CS 345
Data Mining

Online algorithms

Search advertising

Online algorithms

- Classic model of algorithms
  - You get to see the entire input, then compute some function of it
  - In this context, “offline algorithm”
- Online algorithm
  - You get to see the input one piece at a time, and need to make irrevocable decisions along the way
- How is this different from the data stream model?

Example: Bipartite matching

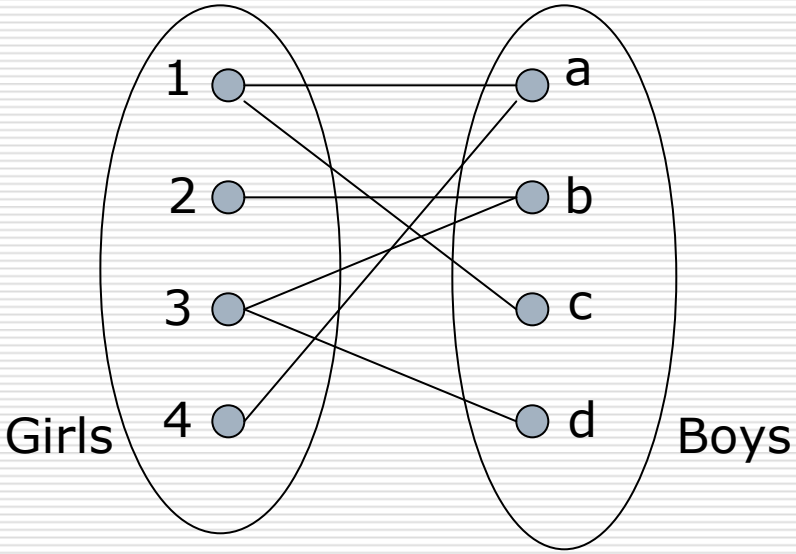

Example: Bipartite matching

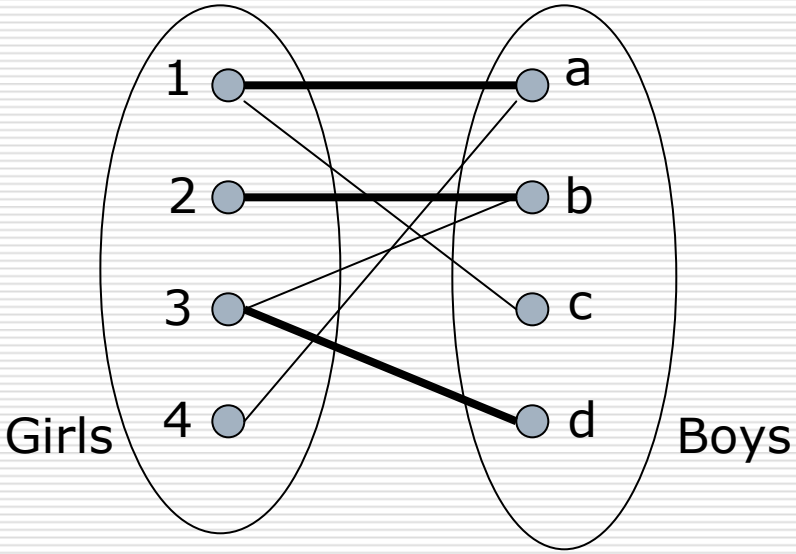

M = {(1,a),(2,b),(3,d)} is a matching
Cardinality of matching = |M| = 3

Example: Bipartite matching

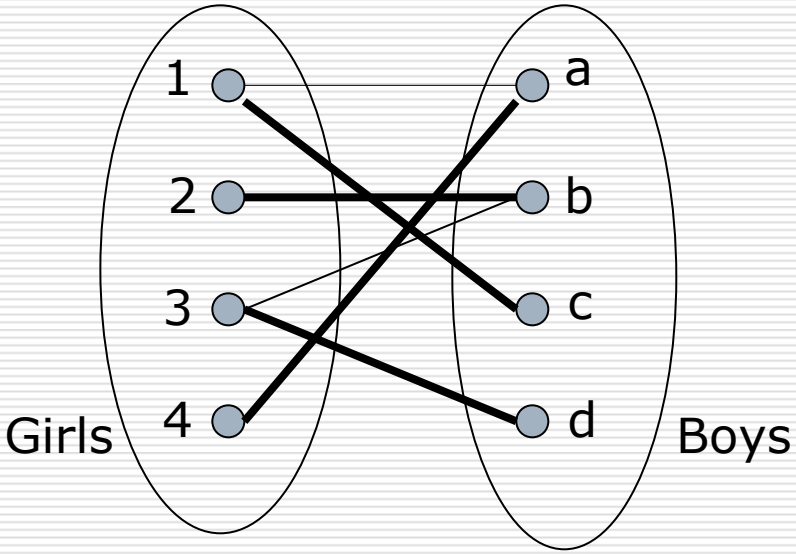

M = {(1,c),(2,b),(3,d),(4,a)} is a **perfect matching**

Matching Algorithm

- Problem: Find a maximum-cardinality matching
  - A perfect one if it exists
- There is a polynomial-time offline algorithm (Hopcroft and Karp 1973)
- But what if we don’t have the entire graph upfront?

Online problem

- Initially, we are given the set Boys
- In each round, one girl's choices are revealed
- At that time, we have to decide to either:
  - Pair the girl with a boy
  - Don't pair the girl with any boy
- Example of application: assigning tasks to servers

Online problem

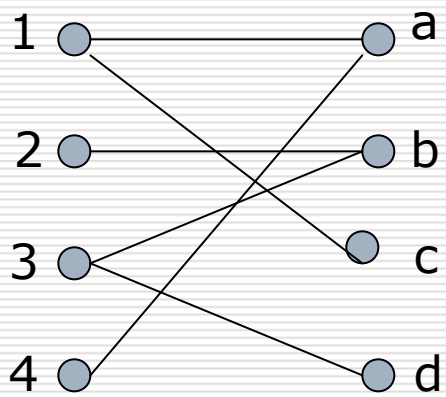

(1,a)

(2,b)

(3,d)

Greedy algorithm

- Pair the new girl with any eligible boy
  - If there is none, don’t pair girl
