The Future of Information Discovery: Content Optimization, Social Networks, Interactivity---It's All Very Cloudy

Raghu Ramakrishnan

Chief Scientist, Search & Cloud Computing

Yahoo! Fellow

Main Points

- The medium is being massaged
  - i.e., the canvas is no longer passive
  - It is also very personal, and `always on'
  - *(With apologies to Marshall McLuhan)*
- Search and browse converging
  - Semantics is the new frontier
  - Social networks the new distribution channel
- Learning from Big Data is a key
  - Value of data is now recognized
  - Compute platforms will be increasingly data-centric

The Medium is the Message

McLuhan, 1964: Understanding Media: The Extensions of Man

Wikipedia:

- “People tend to focus on the obvious, which is the content, to provide us information, but in the process, we largely miss the structural changes in our affairs that are introduced subtly, or over long periods of time.”
- McLuhan frequently punned on the word "message" changing it to "mass age", "mess age", and "massage"; a later book, *The Medium is the Massage* was originally to be titled *The Medium is the Message*.

Yahoo! Data Scale

- **Massive user base and engagement**
  - 650M+ unique users, 11B page visits/month
  - Hundreds of petabytes of storage
  - Hundreds of billions of objects
  - Hundreds of thousands of requests/sec, 200B events/day, 200 PB/day
- **Global**
  - Tens of globally distributed data centers
  - Serving each region at low latencies
- **Challenging Users**
  - Rapidly extracting value from voluminous data
  - Downtime is not an option (outages cost $millions)
  - Variable usage patterns

Content Optimization

Recommended links

+79% clicks
vs. randomly selected

News Interests

+250% clicks
vs. one size fits all

Top Searches

+43% clicks
vs. editor selected

Key Features

Package Ranker (CORE)

Ranks packages by expected CTR based on data collected every 5 minutes

Dashboard (CORE)

Provides real-time insights into performance by package, segment, and property

Mix Management (Property)

Ensures editorial voice is maintained and user gets a variety of content

Package rotation (Property)

Tracks which stories a user has seen and rotates them after user has seen them for a certain period of time

Key Performance Indicators

Lifts in quantitative metrics

Editorial Voice Preserved

Cloud Structured Storage @ Yahoo!

By the Numbers

- 1.5+M read rps
- 4 Key Platforms
- 10,000+ servers
- 10 Geo Zones
- 102B emails/month
- 13+B Ad serves/day
- 11B visits/month

- ~2B User Ids
- ~700M Uniq Users*
- 175+M Users in US
- 285+M Mail users
- 40+ Countries

All Data from July/Aug. Worldwide unless indicated to the contrary
* Yahoo!-branded sites

Hadoop: Stability at Scale

Hadoop powers the Yahoo! Network: must be rock-solid

We fix bugs before you see them

- We run very large clusters
- We have a large QA effort
- We run a huge variety of workloads
- Recent spinoff (HortonWorks)

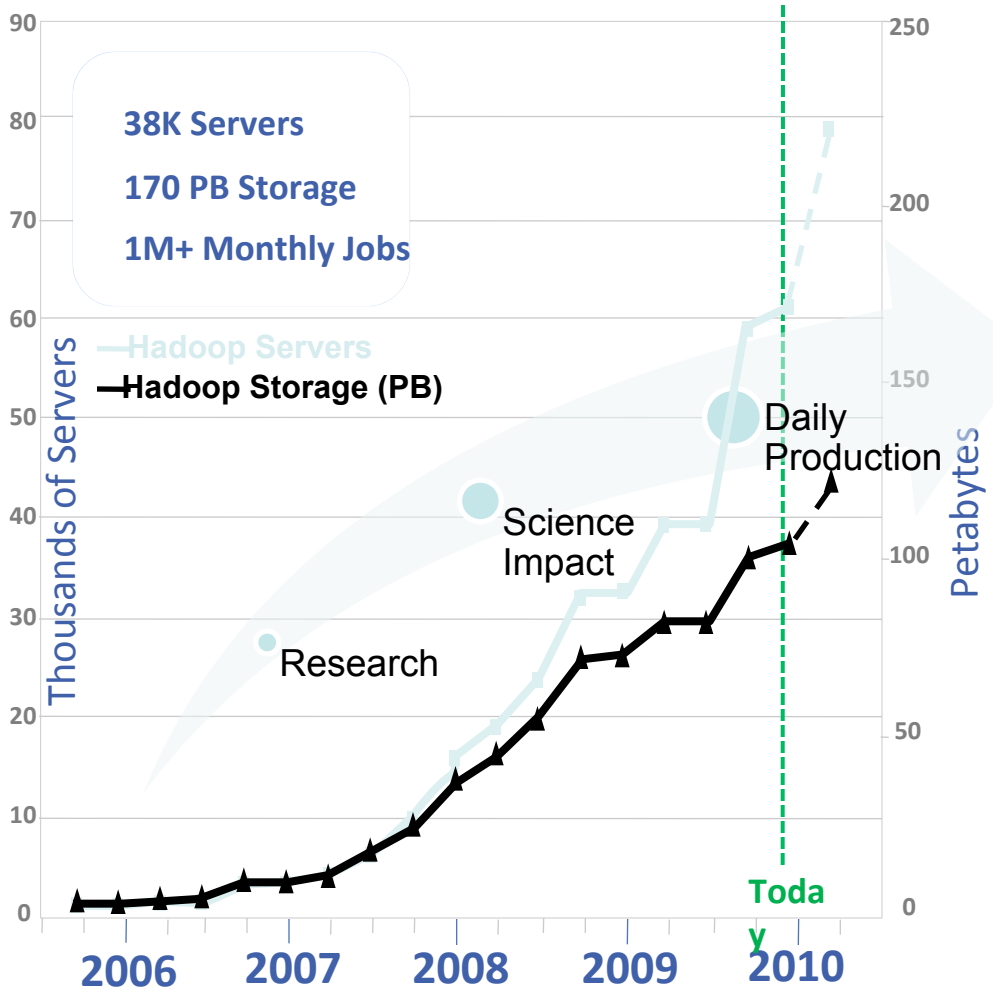

What is PNUTS/Sherpa?

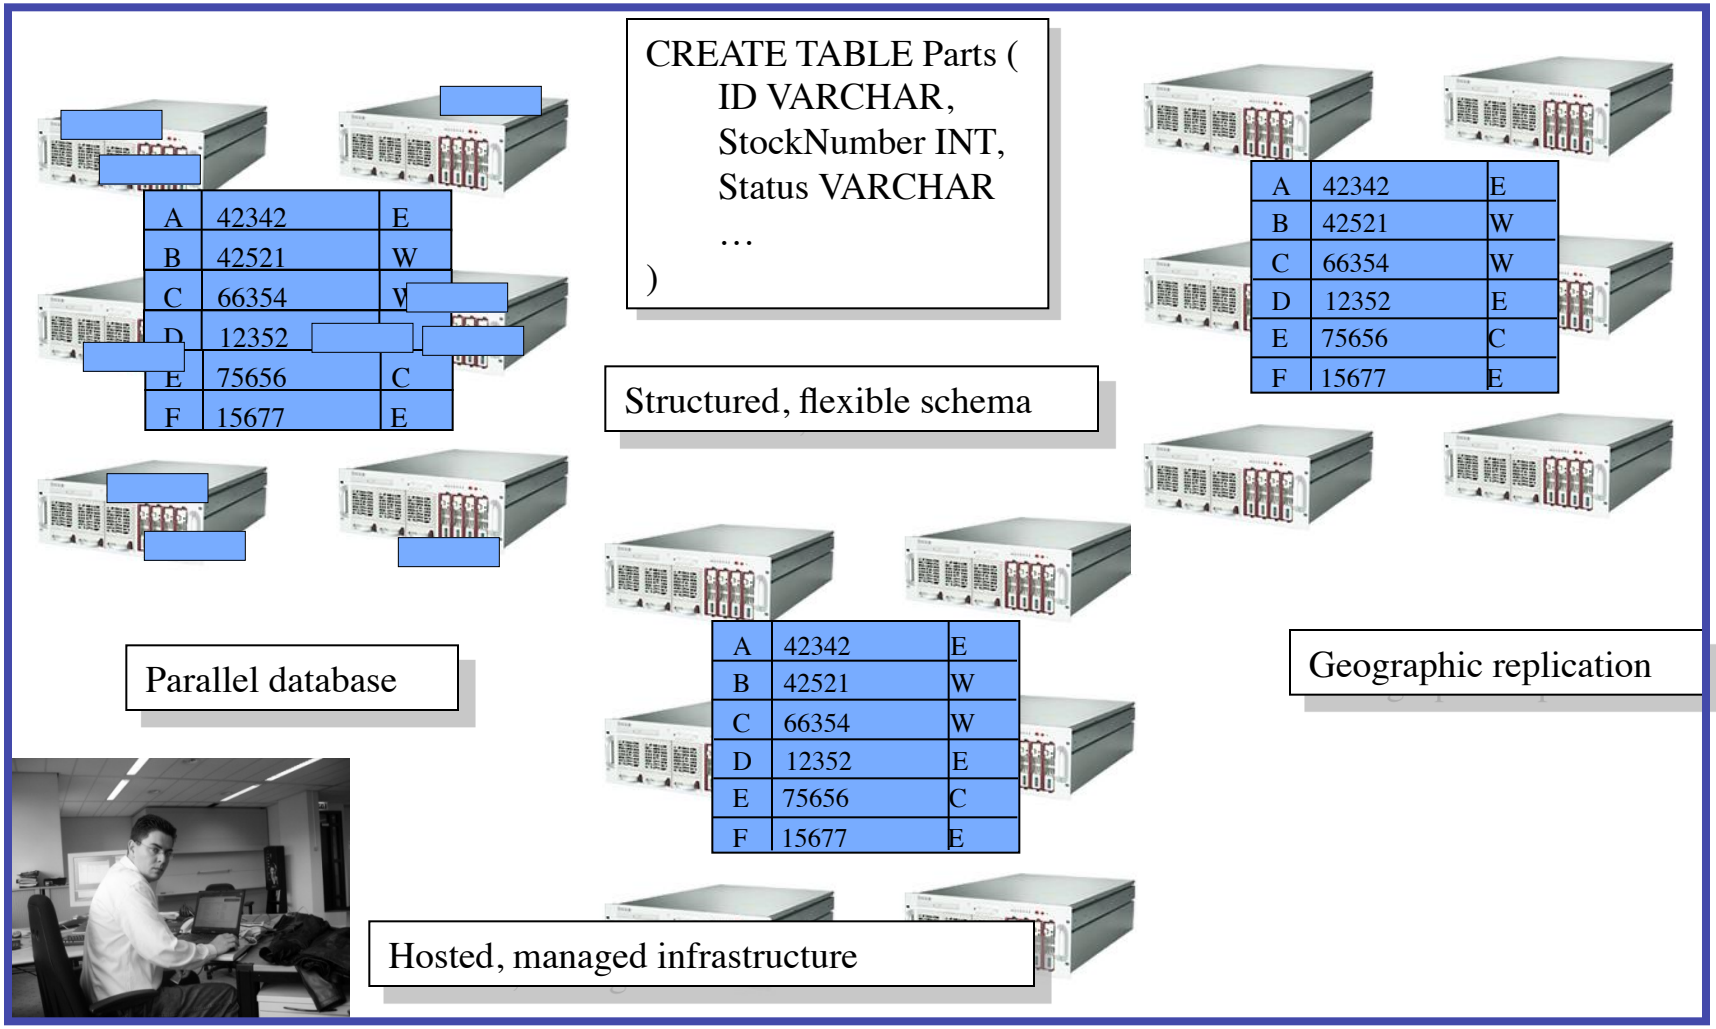

Sherpa Adoption

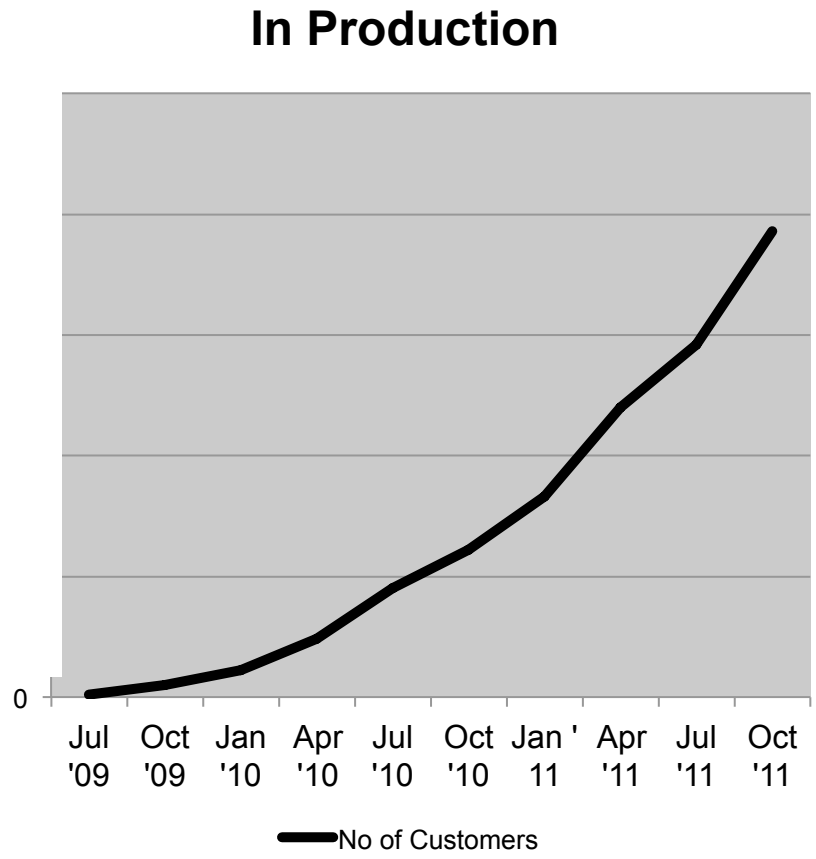

- Abuse Platform
- Answers
- TW Auctions
- COKE
- ConnectedTV
- Groups
- Global Blog Search
- LocDrop
- Mobile Platform
- Mobile Advertising
- Consumer Platform
- Digu
- Search History
- Pipes
- Pulse
- Sports
- Shopping
- UGC Cloud
- YOS Social Directory
- YAP
- YQL
- Aqua (Ad Quality)
- Image Search
- Hosted Search
- And others

