

Large Scale Internet Search at Ask.com



-
- Tao Yang
Chief Scientist and Senior Vice President

InfoScale 2006



Outline

- **Overview of the company and products**
- **Core techniques for page ranking**
 - ExpertRank
- **Challenges in building scalable search services**
 - Neptune clustering middleware.
 - Fault detection and isolation.

Ask.com: Focused on Delivering a Better Search Experience

- **Innovative search technologies.**
- **#6 U.S. Web Property; #8 Global in terms of user coverage**
 - 28.5% reach - Active North American Audience with 48.8 million unique users
 - 133 million global unique users for ASK worldwide sites: USA, UK, Germany, France, Italy, Japan, Spain, Netherlands.
- **A Division of IAC Search and Media (Formally Ask Jeeves)**



Sectors of IAC (InterActiveCorp)

•Retailing	 HSN.com	Cornerstone Brands, Inc.  AMERICA'S STORE	 HSE24	 SHOEbuy.com
•Services	 ticketmaster	 RealEstate.com Inventing a better way	 PRC	 ServiceMagic
•Media & Advertising	 Ask.com	 Citysearch	 evite	
•Membership & Subscriptions	 Interval International	 match.com	 entertainment	
•Emerging Businesses	 gifts.com	 Quiz.tv	 pronto	 Buy

IAC (InterActiveCorp)

- Fortune 500 company
 - Create, acquire and build businesses with leading positions in interactive markets.
- 60 specialized & global brands
- 28K+ employees
- \$5.8 billion – 2005 Revenue
- \$668 million – 2005 OIBA (Profit)
- \$1.5 billion net cash

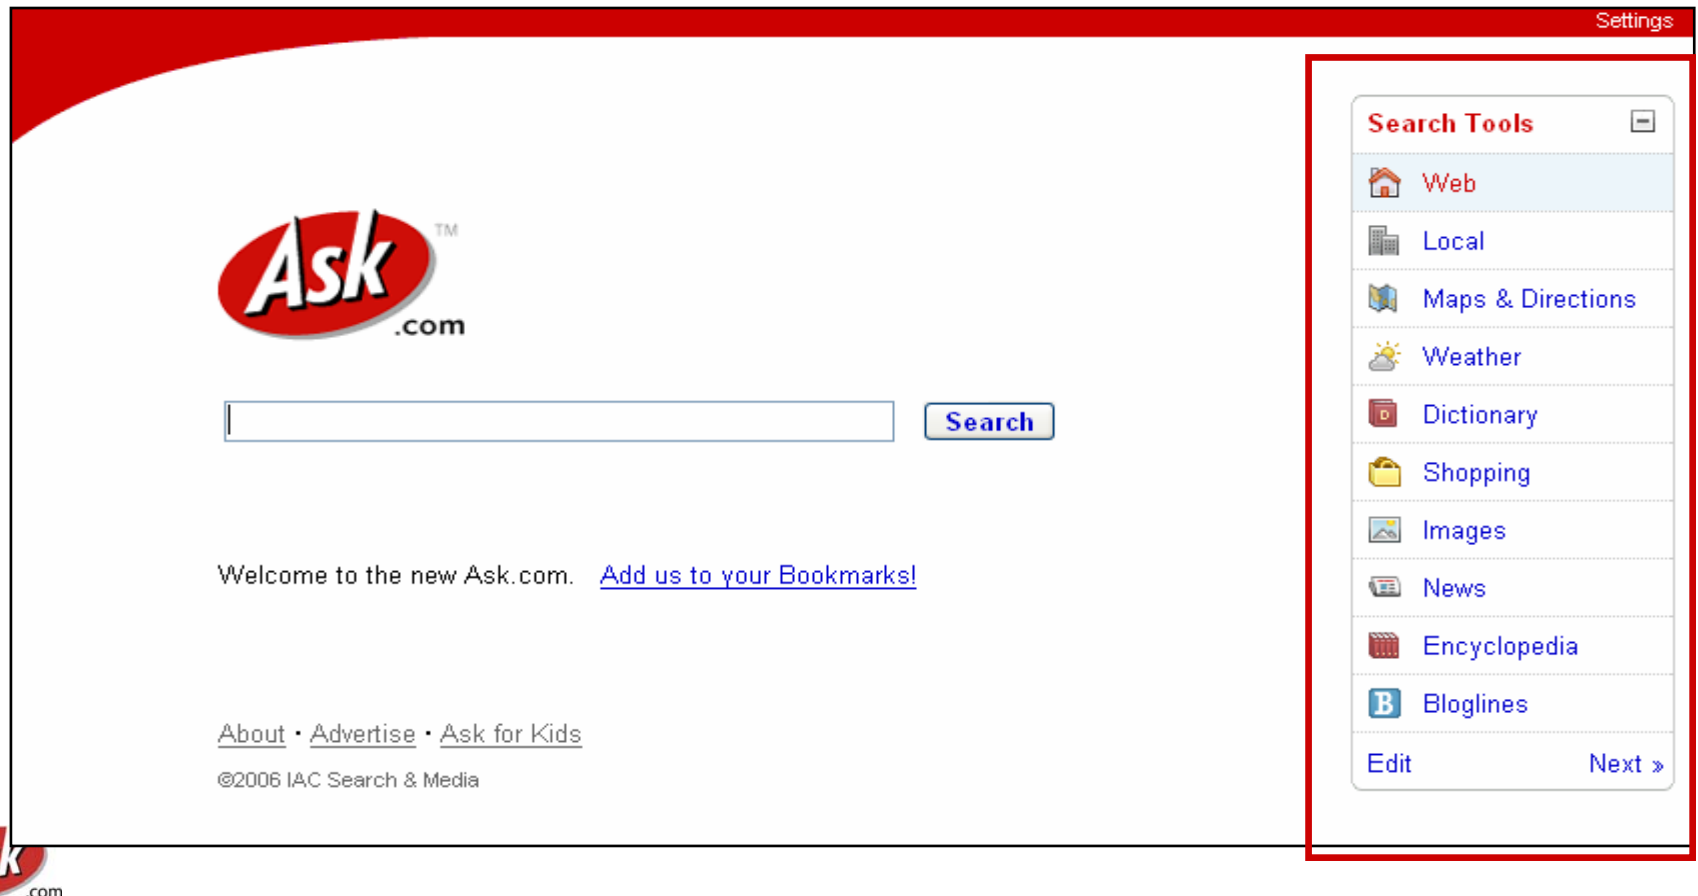


IAC / InterActiveCorp



Ask.com Site Relaunching and branding in Q1 2006

- Cleaner interface with a list of search tools



Site Features: Smart Answer

The screenshot shows a Microsoft Internet Explorer browser window with the title "hot air balloon - Ask.com Search - Microsoft Internet Explorer". The address bar contains the URL "http://www.ask.com/web?q=hot+air+balloon&qsrc=1&o=0". The search bar contains the text "hot air balloon" and a "Search" button. The search results are displayed on a red background with the text "Web Search" and "Showing results 1-10 of 739,300".

The first result is an encyclopedia entry titled "Encyclopedia: Hot air balloon" with a source of Wikipedia. It includes a small image of a hot air balloon and a description: "Hot air balloons are the oldest successful human flight technology, dating back to the Montgolfier brothers' invention in Annonay, France in 1783. The first flight carrying humans was made on November 21, 1783, in Paris by Pilâtre de Rozier and the Marquis d'Arlandes. Hot air ... [More »](#)".

Below the encyclopedia entry are three sponsored results:

- Hot Air Balloon Rides** (Sponsored Results): 200 Locations. 100% Safety Record. \$169.95 No Hidden Fees. Member BBB! www.BalloonRidesAmerica.com
- Hot Air Balloon Flight**: Airborne Gifts at Xperience Days www.XperienceDays.com
- Free Hot Air Balloon Ride**: Experience The Ride Of Your Life! \$500 Value. Complete Our Survey. HotAirRide.Lesure-Offer.Com

At the bottom of the search results is a link for "Ballooning - hot air balloon news and information" with the text "Hot air ballooning's home on the web. Learn where to get a balloon ride, how balloons fly, how to become a hot air pilot or find a balloon festival." and the URL "www.XperienceDays.com".

