nothing* about the applications generating/receiving packets
 - this simplicity is its great strength: provides robustness and speed
- Higher-level protocols are *layered* on top of IP:
 - TCP: for building connections, resending lost packets, etc.
 - http: for the sending and receiving of web pages
 - ssh: for secure remote access to edge computers
 - etc. etc. etc.

Autonomous Systems (ASes)

- Q: So who owns and maintains all these routers?
- A: Networking companies/orgs called "Autonomous Systems"
- ASes come in several different flavors:
 - large, long-haul "backbone" network providers (AT&T, UUNET, Sprint)
 - consumer-facing Internet Service Providers (ISPs) (Comcast, Earthlink)
 - companies/organizations needing to provide Internet access to members (Penn)
- The path of a "typical" packet would usually travel through many ASes
 - email, web page request, Skype call,...
- Q: How do the ASes make money?
- A: Some do, some don't
 - consumers and organizations near the edge pay their ISP/upstream provider
 - ISPs may in turn pay backbone providers
 - backbone providers typically have "peering agreements"
- Let's revisit traceroute...
- Q: How do the ASes coordinate the movement/handoff of traffic?
- A: It's complicated... we'll return to this shortly.

Commercial Relationships in Internet Routing

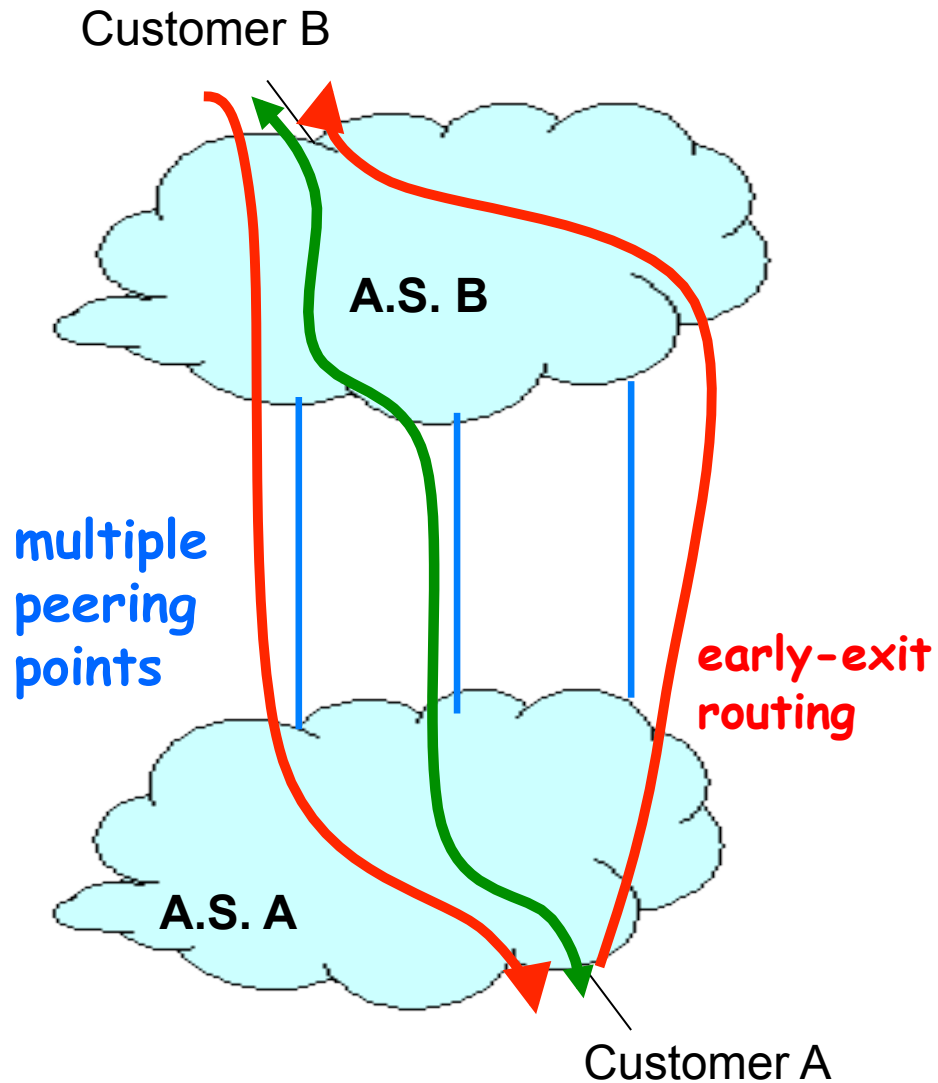
- **Customer-Provider**
 - customer **pays** to send and receive traffic
 - provider transits traffic to the rest of Internet
- **Peer-peer**
 - **settlement free**, under near-even traffic exchanges
 - transit traffic to and from their respective customers
- These are **existing economic realities**
- They create specific economic **incentives** that must co-exist with technology, routing protocols, etc.