QoQ 60% growth in projects using Sherpa

10

Sherpa Footprint Growth Curve

Storage Nodes

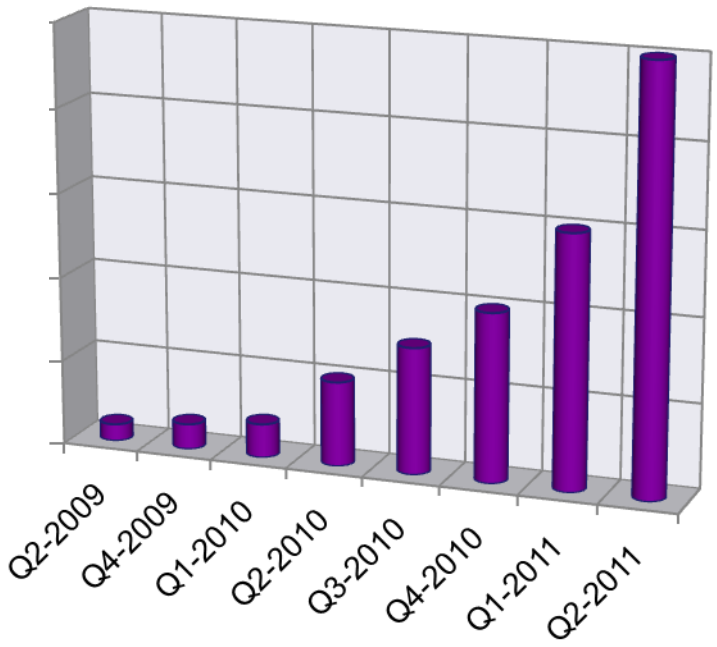

Replicas

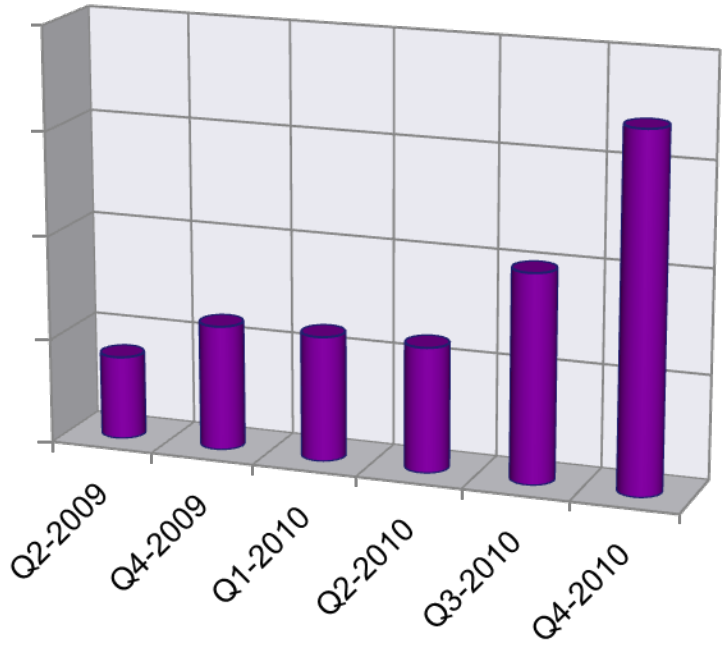

10s of data centers, 1000s of machines, PBs of data stored, 100Bs of records, 100Ks of requests/sec

Selective geo-replication, secondary indexes, hashed and ordered tables, flex schemas, eventual and timeline consistency

CORE:
Content Optimization

Recommender Approaches

Estimate Most Popular (EMP)

"What's most engaging overall?"

Behavioral Affinities

"People who did X, did Y"

Attribute Similarities

"Related items with similar metadata"

Social Recommendations

"What are my trusted connections into?"

Business Optimization

"What generates most business value?"

Personalized Recommendations

"What's most relevant to me based on my interests, attributes and relationships?"

Explore/Exploit

- What is the best strategy for new articles?
  - If we show it and it's bad: lose clicks
  - If we delay and it's good: lose clicks
- Solution: Show it while we don't have much data if it looks promising
  - Classical **multi-armed bandit** type problem
  - Our setup is different than the ones studied in the literature; new ML problem

(Online Models for Content Optimization, NIPS 2008: D. Agarwal, B. Chen, P. Elango, N. Motgi, S. Park, R. Ramakrishnan, S. Roy, J. Zachariah)

CTR Curves for Two Days

Each curve is the CTR of an item in the Today Module over time

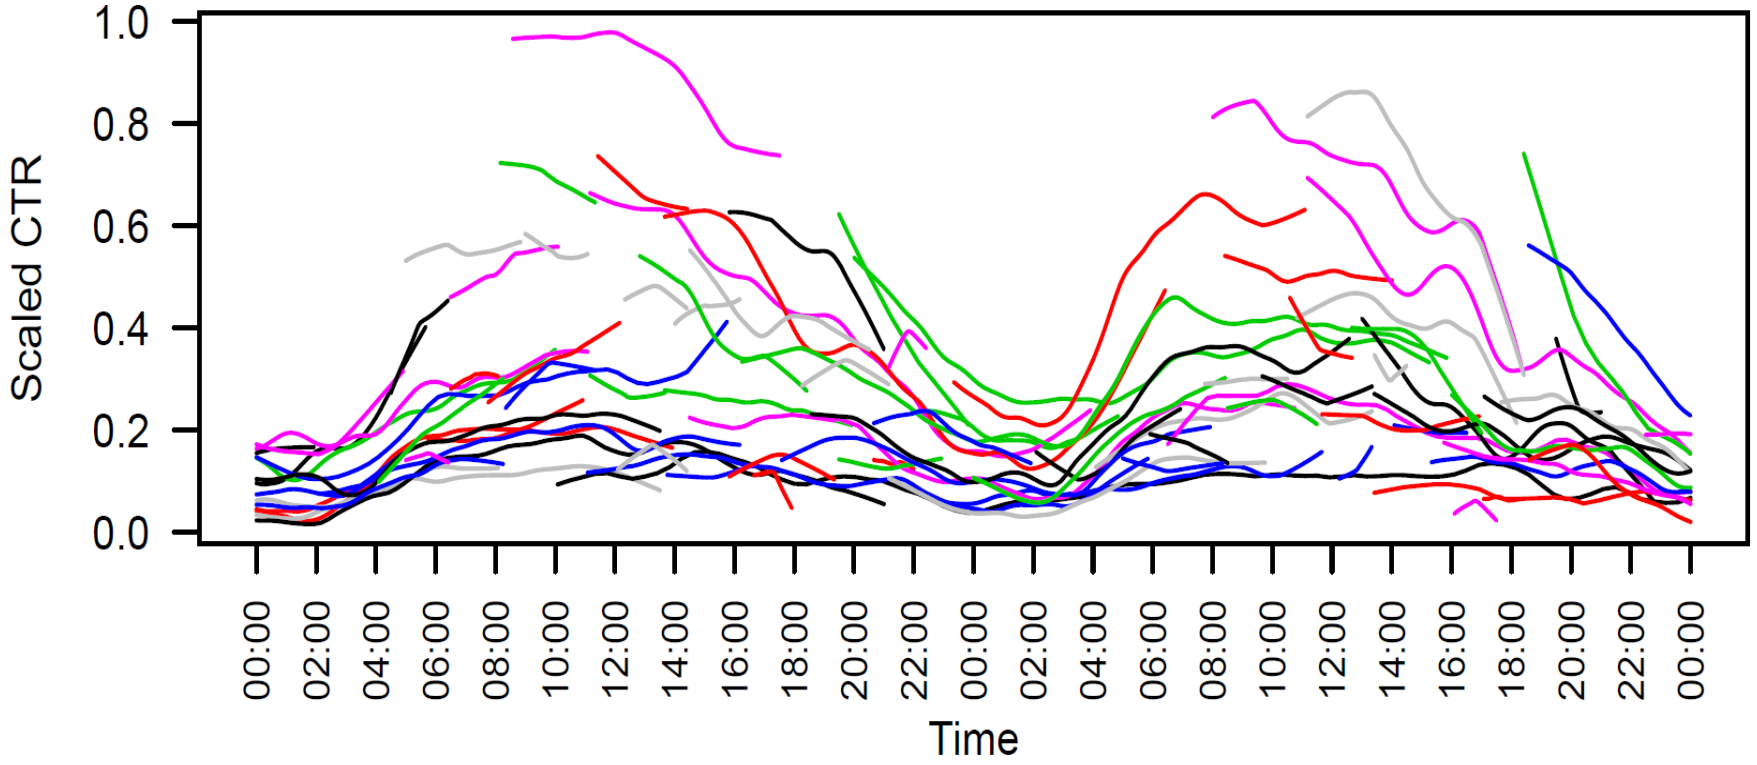

Traffic obtained from a controlled randomized experiment

Things to note:

(a) Short lifetimes, (b) temporal effects, (c) often breaking news stories

Degrees of Personalization

Most Popular

Most engaging overall based on objective metrics

Most Popular + Per User History

Engaging overall, and aware of what I've already seen

Light Personalization

More relevant to me based on my age, gender and property usage

Deep Personalization

Most relevant to me based on my deep interests and relationships

---

Voice and Business Rules

Real-time Dashboard

Business Optimization

CORE Modeling Overview

Online Learning
- Online regression models, time-series models
- Model the temporal dynamics
- Provide fast learning for per-item models

Offline Modeling
- Exploratory data analysis
- Regression, feature selection, collaborative filtering (factorization)
- Seed online models & explore/exploit methods at good initial points
- Reduce the set of candidate items

Near real-time user feedback

Explore/Exploit
- Multi-armed bandits
- Find the best way of collecting real-time user feedback (for new items)

Large amount of historical data (user event streams)