On the right side of the page, there is a "Narrow Your Search" section with links to "Hot Air Balloon Works", "Who Invented the Hot Air Balloon", "Build a Hot Air Balloon Model", "How to Make a Hot Air Balloon", "First Hot Air Balloon Invented", "Flight And Hot Air Balloons", and "Hot Air Balloon Science". Below this is an "Expand Your Search" section with a link to "Airplane".



Topic Zooming with Search Suggestions

hot air balloon - Ask.com Search - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites

Address http://www.ask.com/web?q=hot+air+balloon&qsrc=1&o=312&dm=all Go Links

www.reletterdays.co.uk

Hot Air Balloon Flights
Take a **hot air balloon** flight over Britain. Book online now!
www.balloonsoverbritain.com

Ballooning - hot air balloon news and information
Hot air ballooning's home on the web. Learn where to get a **balloon** ride, how balloons fly, how to become a **balloon** pilot, or find a **balloon** festival...
www.launch.net/ · Cached · Save

Howstuffworks "How Hot Air Balloons Work"
Popular Searches > **Hot Air Balloon** > New York > Personal Watercraft > Pickpocket > Rock Climbing > Roller Coaster > Skydiving
www.howstuffworks.com/hot-air-balloon.htm · Cached · Save

Hot Air Ballooning Pictures, History and Information from...
...eBalloon.org is the online ballooning encyclopedia - a one stop resource for anyone who wants to find out how a **hot air balloon** works, the history...
www.eballoon.org/ · Cached · Save

BalloonZone ~ Hot Air Ballooning ~ hot air balloon festivals &...
Read an abbreviated history of **hot air** ballooning and learn about how the **balloon** is put together .
www.balloonzone.com/ · Cached · Save

BalloonZone ~ Short History of Hot Air Ballooning
...modern **hot-air balloon**. The first man-carrying free flight took place at Bruning, Nebraska, in the Raven prototype 'modern' **hot-air balloon**.
www.balloonzone.com/history.html · Cached · Save

Balloons
Hot Air Balloon Science
More »

Expand Your Search

Airplane
Helicopter
Train
Blimp
Montgolfier Brothers
More »

Related Names

Joseph Montgolfier
Jacques Charles
Richard Branson

Internet

Site Feature: Web Direct Answer

what is the largest city in usa - Ask.com Search - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Home Search Favorites Refresh Mail Print View Source Help

Address <http://www.ask.com/web?q=what+is+the+largest+city+in+usa&qsrc=1&o=0> Go Links

Settings

Ask.com Web · Images · News · Shopping · More »

what is the largest city in usa [Advanced Search](#)

Web Search Showing results 1-10 of 4,784,000

[what is the largest city in usa](#) [Web Answer]
The island makes up one of the five boroughs of **New York City, the USA's largest city.**
en.wikipedia.org/wiki/Grasmere,_Staten_I... · Save · [See 8 more Web Answers](#)

[Population of the 100 Largest Cities and Other Urban Places In The](#)
In 1990, the 20 **largest** cities were the **largest** cities in only 13 of the 20 **largest** urbanized areas. In the 27 cases where a **city** and/or its...
www.census.gov/population/www/documentation/twps0027.html · Cached · Save

[USA City hotels, the largest directory of discount hotels in USA](#)
Welcome to the **USA City** hotels Guide. The **largest** directory of discount hotels in **USA**. The **USA City** hotels guide offers you direct and easy access...
www.usacityhotels.com/ · Cached · Save

[Re: What are the 10 largest cities in the USA](#)
Re: What are the 10 **largest** cities in the **USA** ... cen.) :: 9. Istanbul, Turkey 7,784,169 (1995 est.) :: 10. New York City, USA 7,380,906
www.usahistory.com/trivia/wwwboard/archive/messages/1355.html · Cached · Save

[What are the 10 largest cities in the USA](#)

Narrow Your Search

- [Largest United States Cities](#)
- [USA 50 Largest Cities](#)
- [What Is the Longest River in USA?](#)
- [How Big Is USA](#)
- [What Is the Largest City in the USA in 2004?](#)
- [What Is the Most Populated City in USA](#)
- [USA Facts And Figures](#)
- [Largest Population of US Cities](#)
- [1st Largest US City](#)
- [10 Largest Cities in USA](#)

Expand Your Search

Internet

More Site Features - Binoculars

- Our **Binoculars tool** lets you see what a site looks like before clicking to visit it

[Tour de France : bienvenue sur le site officiel](#)
24-01-06 - Start of the 2007 **Tour de France** from London ... 09-12-05 - Organizers of the three major Tours...
The Store...
 www.letour.fr/indexus.html • Cached • Save
Results from www.letour.fr

[Lance Arm...](#)
The official Lan...
discussion for...
 [www.cyclin...](http://www.lanc...</p><p><a href=)
2006 **Tour de...**
 [VeloNews: ...](http://www.cycl...</p><p><a href=)
Will you see a...
 [trekbikes.co...](http://www.velo...</p><p><a href=)
Official site fe...
maintenance t...
 [... and the USPS team. Guestbook,](http://www.trek...</p><p><a href=)

[...ervice](#)
actions, News, Photos More Major races...

[...lot sure. View Results Archived Polls...](#)

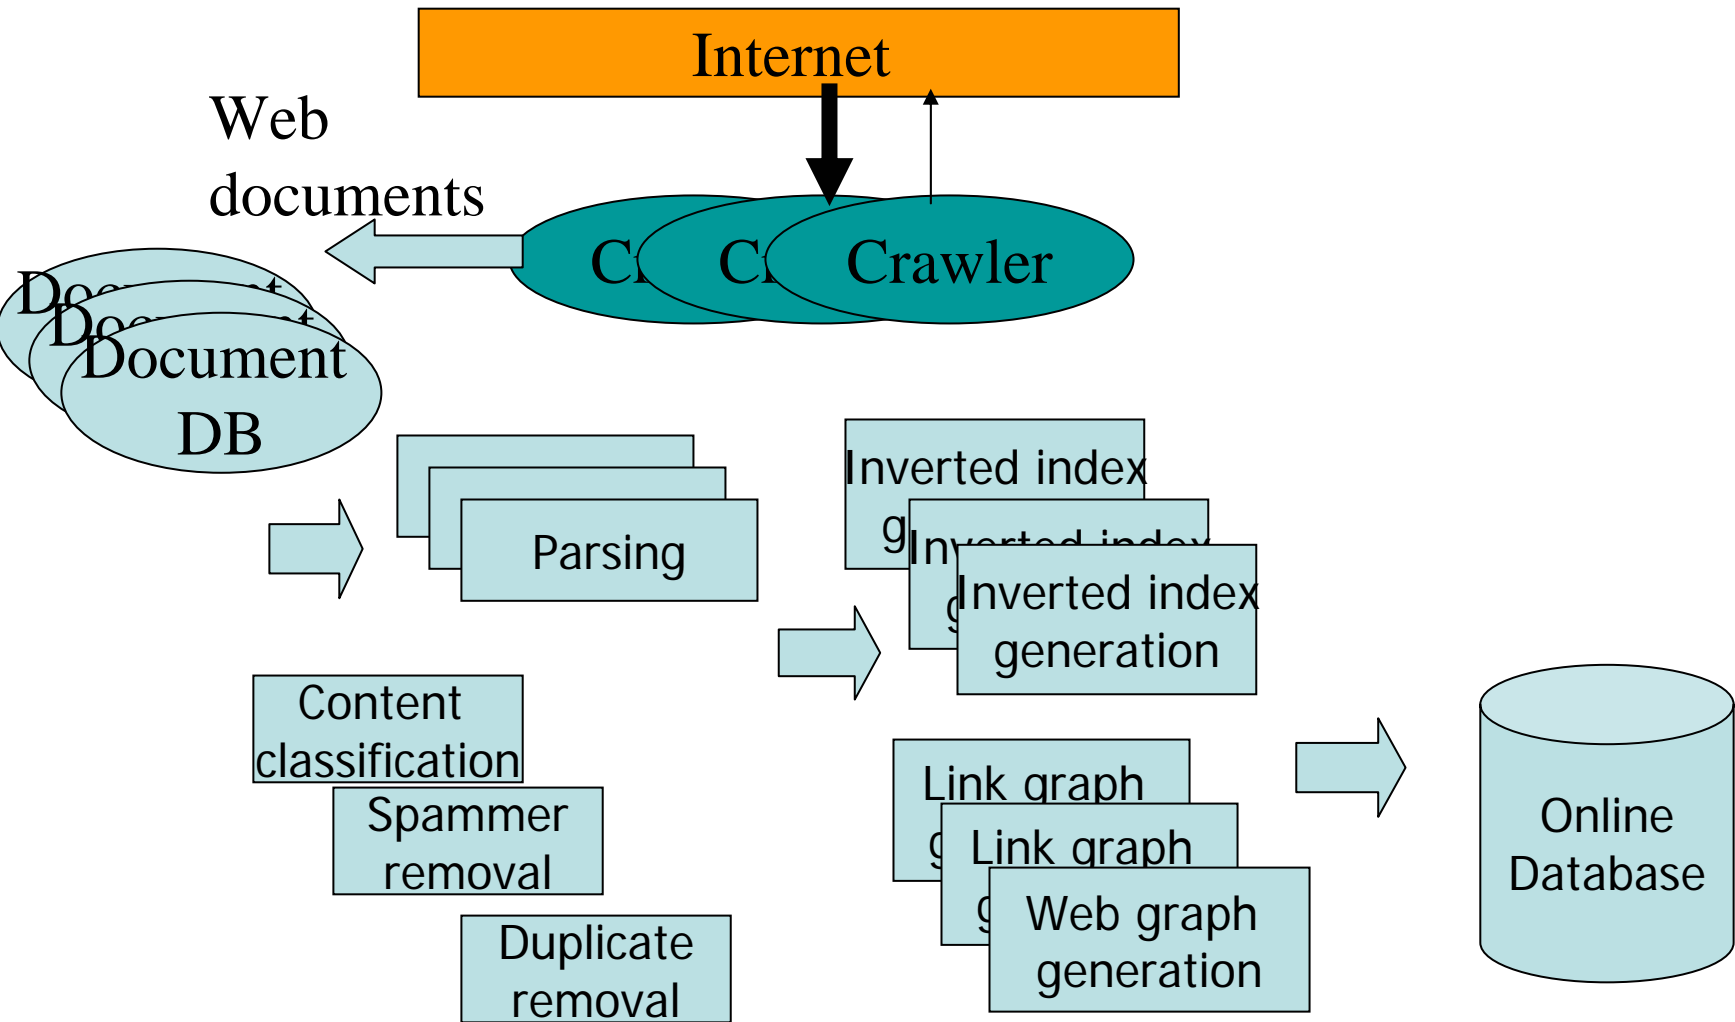
[...e, Lance Armstrong coaching log, training and](#)