Border Gateway Protocol (BGP)

- Within its *own network*, an AS may choose to route traffic as it likes
 - typically might follow a shortest path between the entry router and the exit router
- Interfaces *between* ASes are formed by special *border routers*
 - these are the routers where a packet travels from one AS to the "next"
- Communication at border routers governed by the Border Gateway Protocol:
 - border routers "announce" paths to neighboring ASes
 - e.g. "I have a 13-hop path through my AS to www.cis.upenn.edu"
 - ASes use neighboring announcements to decide where to forward traffic & determine own paths
 - paths actually specify *complete list* of ASes: e.g. 13-hop path Comcast → AT&T → UUNET → Penn
- Fair amount of *trust* and *honesty* expected for effective operation of BGP
- What are the incentives to cheat or deviate from expected behavior?
 - announce false paths to get more traffic
 - announce false paths to omit
 - deliberately avoid shortest announced path (UUNET is my competitor, don't give them traffic)
- Very recent research: try to make announced paths truthful
 - crypto/security approach: monitor/measure announced vs. actual paths
 - very difficult, high overhead
 - alternative approach: game theory
 - establish conditions under which "rational" ASes will announce truthful paths
 - rational: use announced paths which give best route to outbound traffic; announce paths which will maximize revenue


Economic Incentives for Peering



- How to select peers?
 - need to reach some other part of the Internet
 - improve end-to-end customer performance
 - avoid payments to upstream providers
- How to route the traffic?
 - today: **early-exit** routing to use less bandwidth
 - tomorrow: **negotiate** for lower total resource usage?


Game Theory of Internet Routing

- Strong analogy between routing and driving on a network of roads
 - each driver has their own starting (source) point and ending (destination) points
 - each driver (packet flow) wants to minimize their own latency
 - each driver chooses their sequence of roads ("source" vs. default routing)
 - delays on each road depend on how much traffic they carry
- Very similar to navigation problem in social networks, but now:
 - network is **technological** instead of social
 - **many** source/destination pairs instead of one
 - flows are **selfish**
- Formalize as a **game** on a network:
 - network: network of roads or routers
 - players: individual drivers or traffic flows
 - payoff for a player: negative of their total driving time
 - assume delay on each road proportional to traffic
- Huge number of players; huge number of possible actions
 - actions: all possible routes from source to destination
 - still, we know there is a Nash equilibrium...
- What could we hope to say?

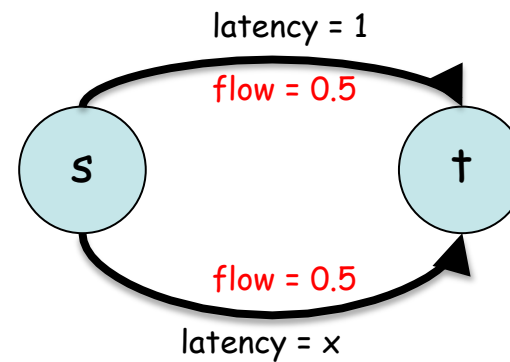


Routing Equilibrium Example


- Suppose we have only two roads/connections in the network:
 - "normal" road: delay/latency is equal to the amount of traffic x
 - "mountain" road: delay/latency is 1 unit no matter how much traffic
- Imagine 1 fully divisible unit of traffic that wants to travel from s to t :



At equilibrium, all traffic takes the normal road and everyone has latency = 1




A better collective solution: half the population has latency 0.5, half has latency 1... But upper flow is envious



Braess's Paradox


Initial Network:



Delay = 1.5


Braess's Paradox

Initial Network:



Delay = 1.5


Augmented Network:



Now what?