Hadoop pipelines

PNUTS

Hadoop

CORE: Scoring + Modeling Framework

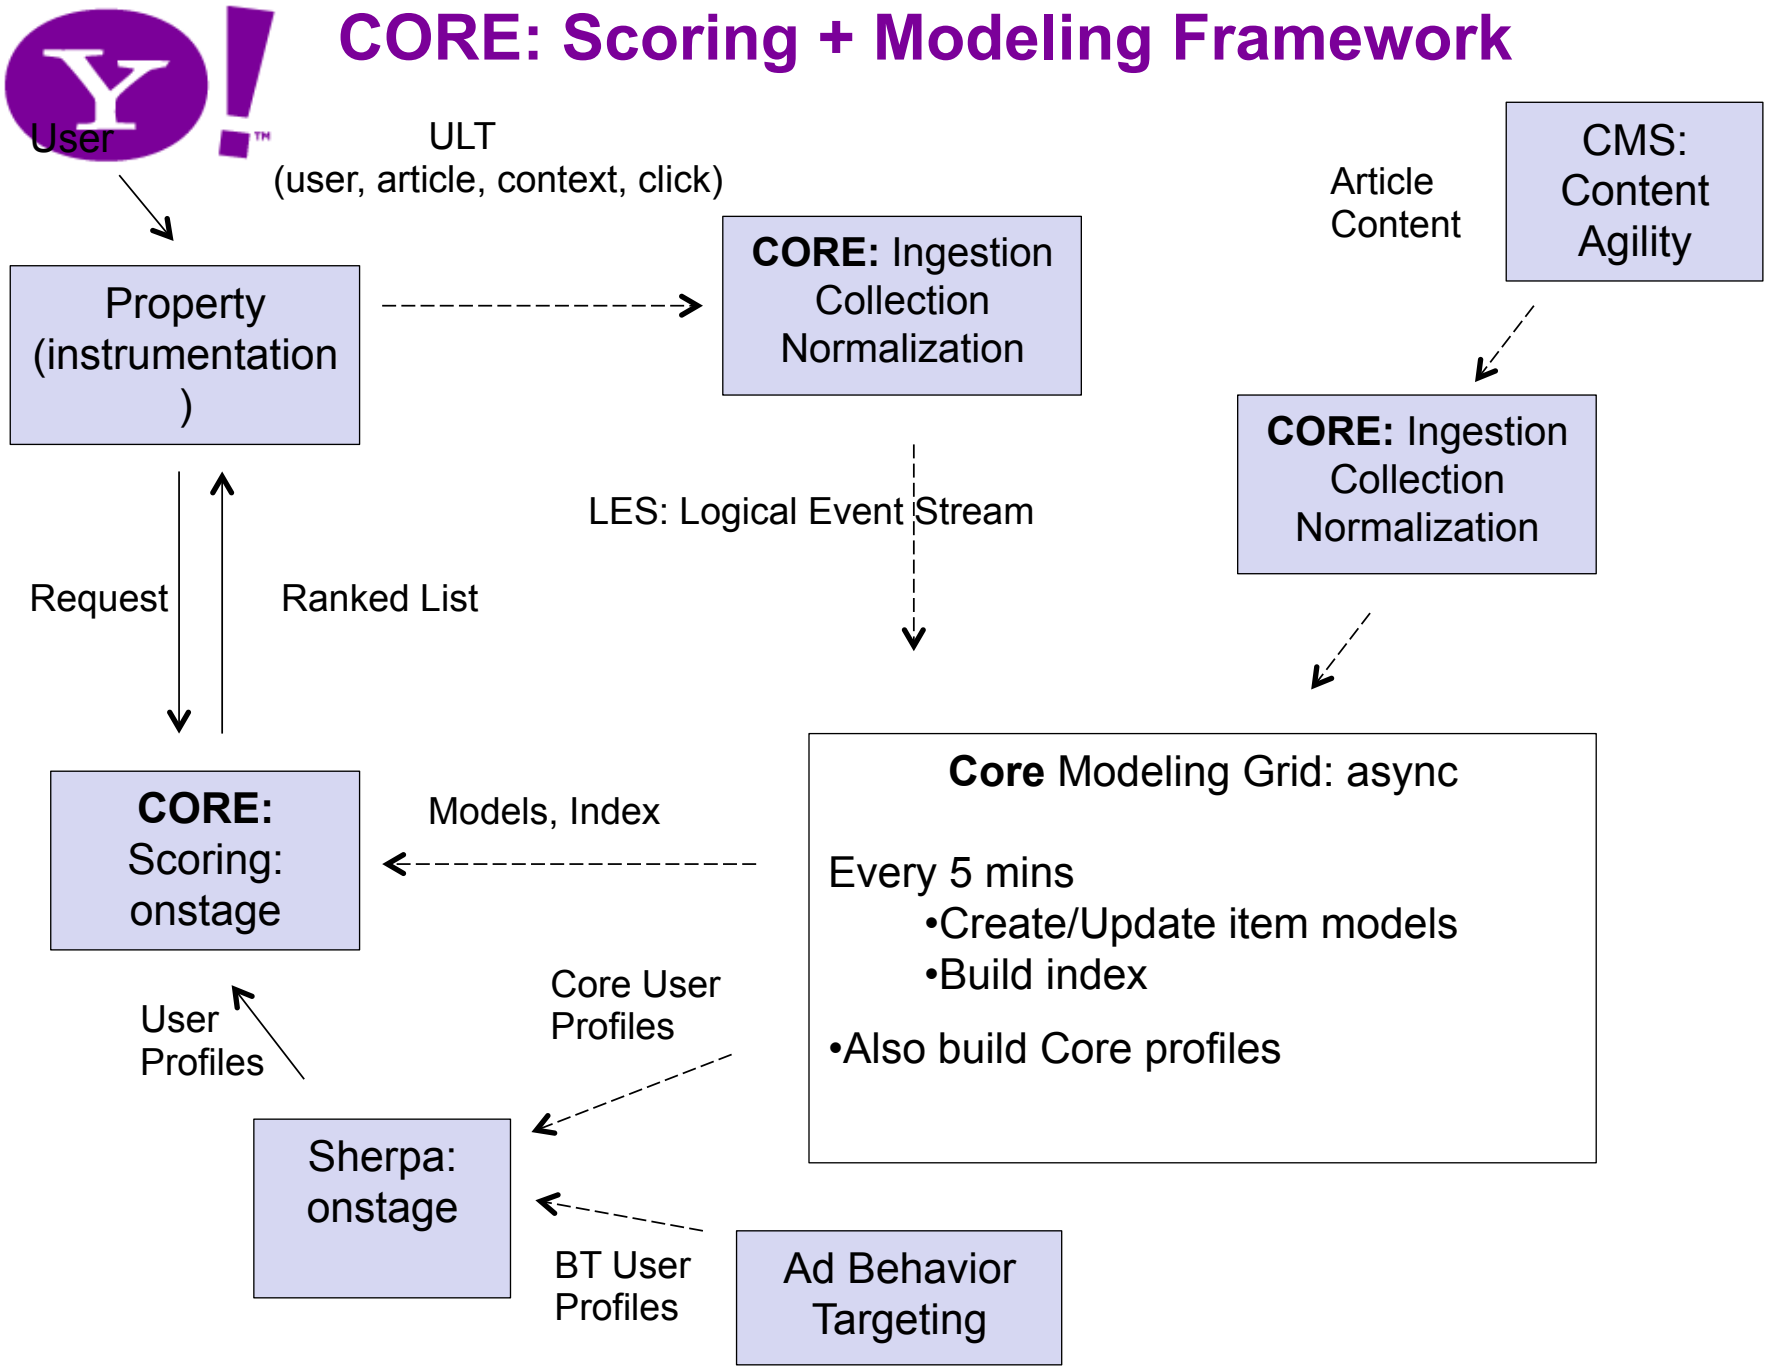

CORE Dashboard: Overall CTR

Compare performance of models and historical benchmarks

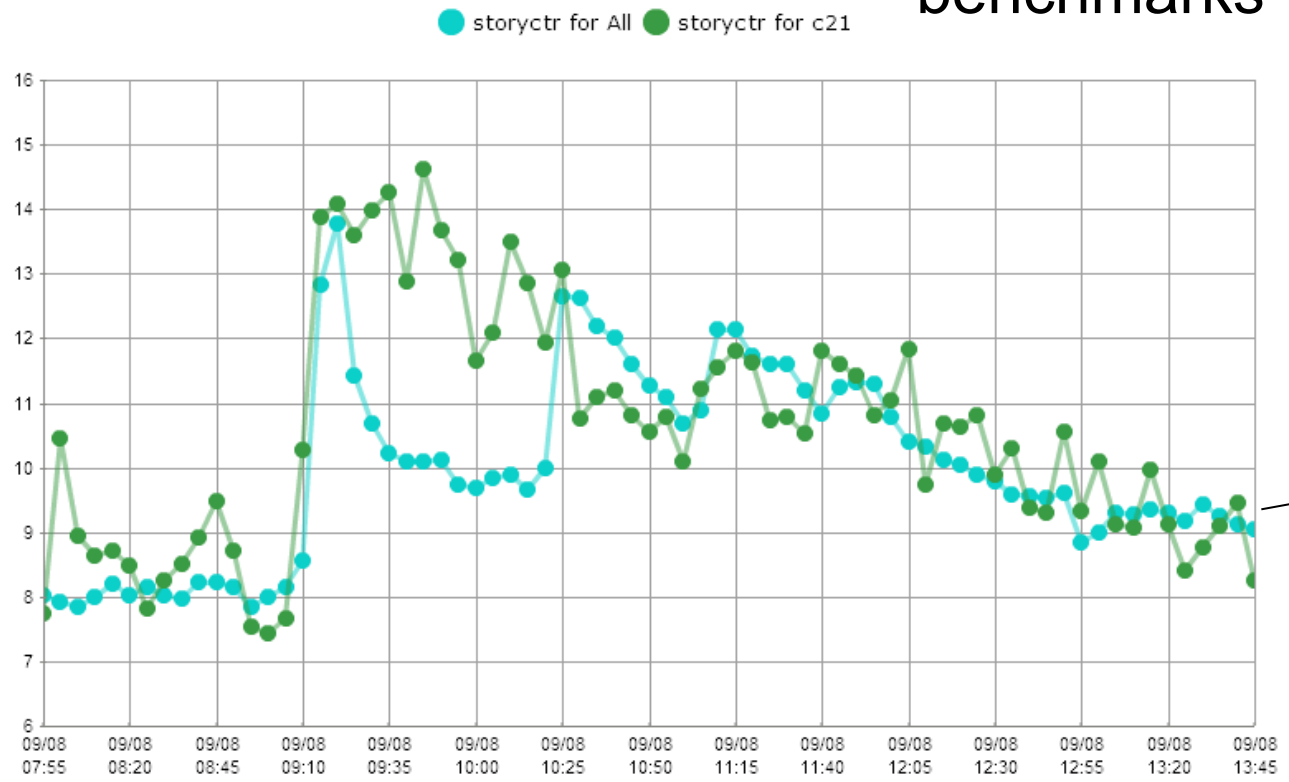

Compare buckets and models over time

c21

All

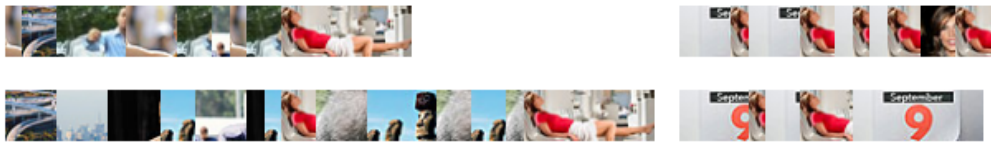

See which content was promoted most across time

Bucket	Page Views,	All Clicks	Story Clicks	Footer Clicks	OverAll CTR	Story CTR	Footer CTR	Lift compared to All
All	43,005,783	13,061,688	4,286,407	8,775,281	30.37	9.97	20.4	0
c21	281,821	84,853	29,777	55,076	30.11	10.57	19.54	6.02

Compare bucket metrics

CORE Dashboard: Segment Heat Map

Package	male	female	OMG	BUAuto	BUEnt	BU Fin	Health	BUSport▲	NBA	BUTrav	ALL
	408,260 18,440 0.0452 8.477	390,404 14,449 0.037 -11.113	270,039 16,940 0.0627 50.661	121,080 7,389 0.061 46.564	270,038 16,940 0.0627 50.661	325,873 20,012 0.0614 47.488	195,796 12,763 0.0652 56.553	350,152 21,454 0.0613 47.152	132,916 9,457 0.0712 70.879	123,388 7,896 0.064 53.691	923,611 38,457 0.0416 0
	1 8,067 852 0.1056 153.654	1 7,657 674 0.088 111.405	1 5,125 720 0.1405 237.406	1 2,382 286 0.1201 188.362	1 5,125 720 0.1405 237.406	1 6,415 858 0.1337 221.221	1 3,769 532 0.1412 239	1 6,750 917 0.1359 226.272	1 2,585 385 0.1489 257.696	1 2,490 330 0.1325 218.294	1 18,131 1,738 0.0958 130.143
	5 9,968 644 0.0646 55.164	3 12,847 777 0.0605 45.256	2 8,569 885 0.1033 148.043	4 3,529 326 0.0924 121.86	2 8,569 885 0.1033 148.043	3 9,744 922 0.0946 127.252	3 6,067 643 0.106 154.537	2 10,187 1,004 0.0986 136.702	5 3,820 420 0.1099 164.058	2 4,037 433 0.1073 157.598	4 25,744 1,595 0.062 48.798
	2 3,326 249 0.0749 79.8	5 3,954 212 0.0536 28.769	5 2,521 231 0.0916 120.066	2 1,004 102 0.1016 143.995	5 2,521 231 0.0916 120.066	5 3,016 276 0.0915 119.782	5 1,860 186 0.1 140.167	3 3,291 310 0.0942 126.229	3 1,141 136 0.1192 186.264	3 1,039 100 0.0962 131.152	3 8,500 541 0.0636 52.859
	11 2,562 133 0.0519 24.677	13 2,004 81 0.0404 -2.926	3 1,250 122 0.0976 134.403	6 629 51 0.0811 94.73	3 1,250 122 0.0976 134.403	4 1,608 151 0.0939 125.53	2 919 103 0.1121 169.175	4 1,669 154 0.0923 121.604	4 655 74 0.113 171.334	4 591 55 0.0931 123.506	10 5,342 252 0.0472 13.295
	3 2,881 206 0.0715 71.727	2 3,242 230 0.0709 70.384	4 2,071 196 0.0946 127.295	3 949 95 0.1001 140.42	4 2,071 196 0.0946 127.295	2 2,614 254 0.0972 133.368	4 1,605 165 0.1028 146.901	5 2,740 239 0.0872 109.489	10 1,036 94 0.0907 117.912	9 958 78 0.0814 95.543	2 7,043 493 0.07 68.114
	6 10,785 649 0.0602 44.523	4 12,768 742 0.0581 39.571	7 8,580 694 0.0809 94.261	7 3,511 283 0.0806 93.584	7 8,580 694 0.0809 94.261	6 9,725 795 0.0817 96.332	6 6,138 550 0.0896 115.204	6 10,670 866 0.0812 94.525	11 3,669 321 0.0875 110.122	5 3,785 339 0.0896 115.104	5 27,331 1,641 0.06 44.2
	10 22,202 1,212 0.0546 31.106	7 23,328 1,200 0.0514 23.543	6 15,593 1,289 0.0827 98.535	5 6,552 533 0.0813 95.374	6 15,593 1,289 0.0827 98.535	7 17,652 1,376 0.078 87.214	8 10,797 915 0.0847 103.532	7 19,050 1,522 0.0799 91.882	9 6,639 604 0.091 118.498	7 6,435 552 0.0858 106.018	6 52,978 2,786 0.0526 26.299
	22 26,685 1,160 0.0435 4.401	10 35,405 1,530 0.0432 3.786	8 19,832 1,572 0.0793 90.371	9 7,844 552 0.0704 69.011	8 19,832 1,572 0.0793 90.371	8 21,743 1,641 0.0755 81.26	7 13,721 1,167 0.0851 104.267	8 22,168 1,743 0.0786 88.836	8 8,249 788 0.0955 129.424	8 8,327 689 0.0827 98.721	18 74,559 3,167 0.0425 2.014
	4 7,745 518 0.0669 60.628	26 7,202 185 0.0257 -38.306	13 4,898 322 0.0657 57.889	15 2,308 148 0.0641 54.007	13 4,898 322 0.0657 57.889	11 6,051 423 0.0699 67.891	19 3,652 235 0.0643 54.544	9 6,436 506 0.0786 88.82	2 2,562 308 0.1202 188.726	12 2,359 169 0.0716 72.057	7 17,235 834 0.0484 16.217
	7 7,699 460 0.0597 43.495	29 7,201 169 0.0235 -43.635	11 4,809 340 0.0707 69.8	10 2,269 158 0.0696 67.239	11 4,809 340 0.0707 69.8	9 6,004 433 0.0721 73.205	14 3,544 243 0.0686 64.674	10 6,247 475 0.076 82.615	6 2,482 257 0.1035 148.682	11 2,329 167 0.0717 72.211	12 17,169 783 0.0456 9.525
	7,688	7,229	4,785	2,280	4,785	6,037	3,501	6,319	2,397	2,312	17,275

Examples

- **ACQUISITION:** A “Star Trek” package was #3 with 18-20 demo, #2 with 21-24 demo, but #9 overall. We can acquire younger audiences with targeted content like this.

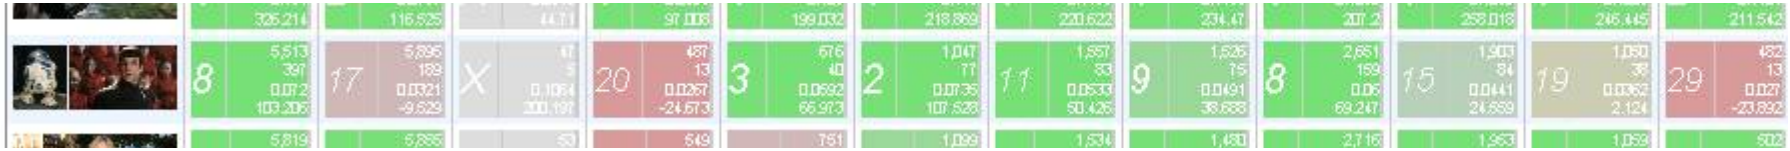

- **ENGAGEMENT:** “Kobe’s astonishing shot” was #25 with women, but #5 with men. We can better engage men (or sports fans) by showing more like this, women by showing less.