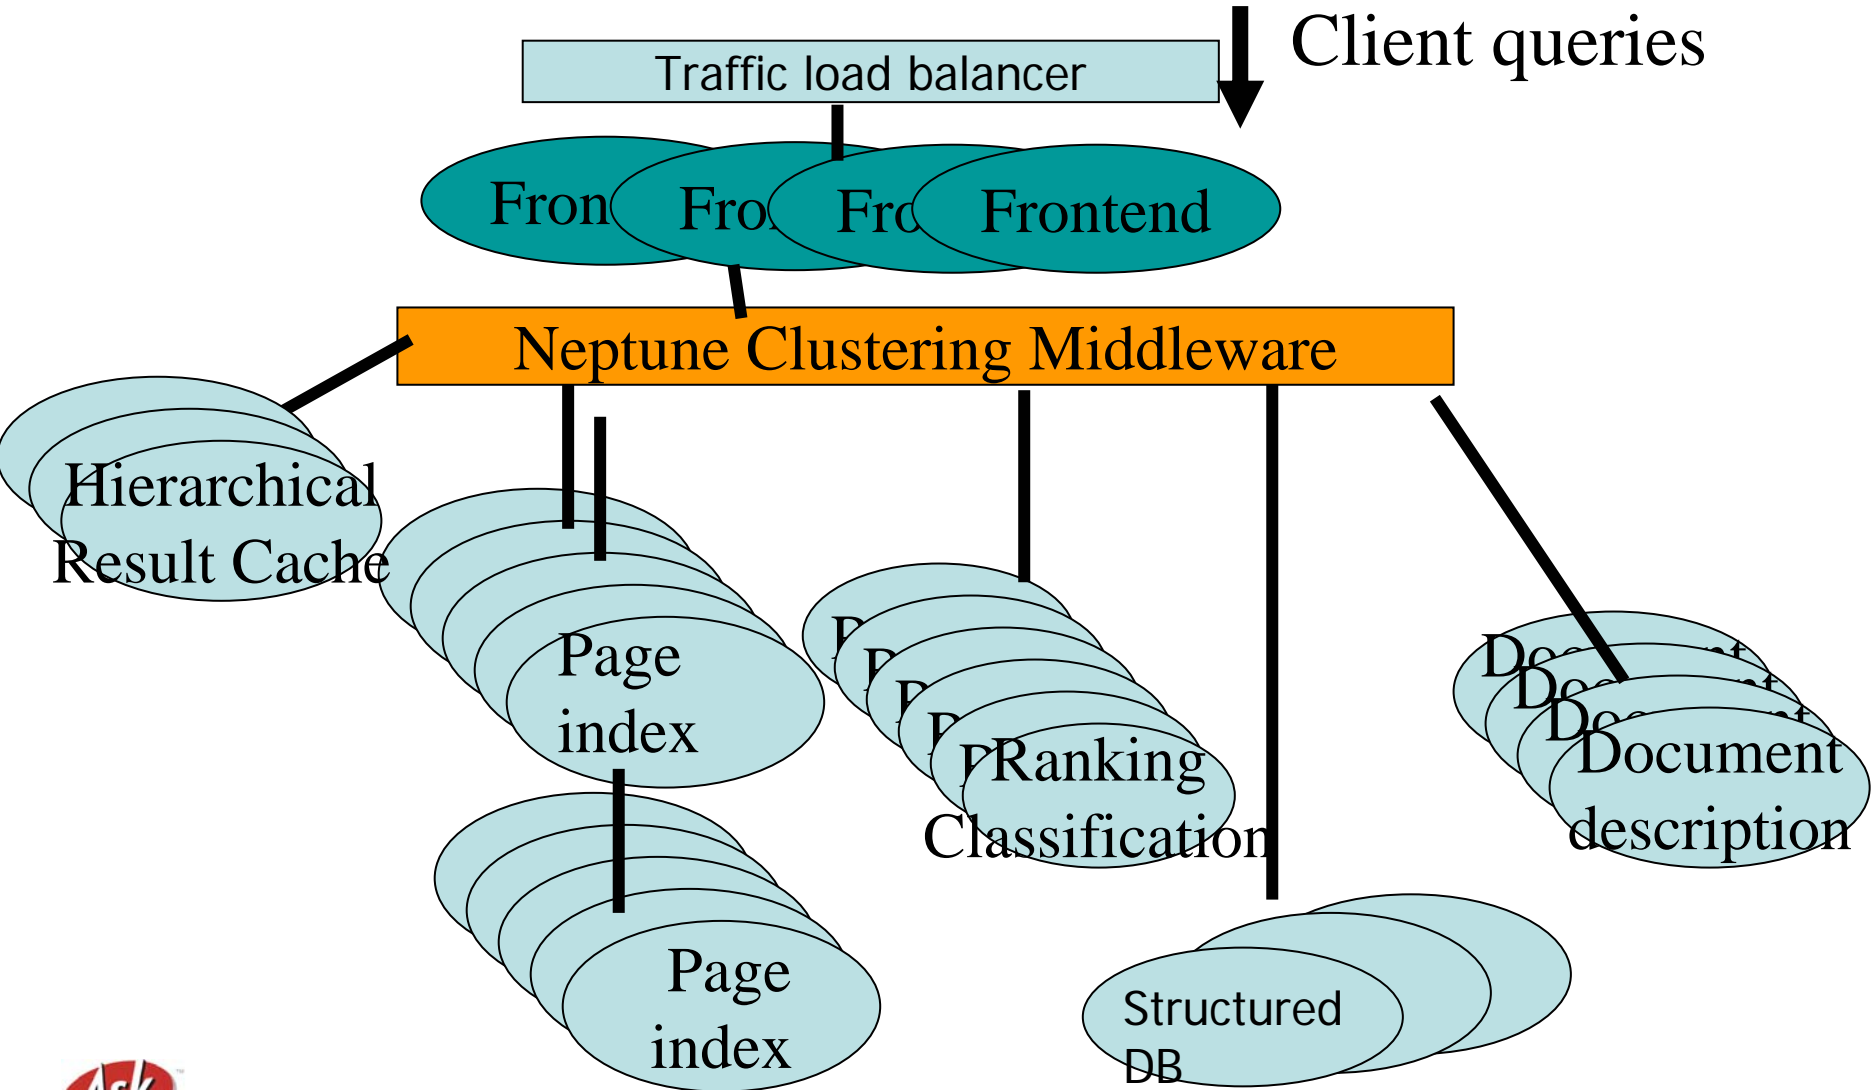
Ask Competitive Strengths

- Deeper topic view of the Internet
 - Query-specific link and text analysis with behavior analysis
 - Differentiated clustering technology
- Natural Language Processing
 - Better understanding/analysis of queries and user behavior
- Integration of structured data with web search.
 - Smart answers