Braess's Paradox

Initial Network:



Delay = 1.5

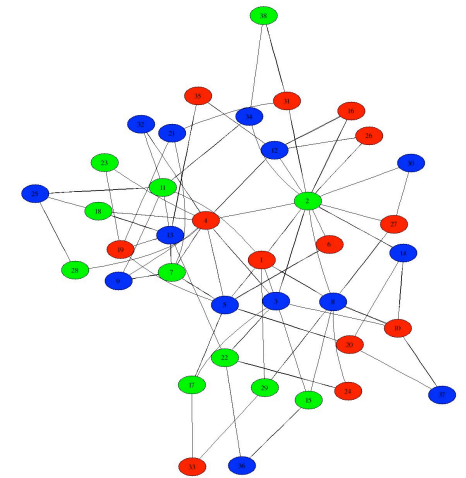
Augmented Network:



Delay = 2

The Price of Anarchy

- In principle (only), could imagine computing a **centralized** solution
 - "Centralized Traffic Authority" assigns each driver/flow their route
 - does so to minimize total population latency; may not be optimal for individuals
 - "maximum social welfare" solution; game-theoretic equilibrium can only be worse
- Surprising result: total latency of Nash equilibrium only 33% worse!
 - no matter how big or complex the network
 - "Price of Anarchy" (selfish, distributed behavior) is relatively small
 - compare to Prisoner's Dilemma
 - network structure irrelevant; contrast earlier results (e.g. networked trading)
 - can be worse than 33% for more complex latency assumptions



Case Study: QoS


- QoS = Quality of Service
 - many varying services and demands on the Internet
 - email: real-time delivery not critical
 - chat: near real-time delivery critical; low-bandwidth
 - voice over IP: real-time delivery critical; low-bandwidth
 - teleconferencing/streaming video: real-time critical; high-bandwidth
 - varying QoS guarantees required
 - email: not much more than IP required; must retransmit lost packets
 - chat/VoIP: two-way connection required
 - telecon/streaming: high-bandwidth two-way connections
- Must somehow be built *on top* of IP
- Whose going to pay for all of this? How much?
 - presumably companies offering the services
 - costs passed on to their customers
- What should the protocols/mechanism look like?
- There are many elaborate answers to these questions...

QoS and the Paris Metro

- Paris Metro (until recently)
 - two classes of service: first (expensive) and coach (cheaper)
 - exact same cars, speed, destinations, etc.
 - people pay for first class:
 - because it is less crowded
 - because the type of person willing/able to pay first class is there
 - etc.
 - self-regulating:
 - if too many people are in first class, it will become less attractive
- Andrew Odlyzko's protocol for QoS:
 - *divide* the Internet into a small number of identical virtual NWs
 - simply charge different prices for each
 - an entirely economic solution
 - California toll roads


Case Study: Sponsored Search

- Organic vs. sponsored web search
- Generalized second price auctions
- Two-sided networked markets



Organic vs. Sponsored Web Search

- Already (briefly) studied **organic** web search:
 - use words in user's query and web sites to rank results
 - other, non-language features also important
 - our emphasis: PageRank algorithm for web site importance
- **Sponsored** web search: a market/auction for ad placement
 - user query may signal "purchasing intent"
 - advertisers bid/compete for attention
- Rules of auction broadly similar across search engines
 - Google, Bing, Yahoo!
- We'll describe these auctions and their properties





philadelphia mountain bike



Sign in

Search

About 2,550,000 results (0.15 seconds)

SafeSearch on



Web

Ad related to philadelphia mountain bike

Why this ad?

[\(484\) 469-7025 - Men's & Women's Mountain Bikes](#)

www.cyclesportmedia.com/

Popular Models To Choose From.

Images

Maps

Videos

News

Shopping

More

Philadelphia, PA

Change location

Show search tools

[\[PMBA\] Philadelphia Mountain Biking Association](#)

www.phillymtb.org/

Saturday, July 14 will be PMBA's support day for Launch **Bike** Park. Come on out and check out all the **mountain** has to offer.... There will be downhill **bikes** to ...

[Forums](#)

A place for PMBA memers to post their stuff for sale. (Must be a ...

[About Us](#)

About Us. PMBA, the local IMBA Chapter for the Philadelphia ...

[Events](#)

Sunday, Monday, Tuesday, Wednesday, Thursday, Friday ...

[Merch](#)

Home · About Us · Forums · Events · News · Merch. Social / RSS ...

[More results from phillymtb.org »](#)

[Philadelphia Mountain Biking Meetup \(Philadelphia, PA\) - Meetup](#)

www.meetup.com/mountainbiking/

We meet to ride, have fun, and help others improve their skills. We ride ...