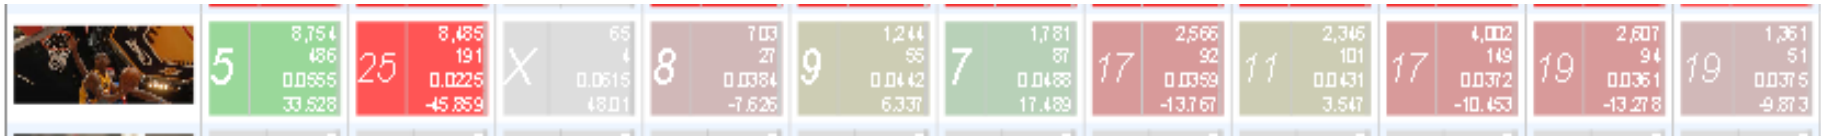

- **REACH:** A package about a hair-pulling soccer player was just plain interesting to everyone (#1-3). We can maintain reach by programming content for the mass audience.

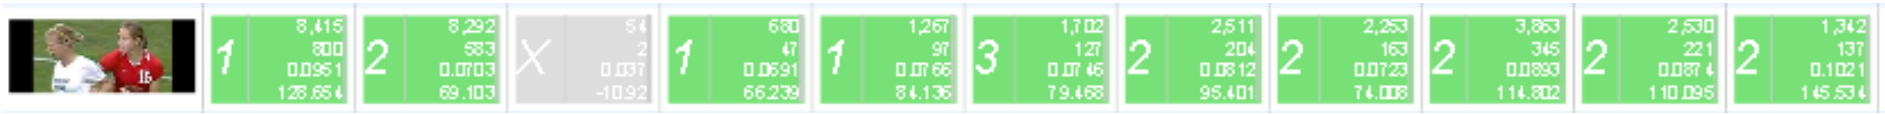

21

SEARCH TRENDS

ry Bookmarks Yahoo! Tools Help

http://search.yahoo.com/search;_ylt=A0oGklHEA_RLAyYBgxJXNyoA?p=eggplant+parmigiana+baltimore&fr2=sb-top&fr=moz35&sao=1

hummus

na baltimore - Yah...

YAHOO!

Web Images Video Local Shopping News More

eggplant parmigiana baltimore

Search

Options

Start typing to see suggestions.

Explore related concepts:

Pizza	mushrooms
veal	tomato sauce
Chicken	Little Italy
Italian Restaurant	Crab Cake

Settin

Search Pad

SearchScan - On

30,300 results for
eggplant parmigiana ...:

Show All

Los Angeles Times

Local Business Sites

Eggplant Parmigiana Restaurants near Baltimore, MD

Neighborhood
All (36) Abell (1) Canton (2) Central Bal... (4) Charles North (1) Chinquapin ... (1) Downtown (1) Federal Hill (1)

1 Ciao Bella - Baltimore ★★★★★ (11)
local.yahoo.com
(410) 685-7733 - 236 S High St, Baltimore, MD
Menu: eggplant parmigiana
4 Reviews | Overview | 2 Photos | Directions

2 Amicci's ★★★★★ (20)
amiccis.com
(410) 528-1096 - 231 S High St, Baltimore, MD
Menu: eggplant parmigiana
14 Reviews | Overview | 23 Photos | Directions

3 Pasticcio ★★★★☆ (8)
local.yahoo.com
(410) 522-7700 - 2400 Boston St, Baltimore, MD
Menu: eggplant parmigiana
5 Reviews | Overview | 3 Photos | Directions

4 Caesar's Den ★★★★☆ (7)
caesarsden.com
(410) 547-0820 - 223 S High St, Baltimore, MD
Menu: eggplant parmigiana
4 Reviews | Overview | 11 Photos | Directions

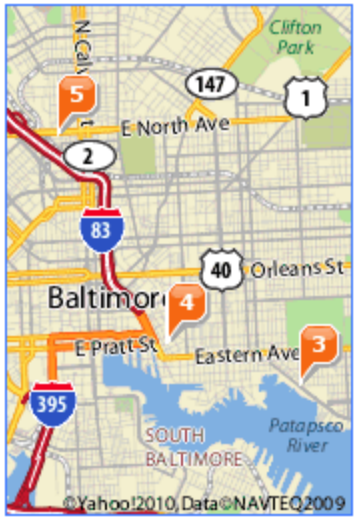

julia roberts - Yahoo! Search Results - Mozilla Firefox

File Edit View History Bookmarks Yahoo! Tools Help

http://search.yahoo.com/search?p=julia+roberts&ei=UTF-8&fr=moz35

shalimar sunnyvale - Yahoo! Search Re... | julia roberts - Yahoo! Search Res...

YAHOO!

julia roberts

Search

julia roberts twins
julia roberts movies
lyle julia roberts
julia roberts babies
julia roberts henry daniel moder

Explore related concepts:
actor
Epi...
Search for julia roberts actor
Pretty Woman
Best Actress

julia roberts photos
Julia Roberts News
julia roberts biography
Julia Fiona Roberts

Search Pad

SearchScan - On

40,500,000 results for
julia roberts

Related People

Scarlett Johansson

Emma Roberts

Hilary Swank

Lindsay Lohan

Tom Hanks

Halle Berry

Julia Roberts - Image Results

News & Photos | Videos | Twitter




more Julia Roberts photos...

Latest News:

Hindus concerned about Julia Roberts' "Eat, Pray, Love" - New Kerala - 6 hours ago
Trailer For 'Eat Pray Love' Starring Julia Roberts - KPBS San Diego - Mar 19 03:18pm
Link Party: Julia Roberts' New Movie Will Teach You How to Live - E! Online - Mar 18 05:48pm

more Julia Roberts news...

Julia Roberts - Wikipedia

Early life | Career | Influence | Personal life

Julia Fiona Roberts is an American actress. She is known for starring in the romantic comedy Pretty Woman opposite Richard Gere, which grossed $463 million worldwide. After receiving...
en.wikipedia.org/wiki/Julia_Roberts - 122k - Cached

Yahoo! Search Direct

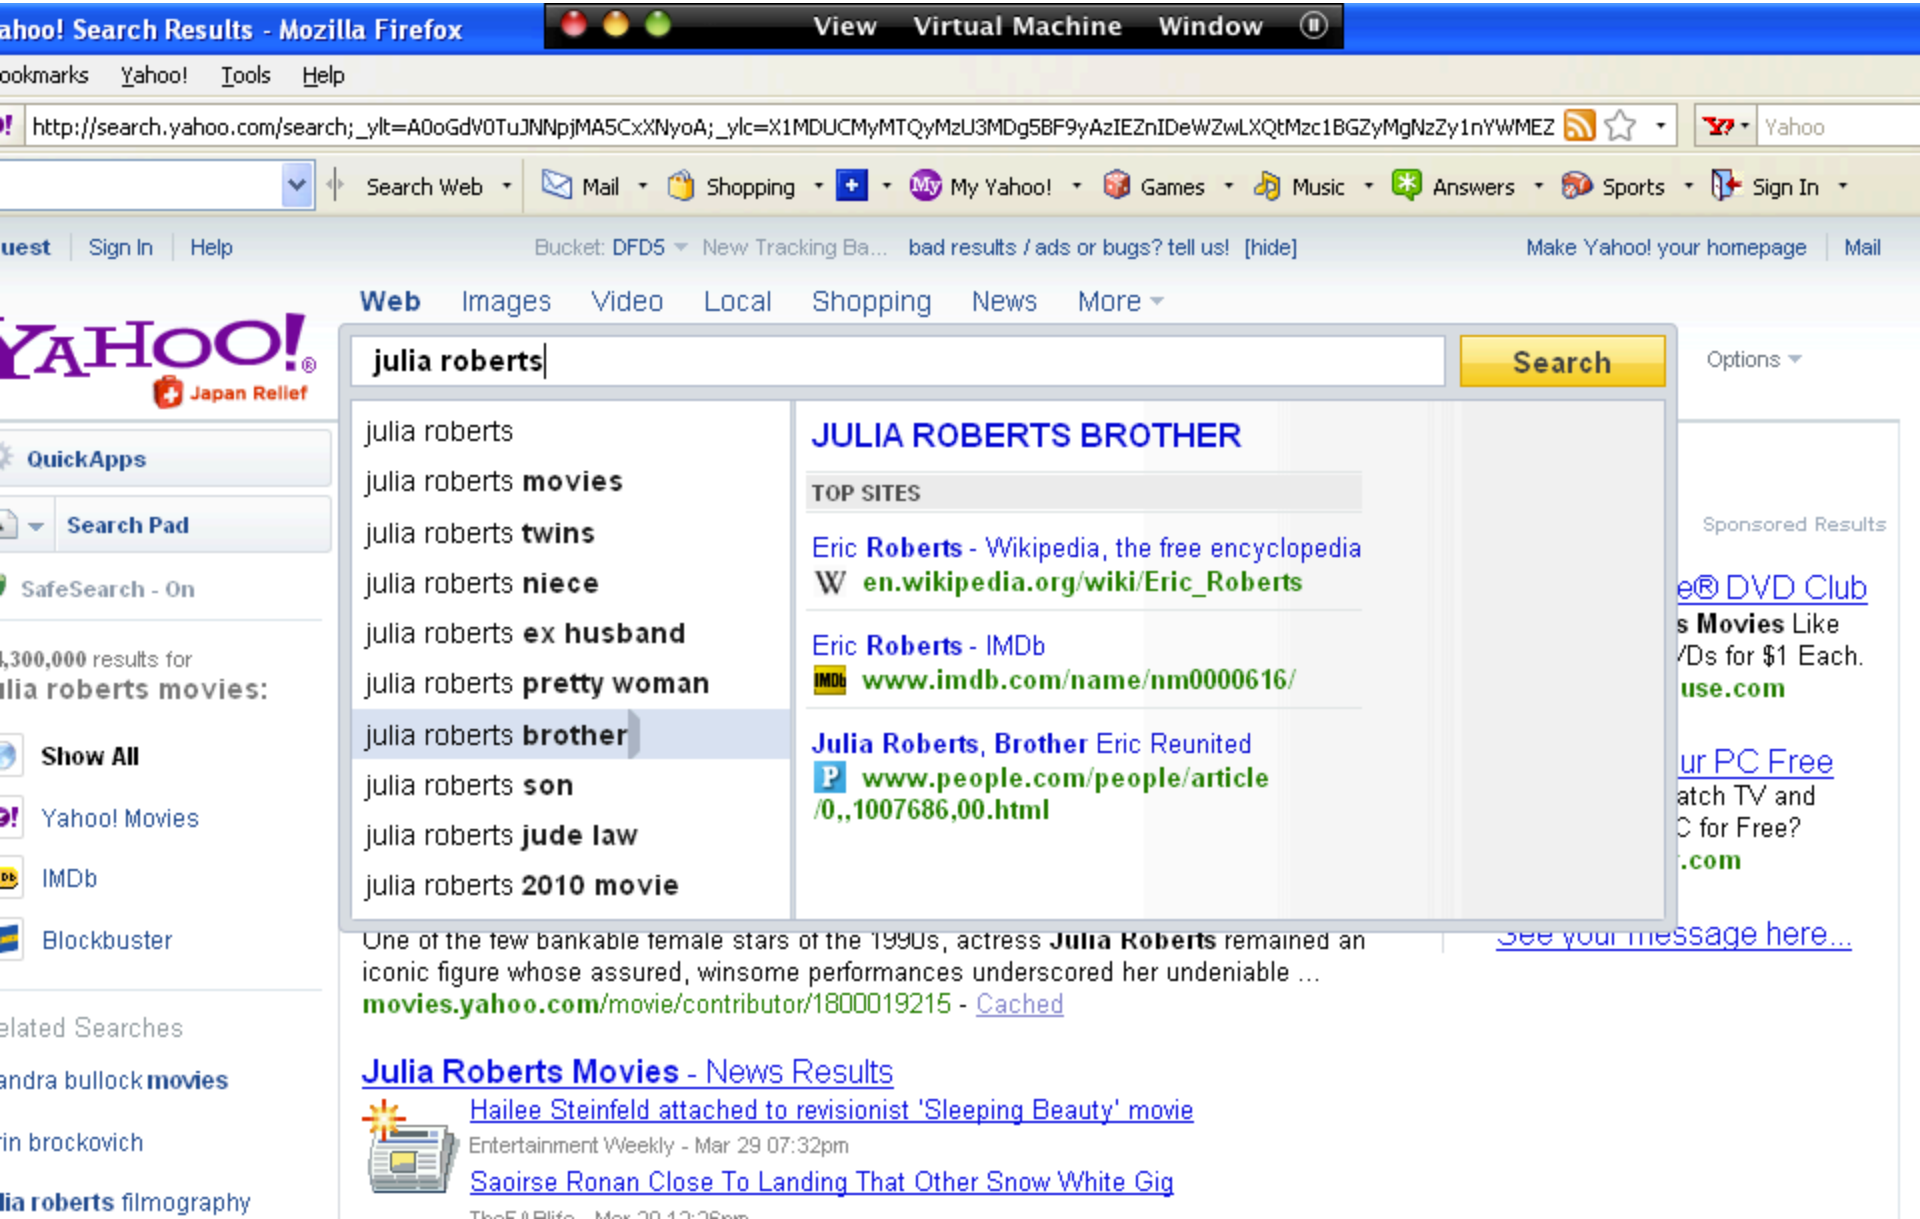


berts - Yahoo! Search Results - Mozilla Firefox
View History Bookmarks Yahoo! Tools Help
http://search.yahoo.com/search?p=julia+roberts&ei=UTF-8&fr=moz35
julia roberts - Yahoo! Search Res...
YAHOO!
julia roberts
Search
julia roberts twins
julia roberts movies
lyle julia roberts
julia roberts babies
julia roberts henry daniel moder
Explore related concepts:
actor
Best Actress
julia roberts photos
Julia Roberts News
julia roberts biography
Julia Fiona Roberts
Search for julia roberts actor
Search Pad
SearchScan - On
40,500,000 results for
julia roberts
Related People
Scarlett Johansson
Emma Roberts
Lindsay Lohan
Tom Hanks
Halle Berry
Julia Roberts - Image Results
News & Photos
Videos
more Julia Roberts news...
Julia Roberts - Wikipedia
Early life | Career | Influence | Personal life
Julia Fiona Roberts is an American actress. She is known for starring in the romantic comedy Pretty Woman opposite Richard Gere, which grossed $463 million worldwide. After receiving...
en.wikipedia.org/wiki/Julia_Roberts - 122k - Cached
Concept-centric
Information
ANSWERS, NOT LINKS