Behind Ask.com: Data Indexing and Mining




Engine Architecture



Concept: Link-based Popularity for Ranking

- **A is a connectivity matrix among web pages. $A(i,j)=1$ for edge from i to j.**



	1	2	3	4	5	6	7	8	9
1									
2					1				1
3					1				1
4					1				
5									1
6									1
7						1			
8									1
9									

- **Query-independent popularity.**
- **Query-specific popularity**

Approaches for Page Ranking

- **PageRank:[Brin/Page'98]** offline computation of query-independent popularity iteratively.
- **HITS:[Kleinberg'98, IBM Clever]**
 - Build a query-based connectivity matrix on the fly. H, R are hub and authority weights of pages.
 - Repeat until H, R converge.
 - $R=A' H=A'A R$;
 - Normalize H, R.
- **ExpertRank:** Compute query-specific communities and ranking in real time.
 - Started from Teoma and evolved at Ask.com

Steps of ExpertRank at Ask.com


1

Search the index for a query



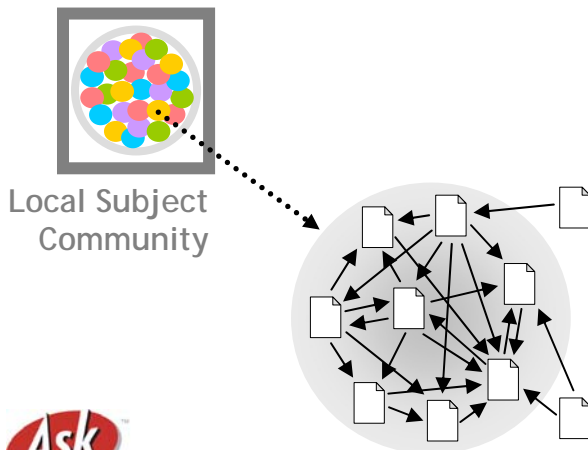
2

Clustering for subject communities for matched results




3

local subject-specific mining



4

Ranking with knowledge and classification



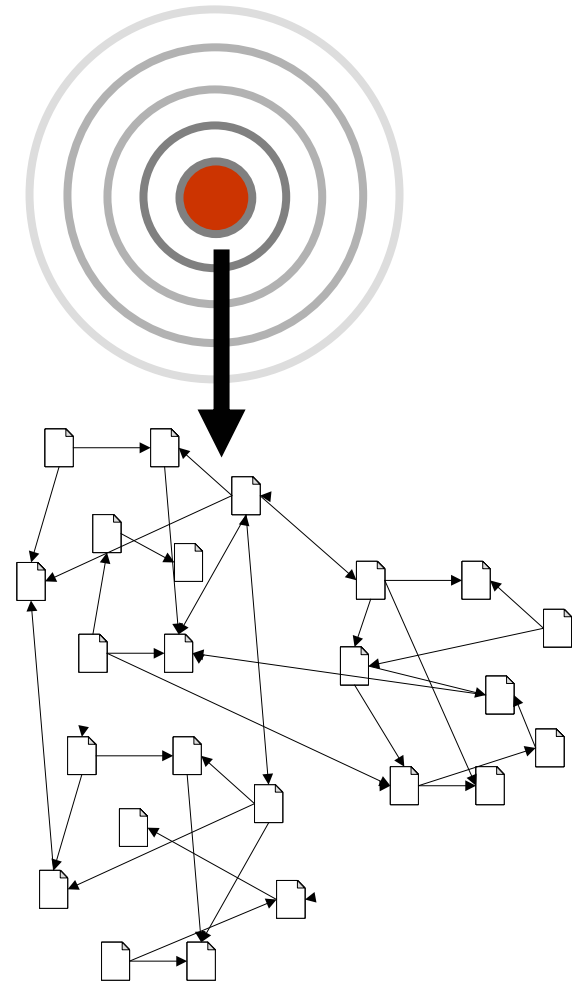
1

Index search and web graph generation

- **Search the index and identify relevant candidates for a given query.**

- Relevant pages, high quality pages, fresh pages.

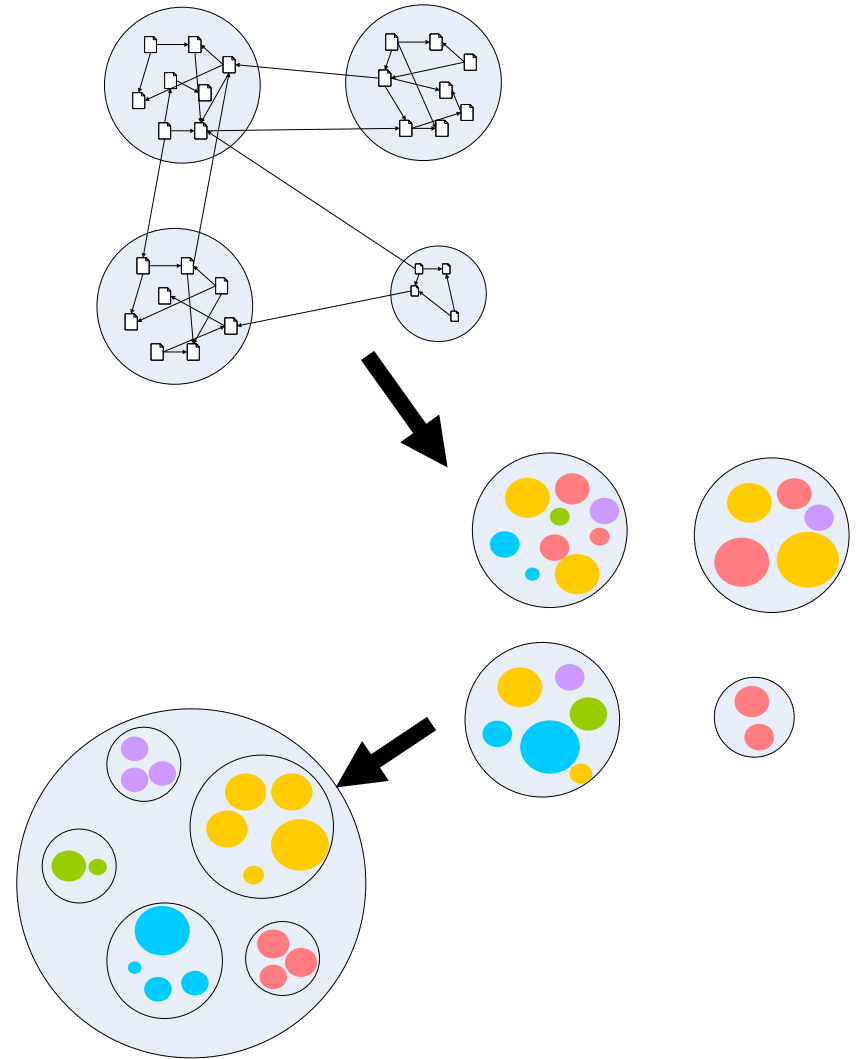
- **Generate a query-specific link graph dynamically.**