Mon, Jul 16 [Pennypack - longer, faster, harder](#)

Tue, Jul 17 [Core Creek "All levels ride ... - Core Creek Park Duchess ...](#)

Wed, Jul 18 [Wissahickon Park - Wissahickon- Livezey Pavillion Henry Ave ...](#)

[Wissahickon Trails, Philadelphia PA - Mountain Bike Bill](#)

www.mountainbikebill.com/pa-wissahickon.htm

Mountain Biking the Wissahickon trails in Philadelphia PA.

[Philadelphia Area Mountain Biking - Mountain Bike Trails](#)

www.trails.com/activity.aspx?area=10200

10+ items -- Get detailed descriptions and maps for Pennsylvania ...

Fairmount Park Forbidden Drive 10 miles

Impoundment Loop 3.5 miles

[philadelphia bicycles - by owner classifieds - craigslist](#)

Ads - Why these ads?

[Philadelphia Bike Trails](#)

www.trailink.com/

Find Detailed Trail Descriptions, Trail Maps, Photos and Reviews!

[Bikes Up To 60% Off List](#)

www.bikesdirect.com/

★★★★★ 637 seller reviews

New Bicycles. Full Factory Warranty Buy Direct. Save Big. Free Shipping

[Mountain Bikes On Sale](#)

www.giantnerd.com/MountainBikes

★★★★★ 248 reviews for giantnerd.com

Free Shipping on **Mountain Bikes**. 365 Day Returns -- Powered by Love!

[Privacy](#)

[Mountain Bike](#)

www.rei.com/Mountain-Bikes

★★★★★ 5,964 reviews for rei.com

Full Suspension -- Hardtail.

We've Got -- 50+ Models.

[Mountain Bikes](#)

www.target.com/

★★★★★ 1,408 reviews for target.com

Mountain Bikes Online.

Shop Target.com.

[Bucks County Bicycle Co.](#)

www.buckscountybikes.com/

Winter Service Special: \$69.99


Tune-up Plus Coode: gpwss

1,005 people +1'd this page

8545 New Falls Road, Levittown (215) 946-7090 - [Directions](#)


How Does It Work?

- Interested advertisers submit their bids for a query
 - \$0.25 for "philadelphia mountain bike", \$0.17 for "philadelphia discount mountain bike"
- Search engine gathers all the bids and determines advertiser ranking
- Advertisers **only pay if a user clicks** on their ad
 - "price per click" (PPC)
 - distinguishes from **display** advertising
- They may pay **less** than what they bid



Generalized Second Price Auctions

- Multiple bidders for a single item
 - each bidder i has a private valuation $v(i)$ for the item
 - each bidder i privately submits a bid $b(i) \leq v(i)$ for the item
- If you give the item to the highest bidder at their bid, everyone will bid less than their valuation
 - bid "shaving"
- If you give the item to the highest bidder, but only make them pay the second highest bid, the optimal strategy is to be "truthful"
 - all $b(i) = v(i)$
- Search engines rank advertisers by their bids
- Advertiser's PPC is the bid **below** them





philadelphia mountain bike



Sign in

Search

About 2,550,000 results (0.15 seconds)

SafeSearch on



Web

Images

Maps

Videos

News

Shopping

More

Philadelphia, PA

Change location

Show search tools

Ad related to philadelphia mountain bike

Why this ad?

[\(484\) 469-7025 - Men's & Women's Mountain Bikes](#)

www.cyclesportmedia.com/

Popular Models To Choose From.

\$0.53

[\[PMBA\] Philadelphia Mountain Biking Association](#)

www.phillymtb.org/

Saturday, July 14 will be PMBA's support day for Launch **Bike** Park. Come on out and check out all the **mountain** has to offer.... There will be downhill **bikes** to ...

[Forums](#)

A place for PMBA memers to post their stuff for sale. (Must be a ...

[About Us](#)

About Us. PMBA, the local IMBA Chapter for the Philadelphia ...

[Events](#)

Sunday, Monday, Tuesday, Wednesday, Thursday, Friday ...

[Merch](#)

Home · About Us · Forums · Events · News · Merch. Social / RSS ...

[More results from phillymtb.org »](#)

[Philadelphia Mountain Biking Meetup \(Philadelphia, PA\) - Meetup](#)

www.meetup.com/mountainbiking/

We meet to ride, have fun, and help others improve their skills. We ride ...

Mon, Jul 16 [Pennypack - longer, faster, harder](#)

Tue, Jul 17 [Core Creek "All levels ride ... - Core Creek Park Duchess ...](#)

Wed, Jul 18 [Wissahickon Park - Wissahickon- Livezey Pavillion Henry Ave ...](#)

[Wissahickon Trails, Philadelphia PA - Mountain Bike Bill](#)

www.mountainbikebill.com/pa-wissahickon.htm

Mountain Biking the Wissahickon trails in Philadelphia PA.