Web of Pages

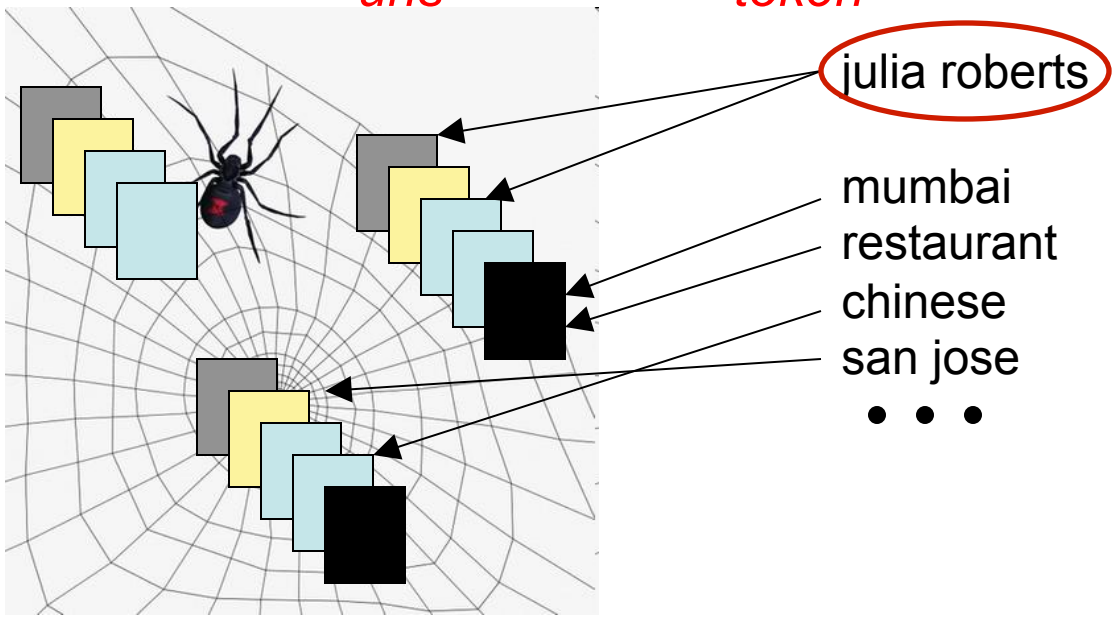

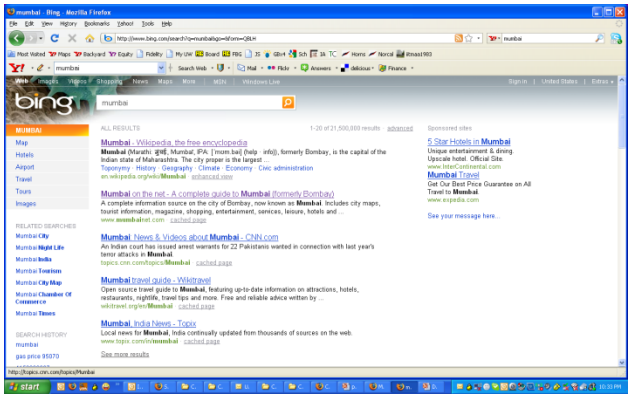

WEB

INDEX

SERP

Web of Concepts

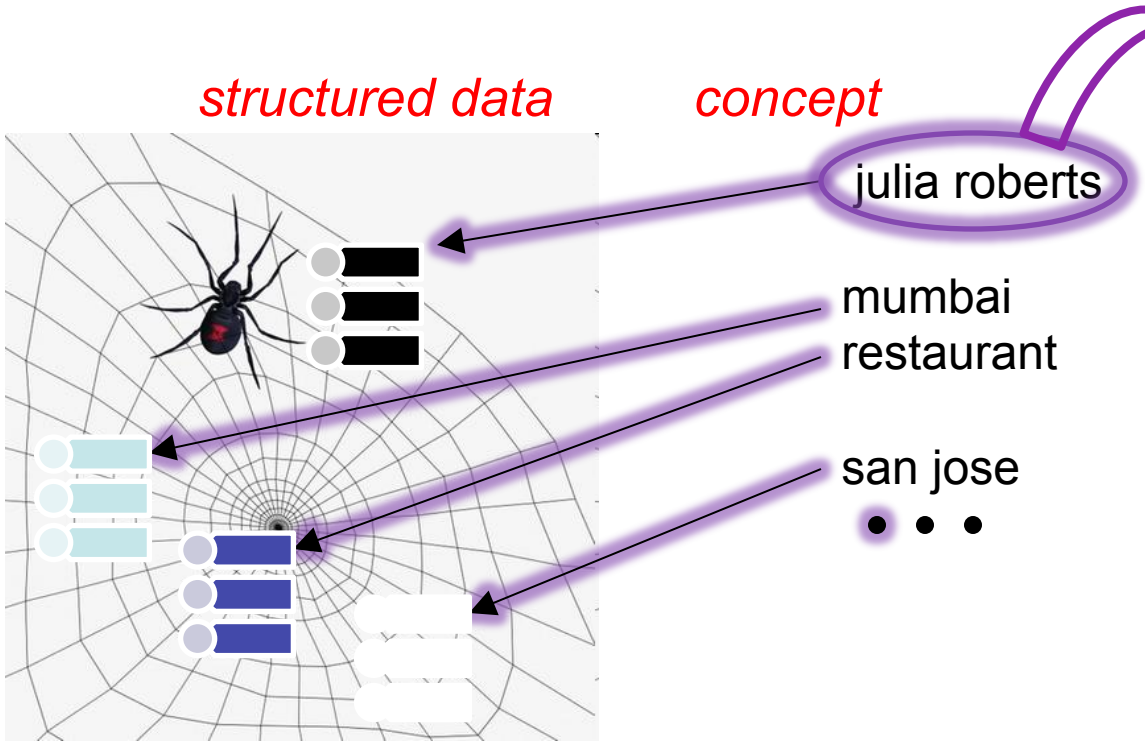

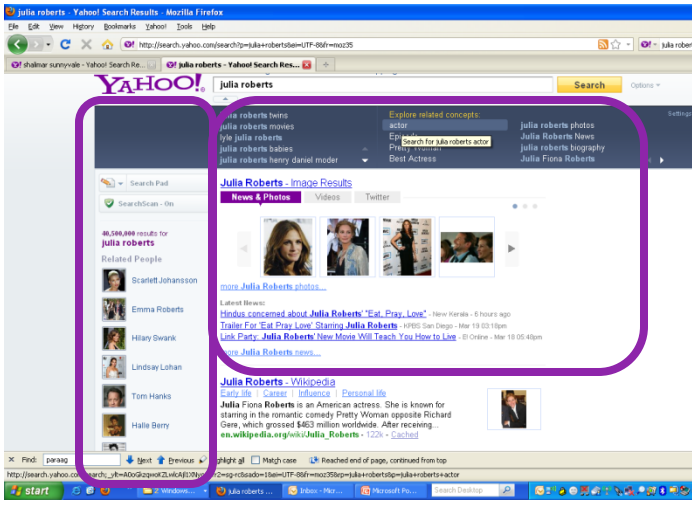

Aggregated KB **INDEX** **SERP**

The “index” is keyed by concept instance, and organizes all relevant information, wherever it is drawn from, in semantically meaningful ways

ANSWERS, NOT LINKS

Web of Concepts

- The web is a source of information about concepts people care about
- The “index” is keyed by concept instance, and organizes all relevant information, wherever it is drawn from, in semantically meaningful ways
  - E.g., Information about Mumbai, or the Dish-Dash restaurant, is organized in terms of meaningful attributes (population, type of cuisine) and can be retrieved and presented to address whatever the user wants to know about Mumbai or Dish-Dash

N. Dalvi et al.: A Web of Concepts, PODS 2009

mumbai - Yahoo! Search Results - Mozilla Firefox

File Edit View History Bookmarks Yahoo! Tools Help

http://search.yahoo.com/search?p=mumbai&ei=UTF-8&fr=moz2

Most Visited Maps Backyard Equity Fidelity My UW Board FBG JS GBx4 Sch IA TC Horns Norcal iitmaa1983

mumbai Search Web Mail Flickr Answers delicious Finance

no bucket bad results / ads or bugs? tell us! Yahoo! My Yahoo! Mail Welcome, sc

Web | Images | Video | Local | Shopping | more

mumbai Search Options Customize

YAHO

mumbai university
mtnl mumbai
mumbai attacks
mumbai mirror
mumbai india

Explore related concepts:

Mumbai Central
South Mumbai
population
tourism

Maharashtra
Mumbai, India
streets
Marathi

1 - 10 of 217,000,000 for mumbai (About) - 0.28 s |

Mumbai, India - Visitor Guide
travel.yahoo.com
Hotels | Restaurant Guide | Flights | Map

Top Rated Things To Do (134)
1 Gateway of India
2 SKS Supercrafts
3 Matheran
More Things To Do...

Mumbai - Wikipedia, the free encyclopedia
Toponomy | History | Geography | Climate
Includes geography, economy, climate, civic administration, utility services, people and culture, media, education, sports and recreation, transport, and history ...
en.wikipedia.org/wiki/Mumbai - 458k - Cached

Mumbai on the net - A complete guide to Mumbai (formerly Bombay)
With news, maps, tourist information, emergency contacts, education, utilities, health, hotels

Visting Mumbai?
Find cheap deals on 4000 hotel sites with 1 search! No booking fees.
www.Kayak.com/Hotels

2009 Top Work At Home
Looking For mumbai?
EricGetsCash.com

Hotels in Mumbai
Looking for Mumbai Hotels? Cheap Deals - 35+ Sites - Reviews.
mumbai.hotelhelpers.com

Mumbai
Find Low Rates Worldwide. Flights, Hotels & More at ORBITZ.
www.ORBITZ.com

Done

start

Mumbai, India

HOTELS, VACATIONS, TOURISM

WEATHER: HAZE, 83F
Annual weather »

Ever since the opening of the Suez Canal in the 1860s, the principal gateway to the Indian subcontinent has been MUMBAI (Bombay) , the city Aldous Huxley famously described as "the most appalling . . . of either hemisphere". Travellers tend to regard time spent here as a rite of passage to be survived rather than savoured. But as the powerhouse of Indian business, industry and trade, and the source of its most seductive media images, the Maharashtrian capital can be a compelling place to kill time. Whether or nor you find the experience enjoyable, however, will depend largely on how well you handle the ... more »




View all (62) »