2

Multi-stage Cluster Refinement with Integrated Link/Topic Analysis

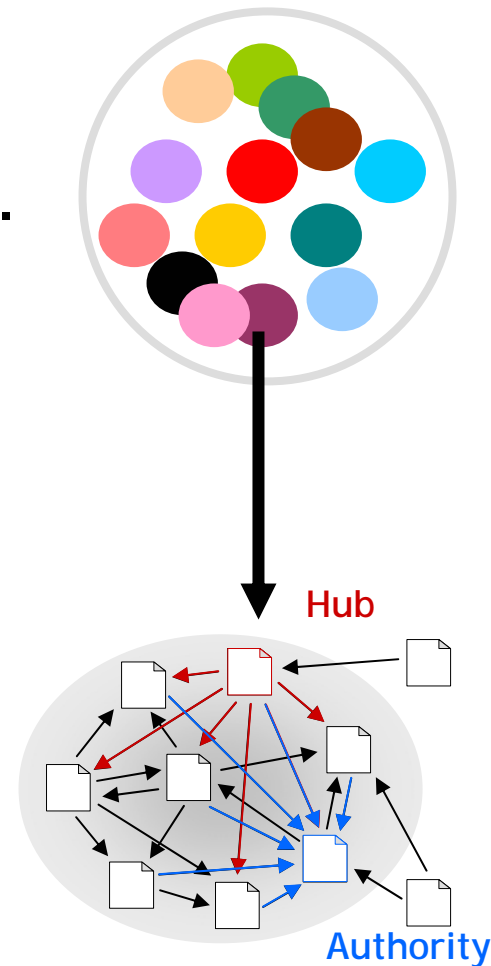
- **Link-guided page clustering**
- **Cluster refinement with content analysis and topic purification**
 - Text classification and NLP
 - Similarity and overlapping analysis



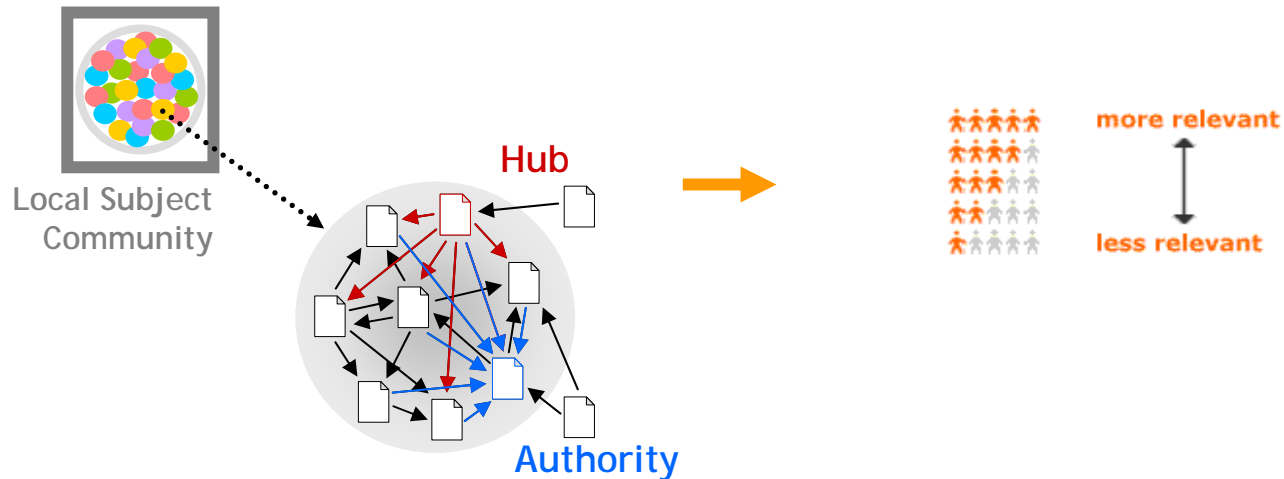
3

Subject-specific ranking

- **Example**
 - “bat”, flying mammals vs. baseball bat.
- **For each topic group, identify experts for page recommendation, and remove spamming links.**
- **Derive local ranking scores**



- **Score weighting from multiple topic groups.**
 - Authoritativeness and freshness assessment.
 - User intention analysis.
 - Result diversification.



Scalability Challenges

- **Data scalability:**
 - From millions of pages to billions of pages.
 - Clean vs. datasets with lots of noise.
- **Infrastructure scalability:**
 - Tens of thousands of machines.
 - Tens of Millions of users
 - Impact on response time, throughput, & availability,
 - data center power/space/networking.
- **People scalability:** From few persons to many engineers with non-uniform experience.

Downtime Costs (per Hour)

• Brokerage operations	\$6,450,000
• Credit card authorization	\$2,600,000
• Ebay (1 outage 22 hours)	\$225,000
• Amazon.com	\$180,000
• Package shipping services	\$150,000
• Home shopping channel	\$113,000
• Catalog sales center	\$90,000
• Airline reservation center	\$89,000
• Cellular service activation	\$41,000
• On-line network fees	\$25,000
• ATM service fees	\$14,000

Source: InternetWeek 4/3/2000 + *Fibre Channel: A Comprehensive Introduction*, R. Kembel 2000, p.8. "...based on a survey done by Contingency Planning Research."

Examples of Scalability Problems

- Mining question answers from web.
- Large-scale spammer detection.
- Computing with irregular data. On-chip cache.
- Large-scale memory management: 32 bits vs. 64 bits.
- Incremental cluster expansion and topology mgmt.
- High throughput write/read traffic. Reliability.
- Fast and reliable data propagation across networks.
- Architecture optimization for low power consumption.
- Update large software & data on a live platform.
- Distributed debugging thousands of machines.

Some of Lessons Learned

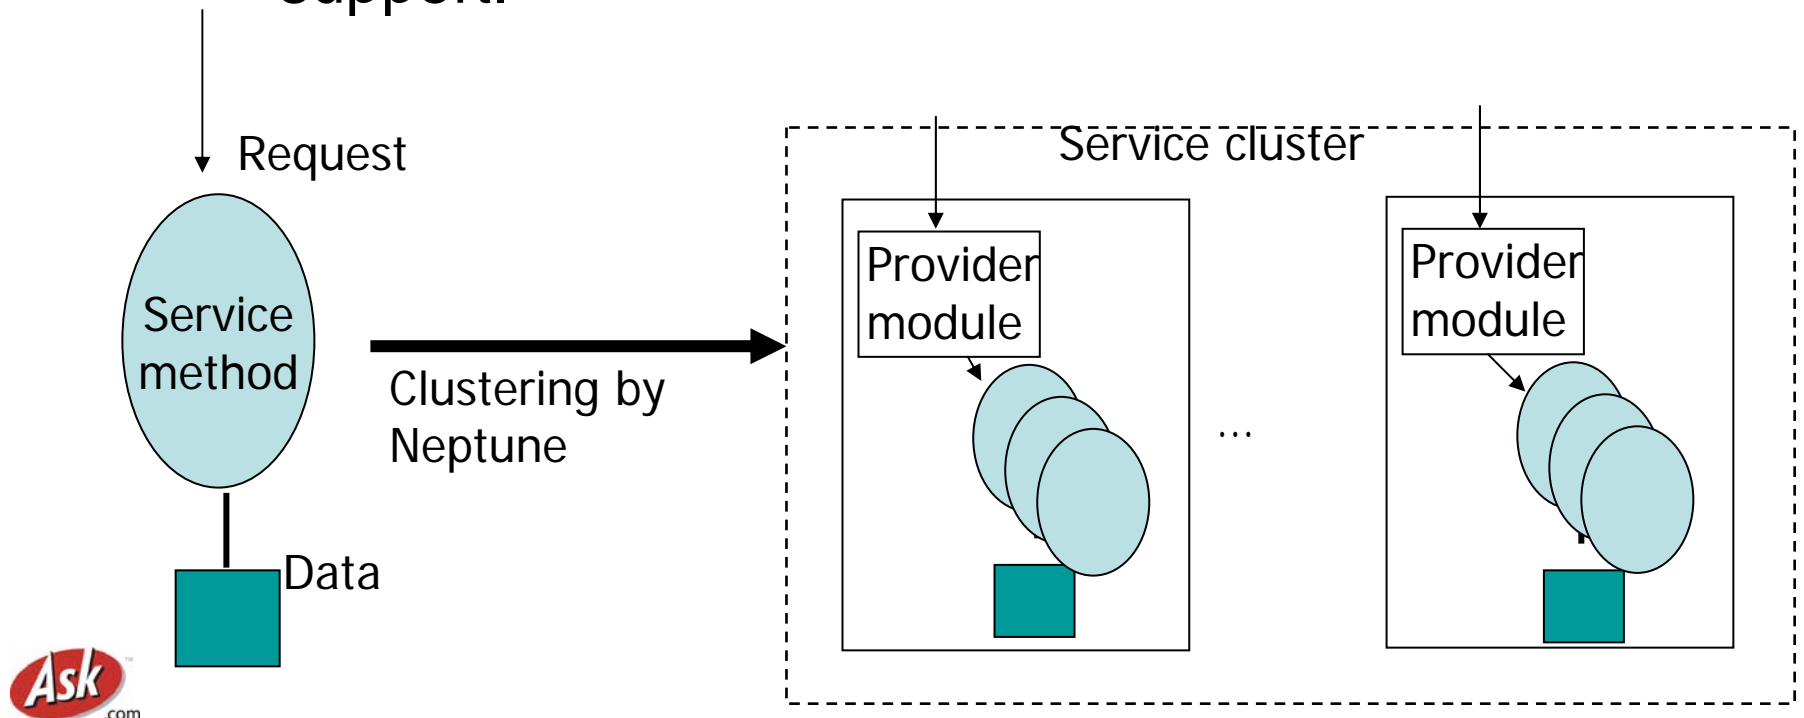
- **Data**
 - Data methods can behave differently with different data sizes/noise levels.
 - Data-driven approaches with iterative refinement.
- **Architecture & Software**
 - Distributed service-oriented architectures
 - Middleware support.
- **Product:**
 - Monitoring is as critical as others.
 - Simplicity