- How good is the algorithm?

Competitive Ratio

- For input I, suppose greedy produces matching M_{greedy} while an optimal matching is M_{opt}

Competitive ratio =

$$\min_{\text{all possible inputs I}} (|M_{greedy}|/|M_{opt}|)$$

Analyzing the greedy algorithm

- Consider the set G of girls matched in M_{opt} but not in M_{greedy}
- Then it must be the case that every boy adjacent to girls in G is already matched in M_{greedy}
- There must be at least |G| such boys
  - Otherwise the optimal algorithm could not have matched all the G girls
- Therefore

$|M_{greedy}| \geq |G| = |M_{opt} - M_{greedy}|$

$|M_{greedy}|/|M_{opt}| \geq 1/2$

Worst-case scenario

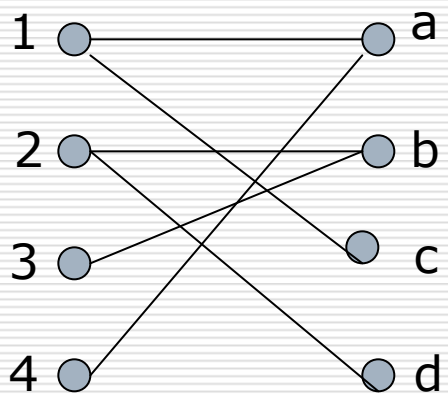

(1,a)

(2,b)

History of web advertising

- Banner ads (1995-2001)
  - Initial form of web advertising
  - Popular websites charged X$ for every 1000 “impressions” of ad
    - Called “CPM” rate
    - Modeled similar to TV, magazine ads
  - Untargeted to demographically tageted
  - Low clickthrough rates
    - low ROI for advertisers

Performance-based advertising

- Introduced by Overture around 2000
  - Advertisers "bid" on search keywords
  - When someone searches for that keyword, the highest bidder's ad is shown
  - Advertiser is charged only if the ad is clicked on
- Similar model adopted by Google with some changes around 2002
  - Called "Adwords"