ADD TO TRIP BEEN THERE WANT TO GO

Understanding Web Content

with a little help from your friends

Generative Model of the Web

The true world database

Site queries

Information loss

Noise addition

Site layout

Surround generation

Document Corpus

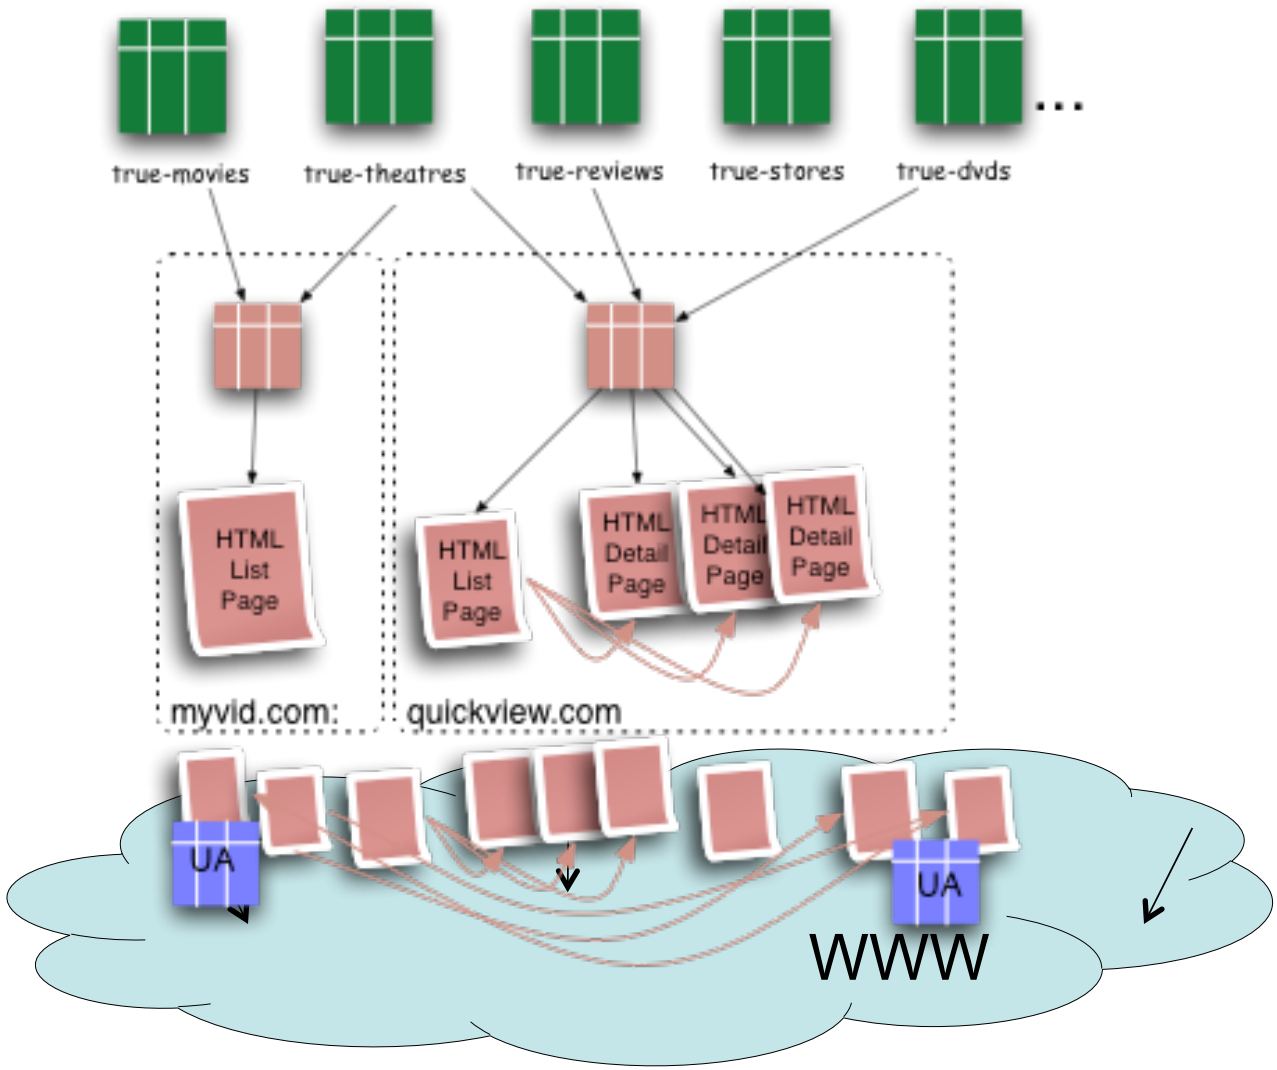

33

Surfacing, Extraction, Integration

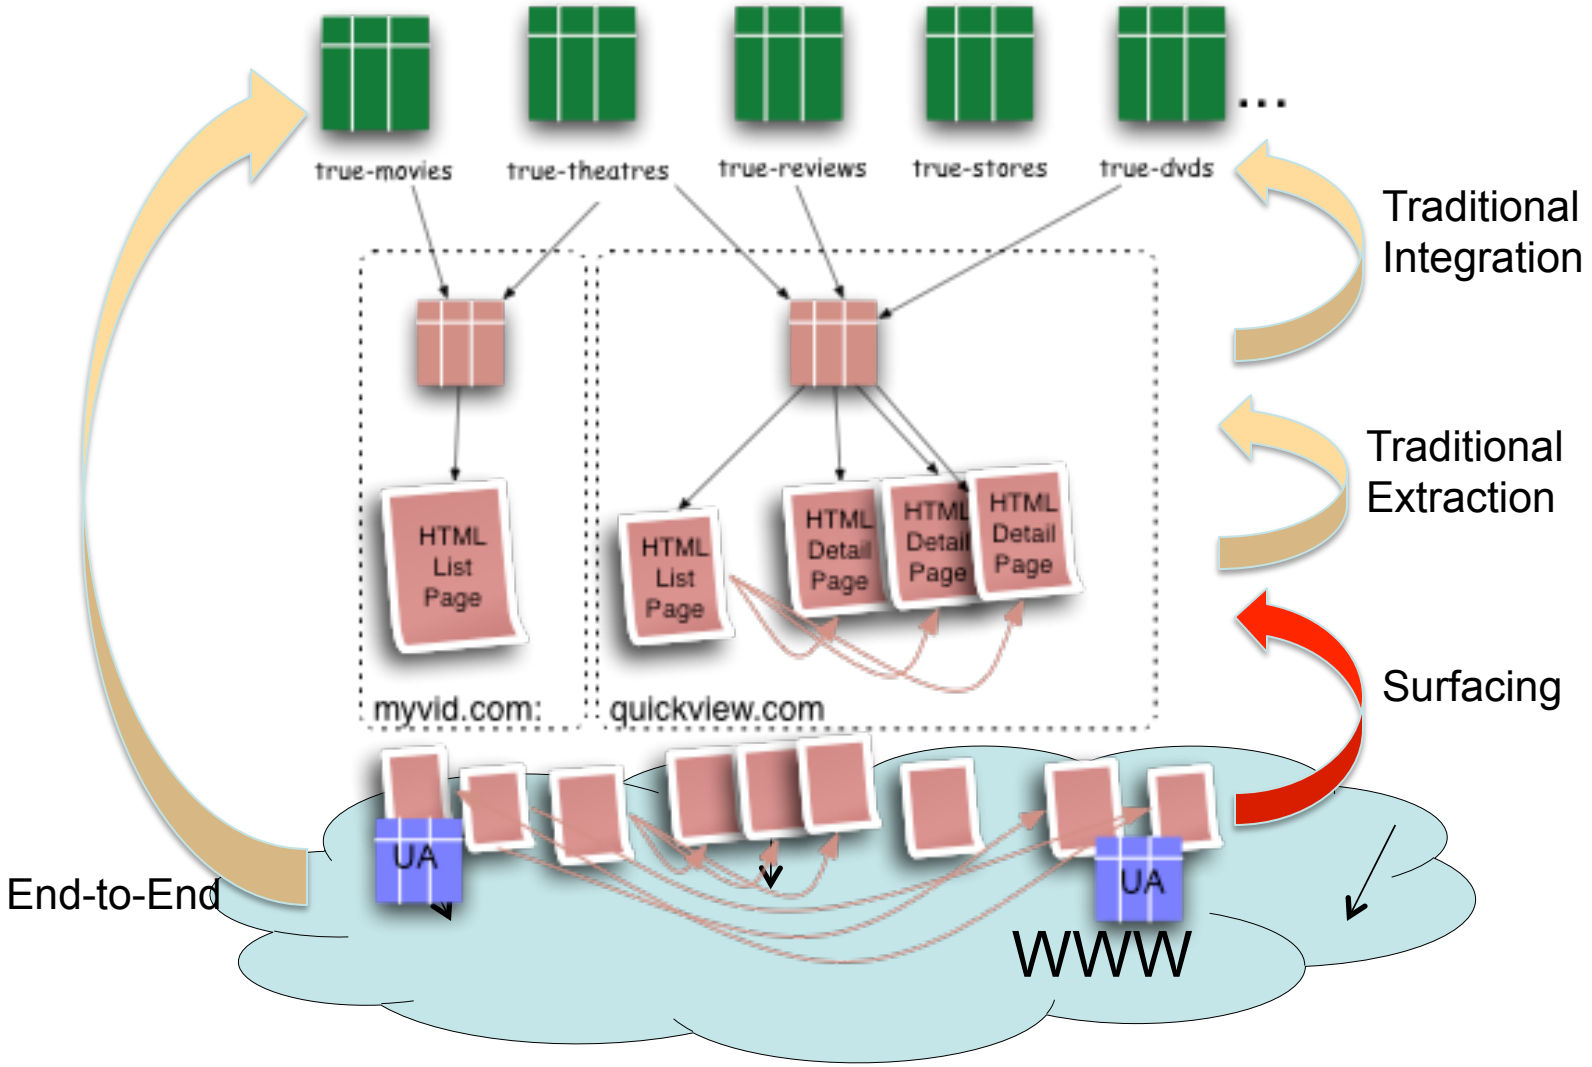

34

DBLife: Community Information Mgmt

- Integrated information about a (focused) real-world community
- Collaboratively built and maintained by the community
  - Semantic web, “bottom-up”
- Joint work with AnHai Doan, Pedro Domingos, Warren Shen and others at Wisconsin

Summary of Gaps

We have observed three key gaps in the state-of-the-art

Gap 1: Finding the pages to extract from may be a significant portion of the supervision cost

Gap 2: Must extract across diverse content types to gather all attributes

Gap 3: Supervision costs for information integration may exceed costs for extraction

36

Integration: Entity Extraction/Deduping

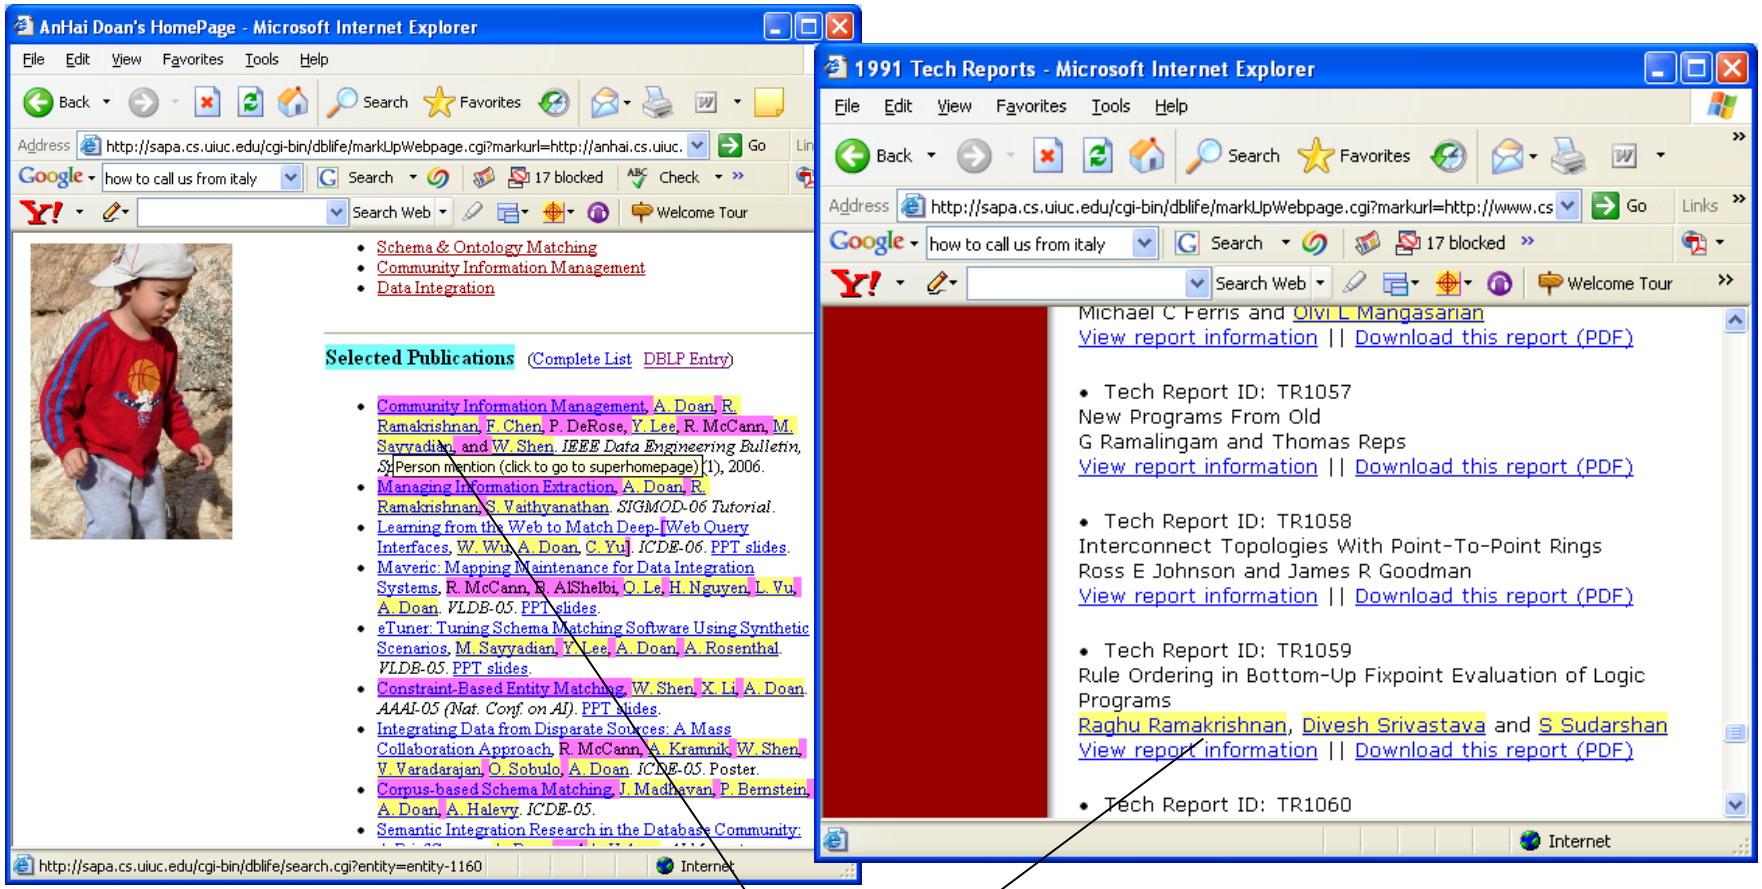

Mass Collaboration for IE

Not David!

Picture is removed if enough users vote “no”.

Mass Collaboration Meets Spam

Some Closely Related Work

- Avatar (IBM Almaden)
- DeepWeb & WebTables (Google)
- DBLife (U. Wisconsin)
- KnowItAll & TextRunner (U. Washington)
- Nell (CMU)
- Rexa (U. Mass)

Social Networking

Social Networks: Distro to Go!