The Neptune Clustering Middleware

- **A simple/flexible programming model**
 - Aggregating and replicating application modules with persistent data.
 - Shielding complexity of service discovery, load balancing, consistency, and failover management
 - Providing inter-service communication.
 - Providing quality-aware request scheduling for service differentiation
- **Started at UCSB. Evolved with Teoma, Ask.com.**


Programming Model and Cluster-level Parallelism/Redudancy in Neptune

- Request-driven processing model.
- SPMD model (single program/multiple data) while large data sets are partitioned and replicated.
- Location-transparent service access with consistency support.

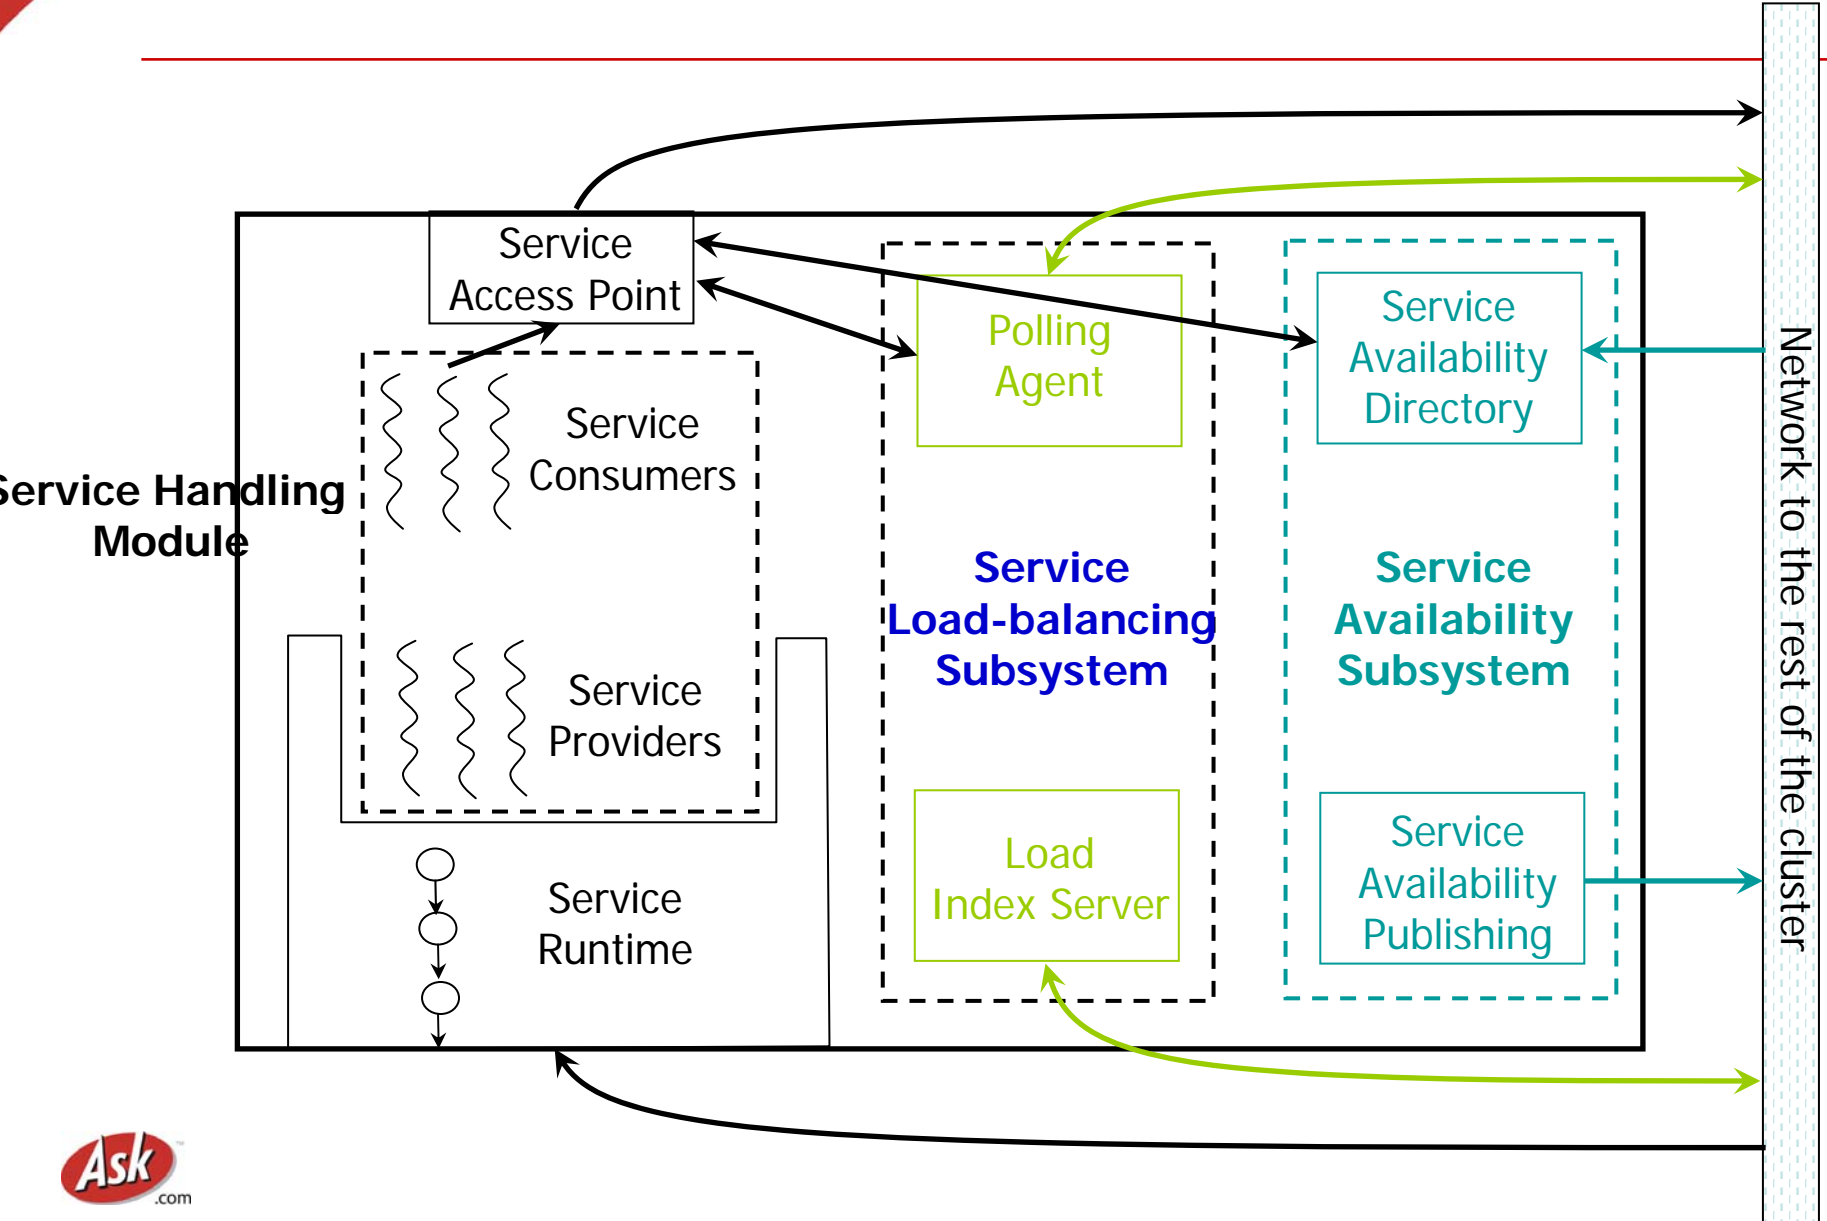


Neptune architecture for cluster-based services

- Symmetric and decentralized:
 - Each node can host multiple services, acting as a service provider.
 - Each node can also subscribe internal services from other nodes, acting as a consumer.
 - *Support multi-tier or nested service architecture*