Ads vs. search results

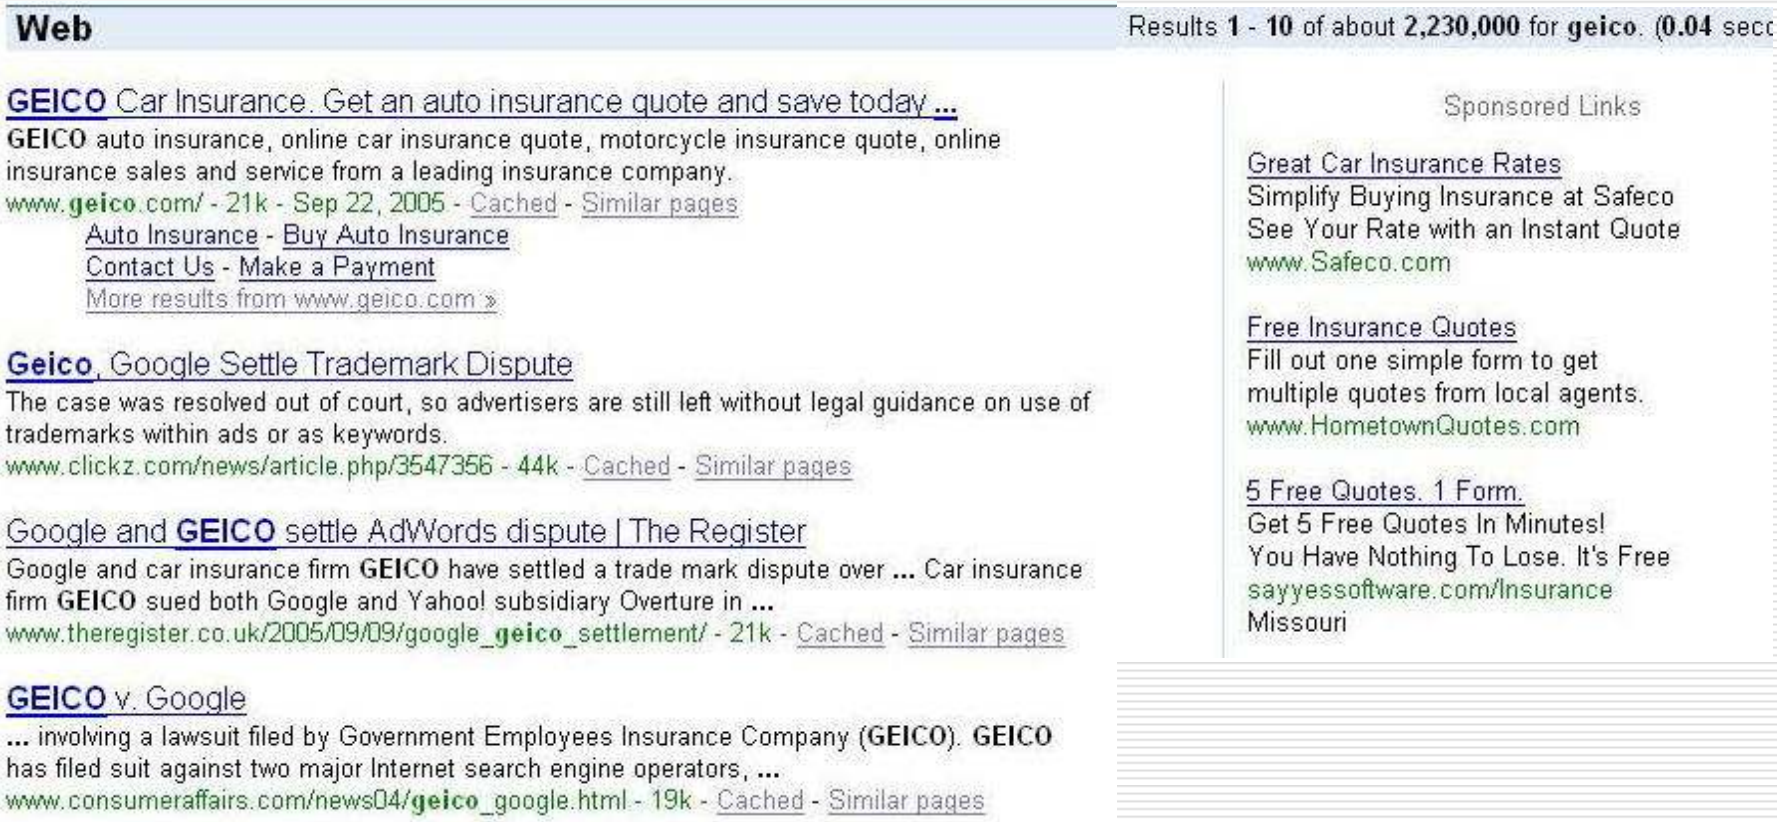

Web Results **1** - **10** of about **2,230,000** for **geico**. (**0.04** secc

GEICO Car Insurance. Get an auto insurance quote and save today ...
GEICO auto insurance, online car insurance quote, motorcycle insurance quote, online insurance sales and service from a leading insurance company.
www.**geico**.com/ - 21k - Sep 22, 2005 - Cached - Similar pages
Auto Insurance - Buy Auto Insurance
Contact Us - Make a Payment
More results from www.geico.com »

Geico, Google Settle Trademark Dispute
The case was resolved out of court, so advertisers are still left without legal guidance on use of trademarks within ads or as keywords.
www.clickz.com/news/article.php/3547356 - 44k - Cached - Similar pages

Google and **GEICO** settle AdWords dispute | The Register
Google and car insurance firm **GEICO** have settled a trade mark dispute over ... Car insurance firm **GEICO** sued both Google and Yahoo! subsidiary Overture in ...
www.theregister.co.uk/2005/09/09/google_**geico**_settlement/ - 21k - Cached - Similar pages

GEICO v. Google
... involving a lawsuit filed by Government Employees Insurance Company (**GEICO**). **GEICO** has filed suit against two major Internet search engine operators, ...
www.consumeraffairs.com/news04/**geico**_google.html - 19k - Cached - Similar pages

Sponsored Links

Great Car Insurance Rates
Simplify Buying Insurance at Safeco
See Your Rate with an Instant Quote
www.Safeco.com

Free Insurance Quotes
Fill out one simple form to get multiple quotes from local agents.
www.HometownQuotes.com

5 Free Quotes. 1 Form.
Get 5 Free Quotes In Minutes!
You Have Nothing To Lose. It's Free
sayyessoftware.com/Insurance
Missouri

Web 2.0

- Search advertising is the revenue model
  - Multi-billion-dollar industry
  - Advertisers pay for clicks on their ads
- Interesting problems
  - What ads to show for a search?