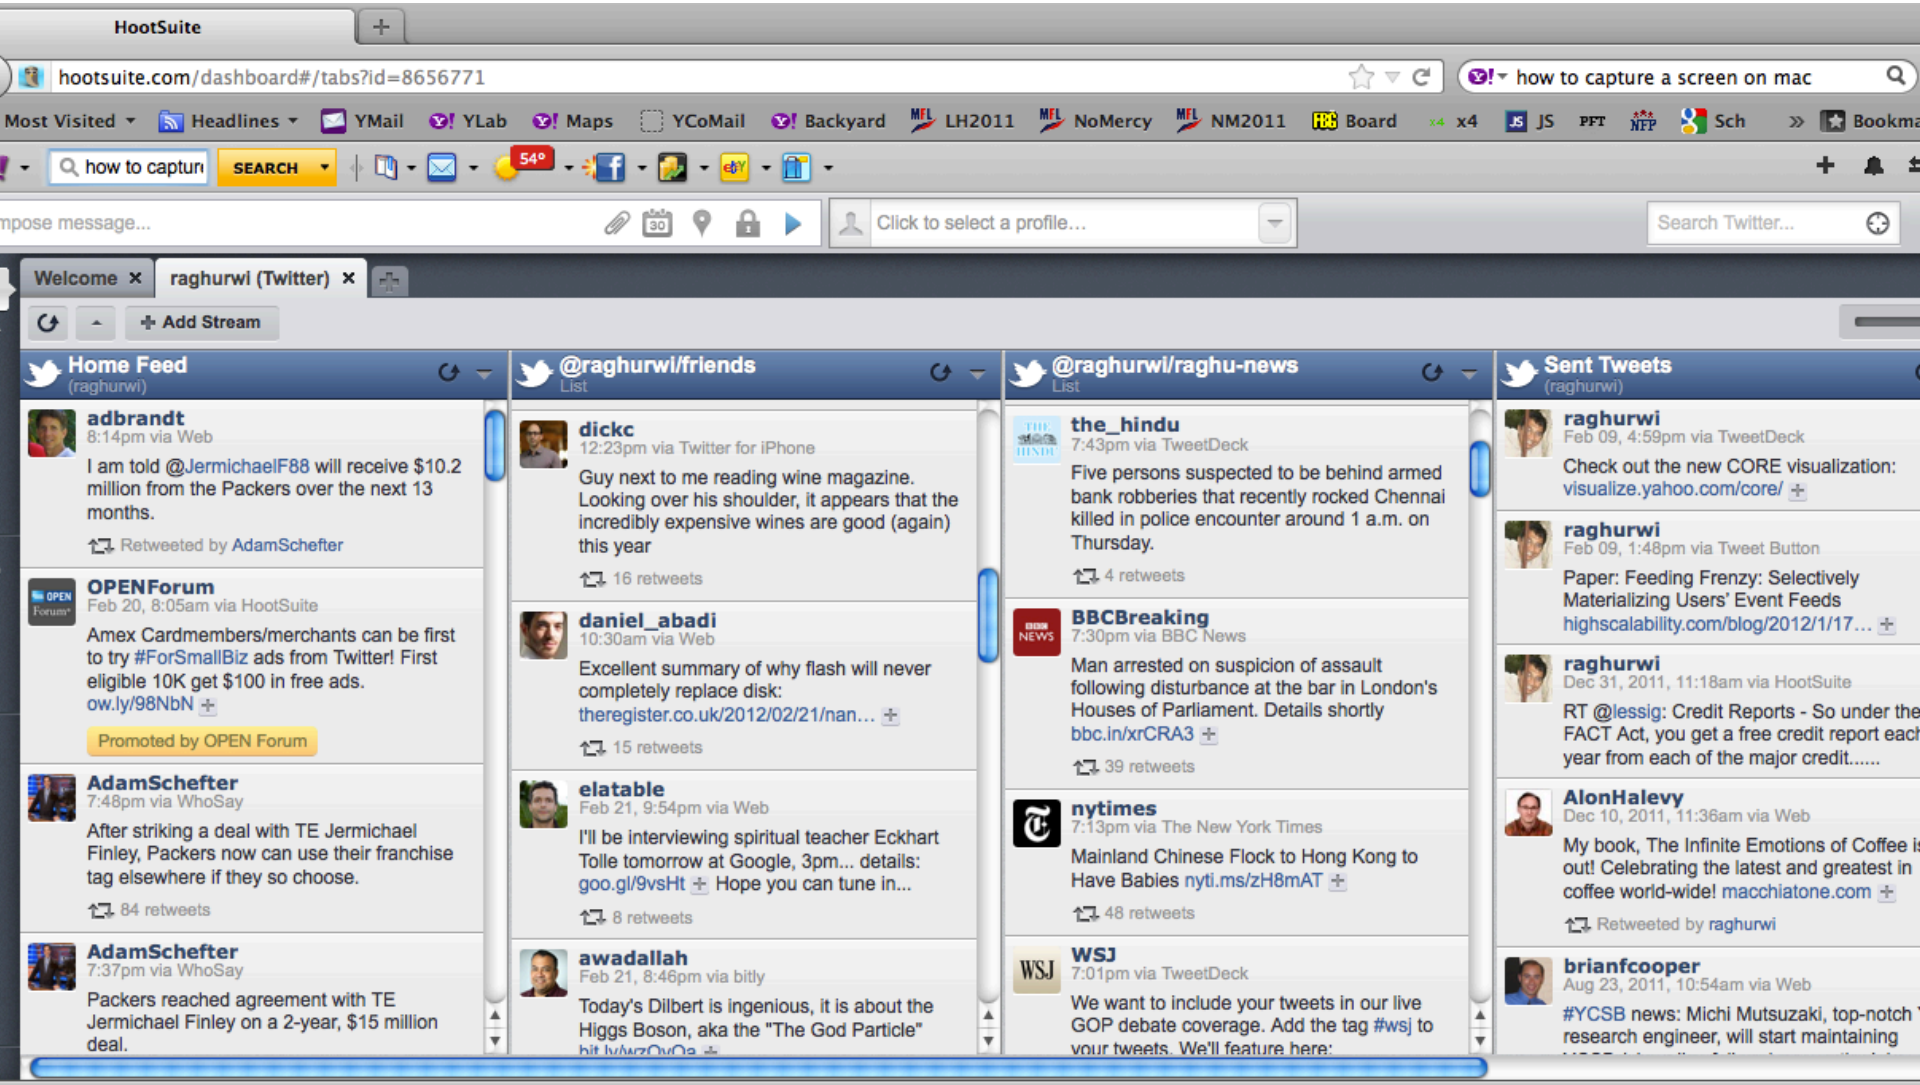

IntoNow:

Connect with your friends around the shows you love

IntoNow from Yahoo! makes engaging with your friends around your favorite television shows easy and fun. Just tap the green button when you're watching, and IntoNow will identify the show, right down to the episode. Once identified it's easy to share with your friends on Twitter or Facebook.

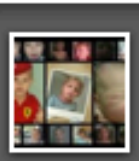

Candace T is watching **SpongeBob SquarePants**
Season 1, Episode 21

John A is watching **The Walking Dead**
Season 2, Episode 9

David S is watching **Parenthood**
Season 3, Episode 17

Discover what your friends are into

Use IntoNow to discover new shows, discuss your favorites with friends, and learn more about whatever you're into! IntoNow is a great source of recommendations for what's on right now for you, based on what your friends are into. Find out what shows you share in common, get notified

Social Advertising

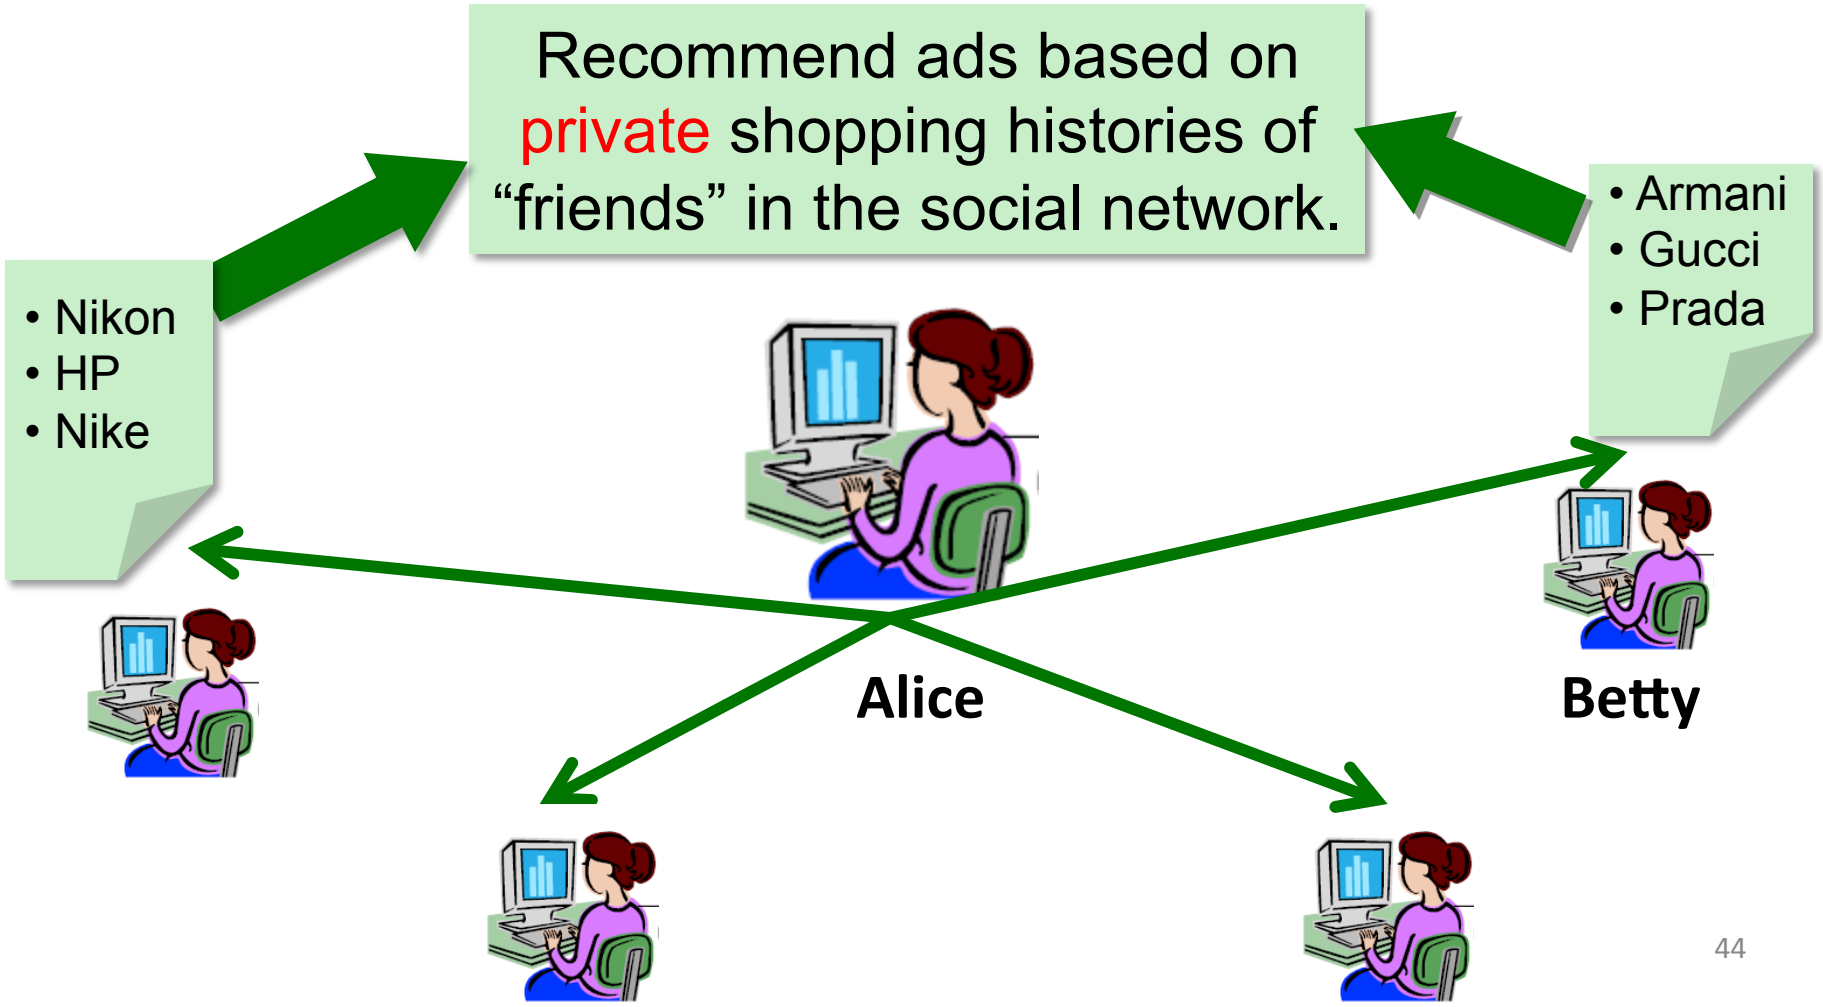

(Personalized Social Recommendations: Machanavajjhala et al., VLDB 2011)

Privacy in Social Advertising

Items (products/people) liked by Alice's friends are better recommendations for Alice

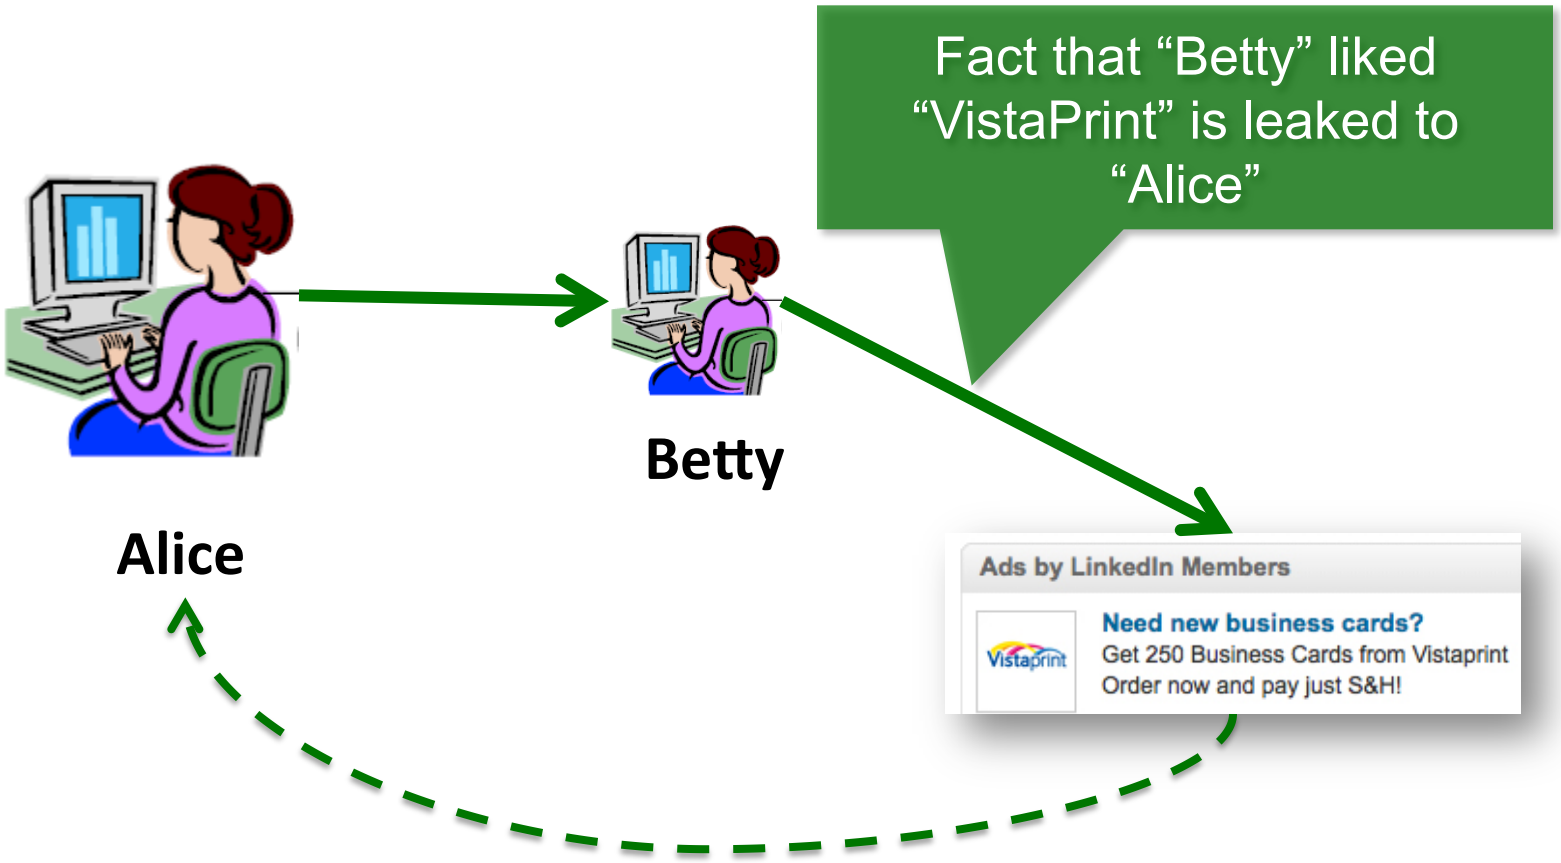

45

Privacy in Social Advertising

Alice is recommended ‘X’