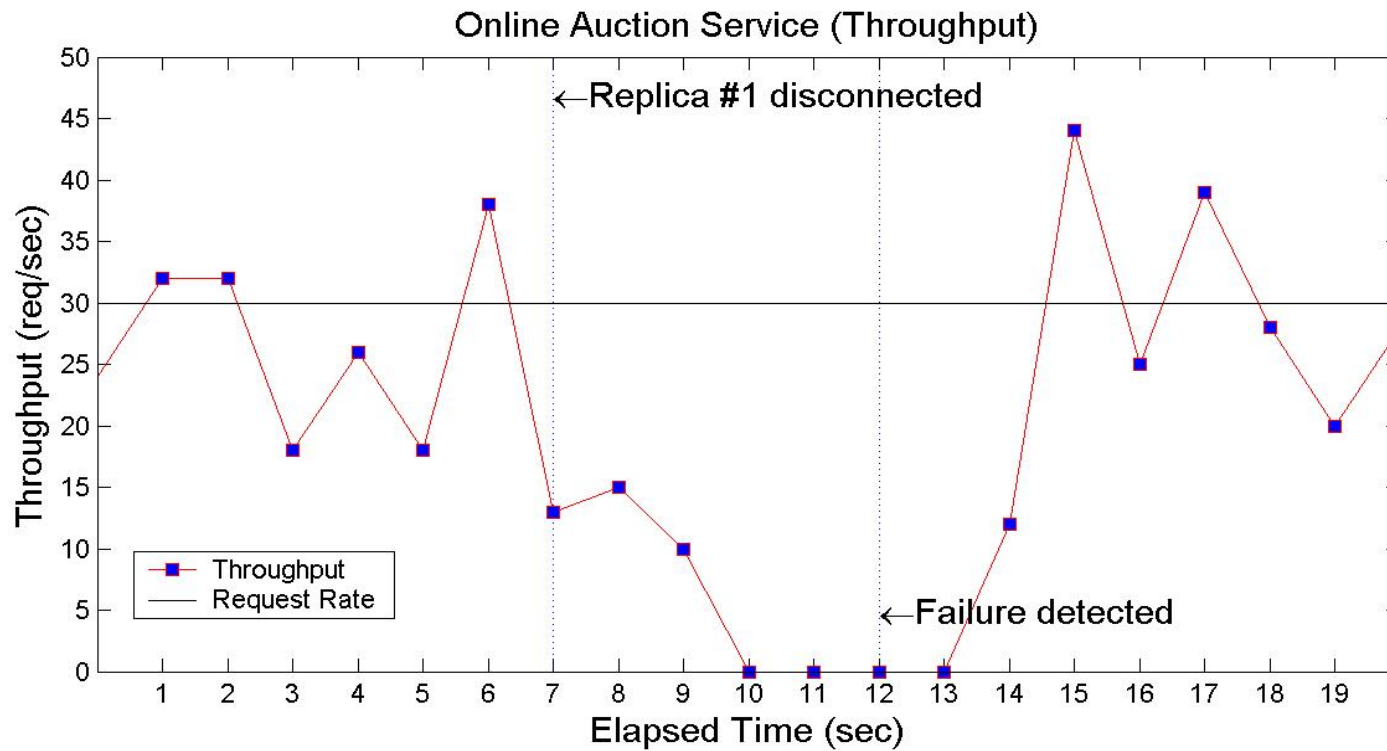
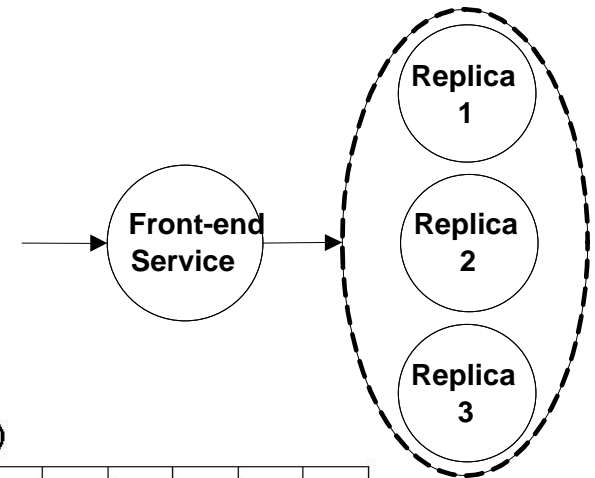


Inside a Neptune Server Node



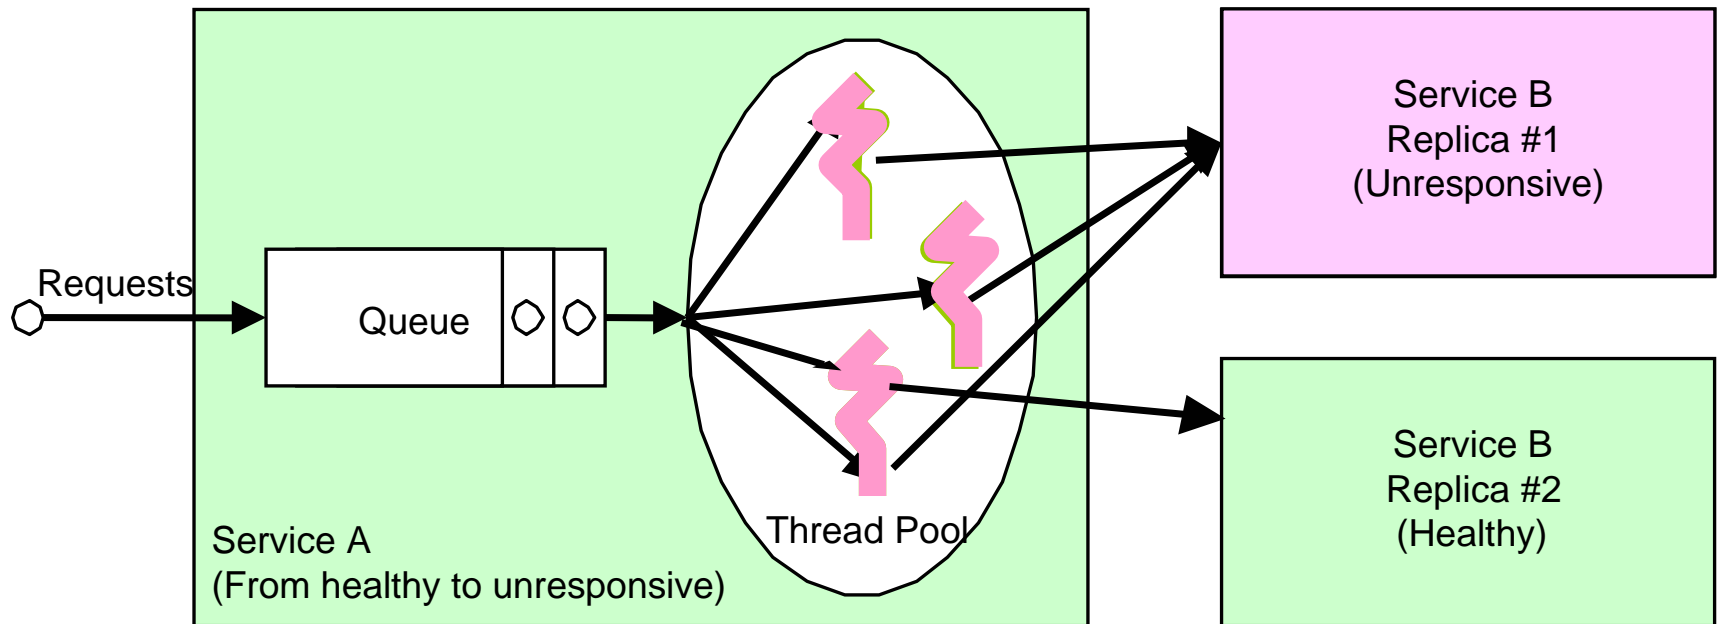
Impact of Component Failure in Multi-tier services

- Failure of one replica: 7s - 12s
- Service unavailable: 10s - 13s



Problems that affect availability

- Threads are blocked with slow service dependency.
- Fault detection speed.