[Philadelphia Area Mountain Biking - Mountain Bike Trails](#)

www.trails.com/activity.aspx?area=10200

10+ items -- Get detailed descriptions and maps for Pennsylvania ...

Fairmount Park Forbidden Drive 10 miles

Impoundment Loop 3.5 miles

[philadelphia bicycles - by owner classifieds - craigslist](#)

Ads - Why these ads?

[Philadelphia Bike Trails](#)

www.trailink.com/

Find Detailed Trail Descriptions, Trail Maps, Photos and Reviews!

\$0.47

[Bikes Up To 60% Off List](#)

www.bikesdirect.com/

★★★★★ 637 seller reviews

New Bicycles. Full Factory Warranty Buy Direct. Save Big. Free Shipping

\$0.42

[Mountain Bikes On Sale](#)

www.giantnerd.com/MountainBikes

★★★★★ 248 reviews for giantnerd.com

Free Shipping on **Mountain Bikes**.

365 Day Returns -- Powered by Love!

\$0.25

[Privacy](#)

[Mountain Bike](#)

www.rei.com/Mountain-Bikes

★★★★★ 5,964 reviews for rei.com

Full Suspension -- Hardtail.

We've Got -- 50+ Models.

\$0.24

[Mountain Bikes](#)

www.target.com/

★★★★★ 1,408 reviews for target.com

Mountain Bikes Online.

Shop Target.com.

\$0.11

[Bucks County Bicycle Co.](#)

www.buckscountybikes.com/

Winter Service Special: \$69.99

Tune-up Plus Coode: gpwss


1,005 people +1'd this page

\$0.09

8545 New Falls Road, Levittown (215) 946-7090 - [Directions](#)


Other Details

- Actually order advertisers by combination of bids and “quality scores”
 - e.g. incorporate click-through rates (CTRs); higher CTRs boosted in ranking
 - prevents display of high bidders who never receive clicks
 - reduces irrelevant advertisers
- Search engines sometimes employ reserve prices
 - e.g. minimum bid for “philadelphia mountain bike” is \$0.05
 - balancing revenue with ad clutter
- Exact match vs. broad match
 - “philadelphia mountain bike” vs. “mountain bike” vs. “bike” vs. “philadelphia”
- Permit advertisers to condition bid on other information about user
 - e.g. geotargeting using user location
- Running a sponsored search advertising campaign is complex
 - all these decisions for a large portfolio of search phrases
- Associated industries/services:
 - Search Engine Optimization (SEO): improve organic ranking
 - e.g. optimize landing page, improve PageRank
 - Search Engine Marketing (SEM): improved sponsored ranking
 - e.g. optimize phrases, bids, quality score




Where's the Network?

- Market is a two-sided network:
 - users and their various interests determine which advertisers they will click on
 - advertisers and their products/services determine which users they want to reach
 - bipartite network with overlapping neighbor sets
 - cosmetically similar to our networked trading model
- Rich Get Richer aspects of two-sided markets:
 - advertisers most want to be on that search engine with the most users
 - users want to be on that search engine with the best search results
 - the more advertisers and users a search engine has, the more data
 - better estimates of advertiser quality, CTRs, good results for rare queries
- The “long tail of search”



Case Study: FCC Incentive Auction

- Problem: Repurpose broadcast TV spectrum for mobile communications
- “Reverse” auction: pay (some) broadcasters to go off the air
- “Forward” auction: mobile carriers purchase vacated spectrum
- Closing condition: forward revenues must cover reverse expenditures
- Many conceptual and technical challenges:
 - “repacking” constraints on remaining broadcasters: network of forbidden adjacencies
 - computing set of repackable broadcasters with highest bids is intractable
 - must keep auction rules as simple as possible for broadcasters
 - some carriers want national footprint → exposure problems



Summary

- Internet: distributed, self-interested behavior; competing incentives
- Leads to economic/game-theoretic situations:
 - routing, sponsored search, Quality of Service, spam, peer-to-peer systems
- Can seek **economic** as well as technological solutions:
 - auction rules in sponsored search; pricing schemes for QoS, spam, etc.
 - payments could be real or virtual
- Sometimes the game-theoretic behavior may not be an issue
 - Price of Anarchy for routing