  - If I'm an advertiser, which search terms should I bid on and how much to bid?

Adwords problem

- A stream of queries arrives at the search engine
  - q1, q2,…
- Several advertisers bid on each query
- When query q_i arrives, search engine must pick a subset of advertisers whose ads are shown
- Goal: maximize search engine's revenues
- Clearly we need an online algorithm!

Greedy algorithm

- Simplest algorithm is greedy
- It’s easy to see that the greedy algorithm is actually optimal!

Complications (1)

- Each ad has a different likelihood of being clicked
  - Advertiser 1 bids $2, click probability = 0.1
  - Advertiser 2 bids $1, click probability = 0.5
  - Clickthrough rate measured historically
- Simple solution
  - Instead of raw bids, use the "expected revenue per click"

Complications (2)

- Each advertiser has a limited budget
  - Search engine guarantees that the advertiser will not be charged more than their daily budget

Simplified model

- Assume all bids are 0 or 1
- Each advertiser has the same budget B
- Let's try the greedy algorithm
  - Arbitrarily pick an eligible advertiser for each keyword

Bad scenario for greedy

- Two advertisers A and B
- A bids on query x, B bids on x and y
- Both have budgets of $4
- Query stream: xxxxyyyy
  - Worst case greedy choice: BBBB____
  - Optimal: AAAABBBB
  - Competitive ratio = ½
- Simple analysis shows this is the worst case

BALANCE algorithm [MSVV]

- [Mehta, Saberi, Vazirani, and Vazirani]
- For each query, pick the advertiser with the largest unspent budget
  - Break ties arbitrarily

Example: BALANCE

- Two advertisers A and B
- A bids on query x, B bids on x and y
- Both have budgets of $4
- Query stream: xxxxyyyy
- BALANCE choice: ABABBB__
  - Optimal: AAAABBBB
- Competitive ratio = ¾

Analyzing BALANCE

- Consider simple case: two advertisers, A_1 and A_2, each with budget B (assume B À 1)
- Assume optimal solution exhausts both advertisers' budgets