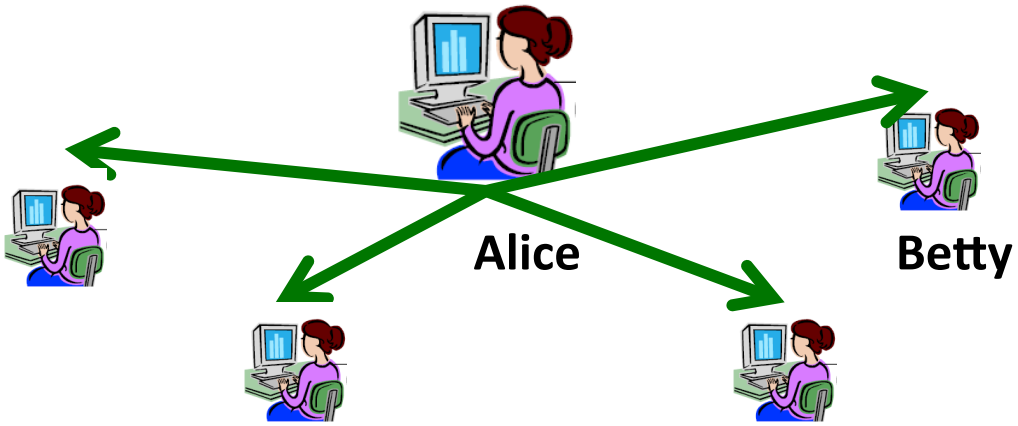

Can we provide *accurate* recommendations to Alice
while ensuring that
Alice *cannot* deduce that Betty likes ‘X’ ?

46

Takeaway …

- "For majority of the nodes in the network, recommendations must either be inaccurate or violate differential privacy!"
  - Maybe this is a "bad idea"
  - Or, Maybe **differential privacy is too strong a privacy definition to shoot for**.

47

Related Work

- Amazon, Netflix, Y! Music, etc.:
  - Collaborative filtering with large content pool (see KDD Cup, 2011)
  - Achieve lift by eliminating bad articles
  - We have a small number of high quality articles
- Search, Advertising
  - Matching problem with large content pool
  - Match through feature based models

Mobius: Unified Messaging and Data for Mobile Apps

Byung-Gon Chun, Carlo Curino, Russell Sears, Alexander Shraer, Samuel Madden*, Raghu Ramakrishnan

Yahoo! Research, *MIT

mobicloud@yahoo-inc.com

Mobile Apps Connect to Cloud

Access, Create, Share Data/Messages

84% of apps are “cloud-connected”

(source Appcelerator)

Mobile App Developers' Challenges

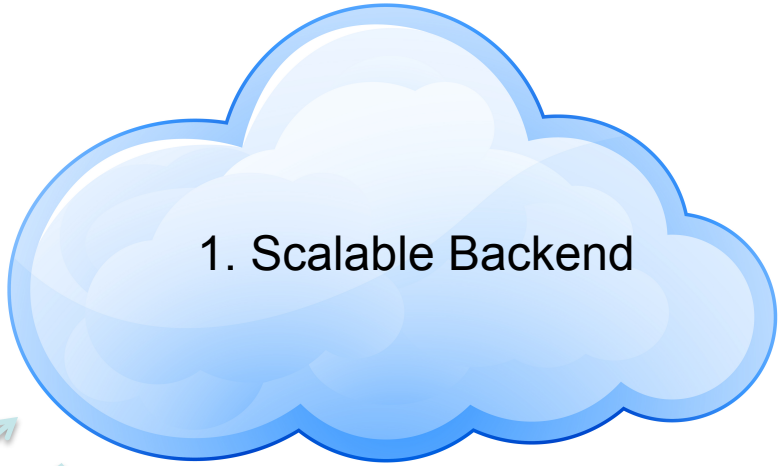

Pull

News

Mobile App Developers' Challenges

2. Speed: wireless networks
- Caching
- Prefetching

Pull

News

Mobile App Developers' Challenges

3. Notification

Push Alerts

Mobile App Developers' Challenges

4. Messaging

Mobile App Developers' Challenges

5. Disconnection

- Read news
- Post on Facebook

Mobius

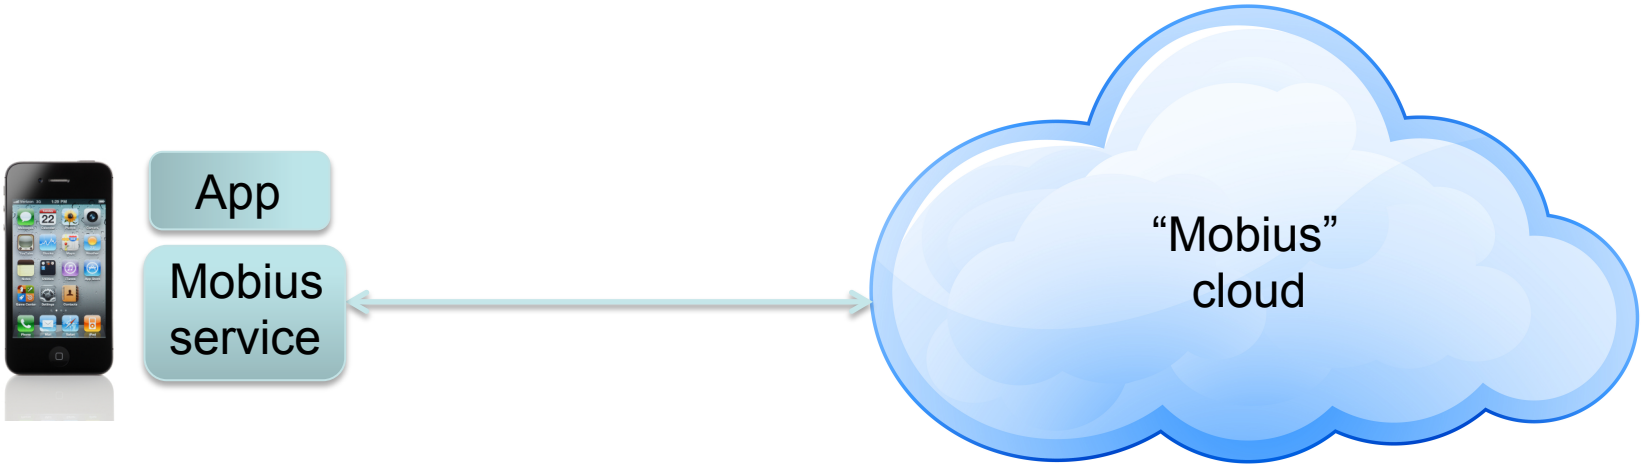

- MUD API: (continuous)read/write of data
- Protocols to handle disconnection (and svn-like conflict resolution)
- Predicate-based notification (receiver and sender predicates)
- Automatic caching/prefetching (global policy generated by the cloud)

Mobius Architecture

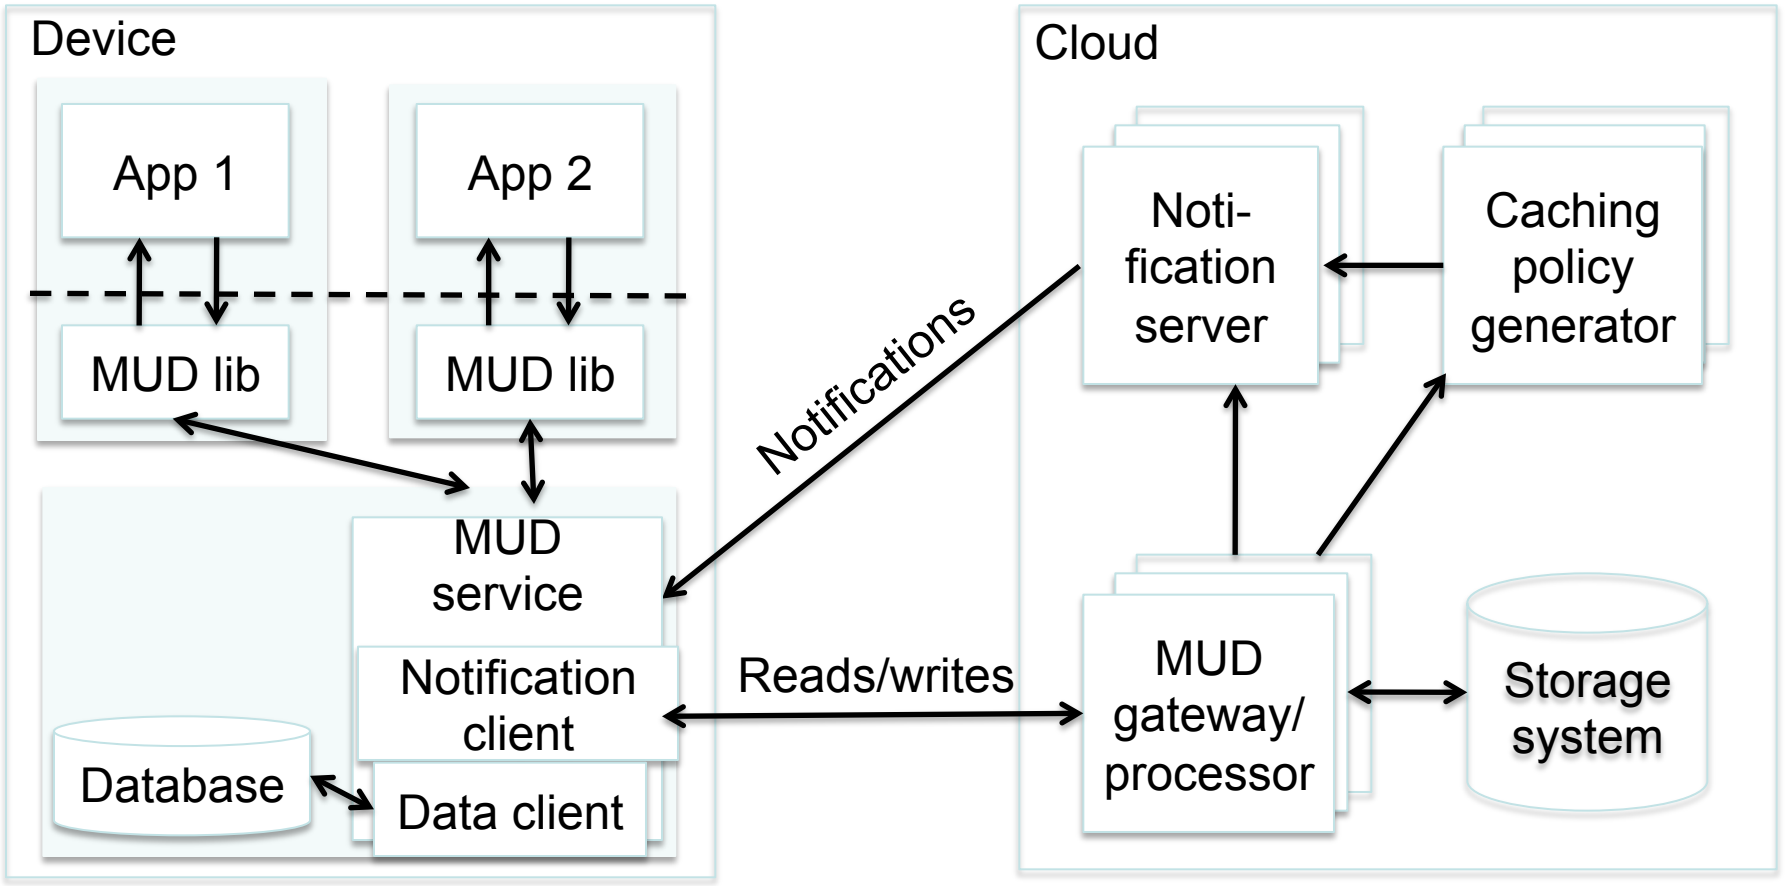

Main Points

- The medium is being massaged
  - i.e., the canvas is no longer passive
  - It is also very personal, and `always on'
  - *(With apologies to Marshall McLuhan)*
- Search and browse converging
  - Semantics is the new frontier
  - Social networks the new distribution channel
- Learning from Big Data is a key
  - Value of data is now recognized
  - Data marketplaces will emerge
  - Compute platforms will be increasingly data-centric