Analyzing Balance

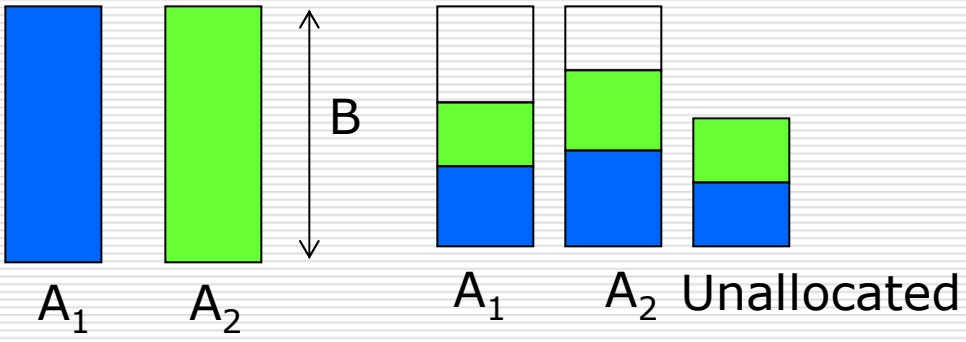

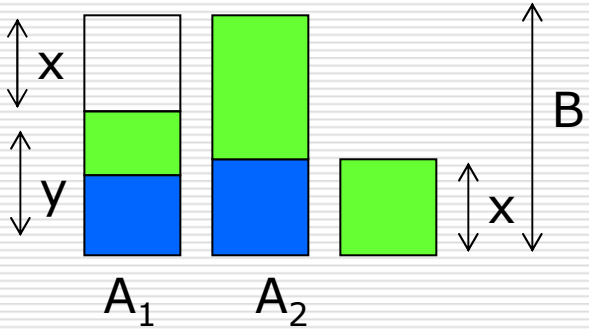

Opt revenue = 2B
Balance revenue = 2B-x = B+y

We have y ¸ x
Balance revenue is minimum for x=y=B/2
Minimum Balance revenue = 3B/2
Competitive Ratio = 3/4

General Result

- In the general case, worst competitive ratio of BALANCE is $1-1/e$ = approx. 0.63
- Interestingly, no online algorithm has a better competitive ratio
- Won't go through the details here, but let's see the worst case that gives this ratio

Worst case for BALANCE

- N advertisers, each with budget B À N À 1
- NB queries appear in N rounds of B queries each
- Round 1 queries: bidders A_1, A_2, …, A_N
- Round 2 queries: bidders A_2, A_3, …, A_N
- Round i queries: bidders A_i, …, A_N
- Optimum allocation: allocate round i queries to A_i
  - Optimum revenue NB

BALANCE allocation

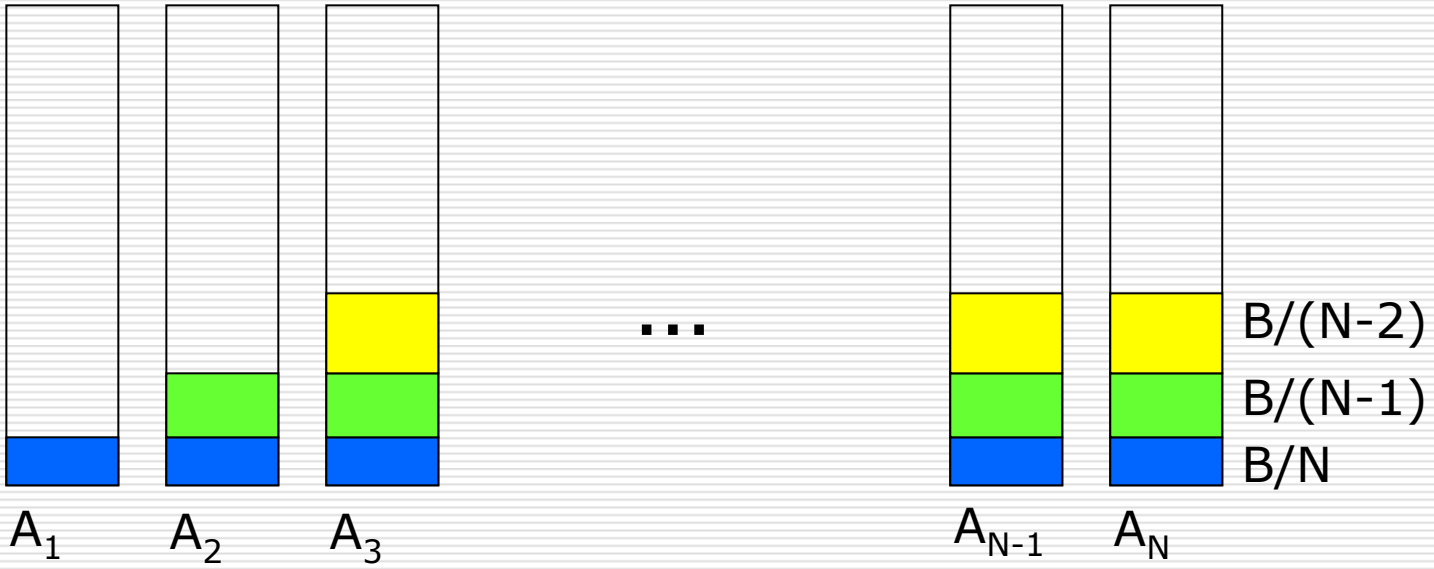

The sum of the allocations to a bin k is given by:
$S_k = \min(B, \sum_{1 \cdot 1 \cdot k} B/(N-i+1))$

BALANCE analysis

B/1 B/2 B/3 B/4 … B/k … B/(N-1) B/N

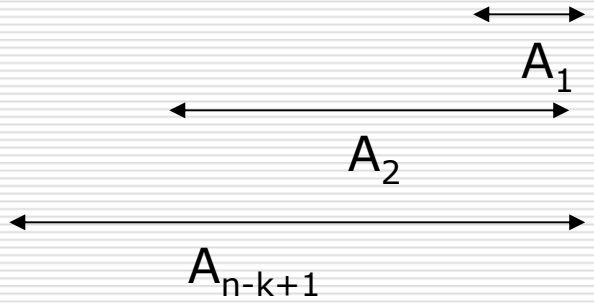

BALANCE analysis

- Fact: $H_n = \sum_{1 \cdot i \cdot n} 1/i$ = approx. log(n) for large n
  - Result due to Euler
- So if H_k = log(N)-1, k=N/e

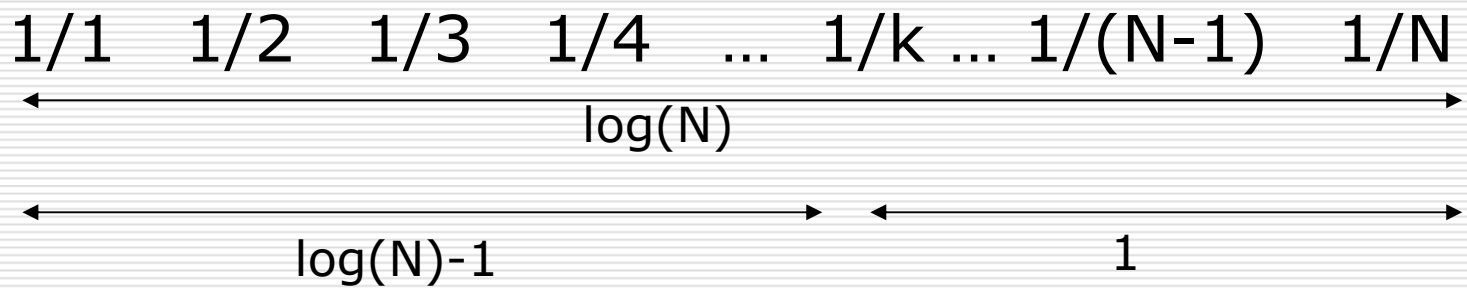

BALANCE analysis

- So after the first N(1-1/e) rounds, we cannot allocate a query to any advertiser
- Revenue = BN(1-1/e)
- Competitive ratio = 1-1/e

General version of problem

- MSVV also provides an algorithm for the general case with arbitrary bids
  - Same competitive ratio

Sidebar: What's in a name?

- Geico sued Google, contending that it owned the trademark "Geico"
  - Thus, ads for the keyword geico couldn't be sold to others
- Court Ruling: search engines can sell keywords including trademarks
- No court ruling yet: whether the ad itself can use the trademarked word(s)