Dependency Isolation

- **Per-dependency management with capsules.**



- Isolate their performance impact.
- maintain dependency-specific feedback information for QoS control.

- **Programming support with automatic recognition of dependency states.**



Fast Fault Detection and Information Propagation for Large-Scale Cluster-Based Services

- **Complex 24x7 network topology in service clusters.**
- **Frequent events: failures, structure changes, and new services.**
 - Yellowpage directory
 - discovery of services and their attributes
 - Server aliveness




TAMP: Topology-Adaptive Membership Protocol

- **Highly Efficient: Optimize bandwidth, # of packets**
- **Topology-aware:**
 - Form a hierarchical tree according to network topology
 - Localize traffic within switches and adaptive to changes of switch architecture.
- **Topology-adaptive:**
 - Network changes: switches
- **Scalable: scale to tens of thousands of nodes. Easy to operate.**

Hierarchical Tree Formation Algorithm

- **Exploiting TTL count in IP packet for topology-adaptive design.**
- **Each multicast group with a fixed TTL value performs an election;**
- **Group leaders form higher level groups with larger TTL values;**
- **Stop when max. TTL value is reached; otherwise, goto *Step 2*.**

An Example of Hierarchical Tree Formation



Scalability Analysis

- **Basic performance factors**
 - Failure detection time (T_{fail_detect})
 - View convergence time ($T_{converge}$)
 - Communication cost in terms of bandwidth (B)
- **Two metrics**
 - $BDP = B * T_{fail_detect}$, lower failure detection time with low bandwidth is desired
 - $BCP = B * T_{converge}$, lower convergence time with low bandwidth is desired

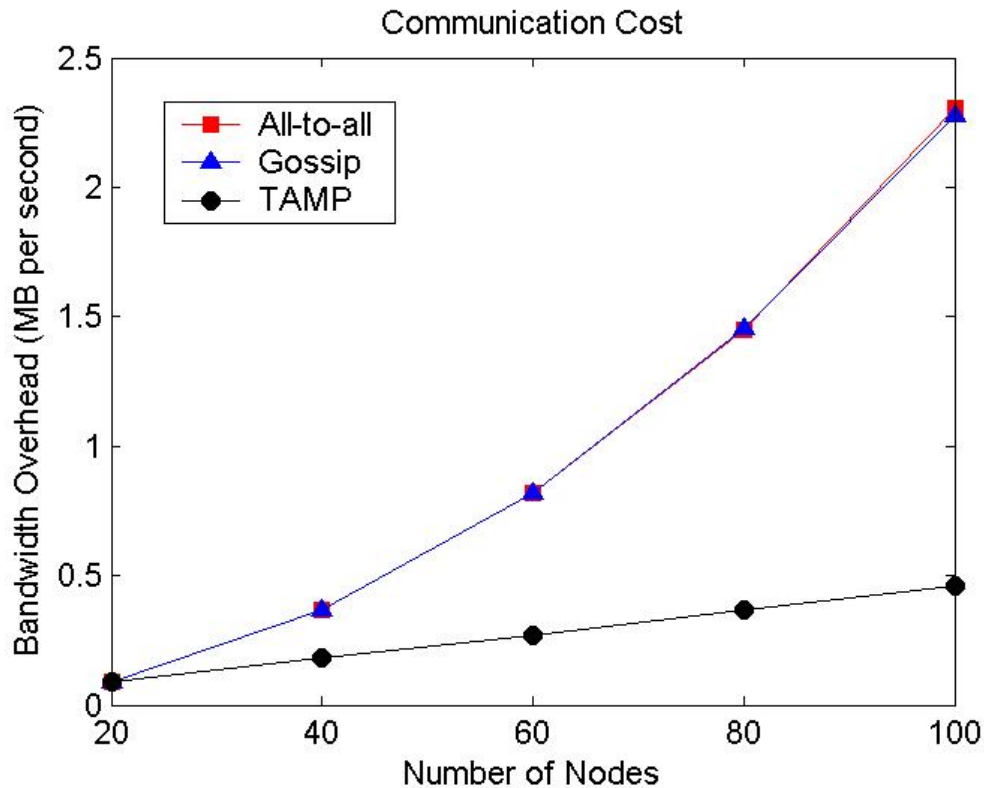
A scalability comparison of three methods

	<i>Failure Detection Time x Bandwidth</i>	<i>Convergence Time x Bandwidth required</i>
All-to-all	$O(n^2)$	$O(n^2)$
Gossip	$O(n^2 \log n)$	$O(n^2 \log n)$
TAMP	$O(n)$	$O(n \log_k n)$

n: total # of nodes

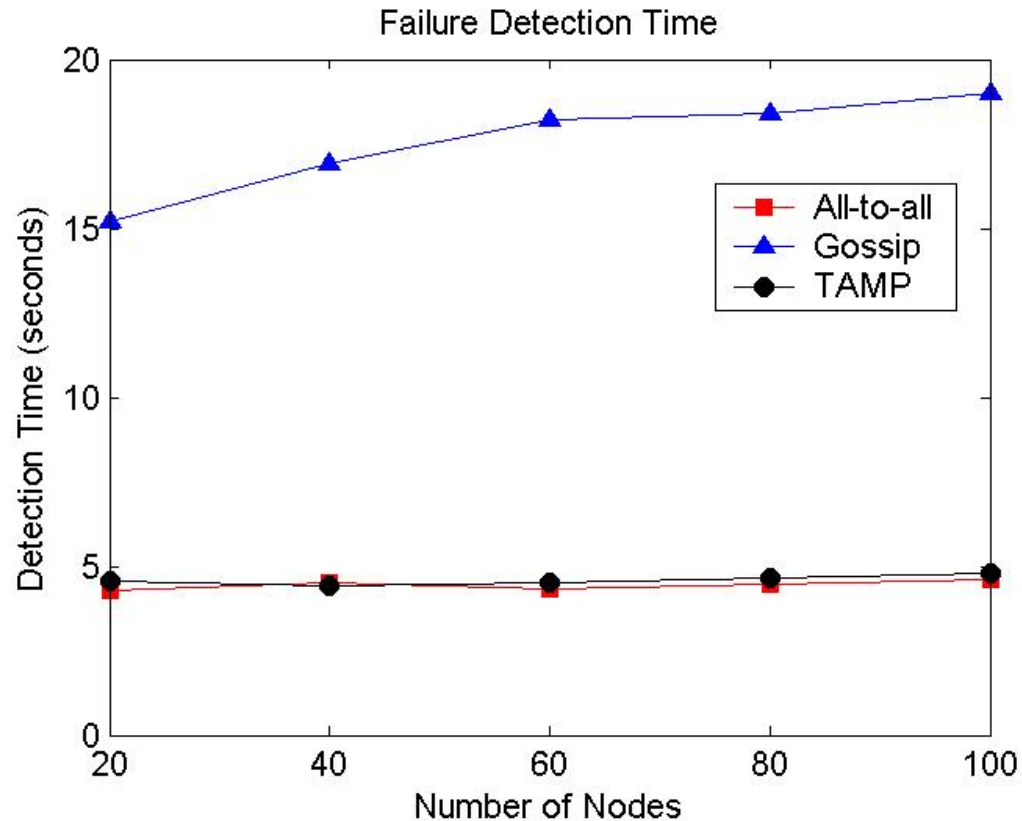
k: each group size, a constant

Bandwidth Consumption



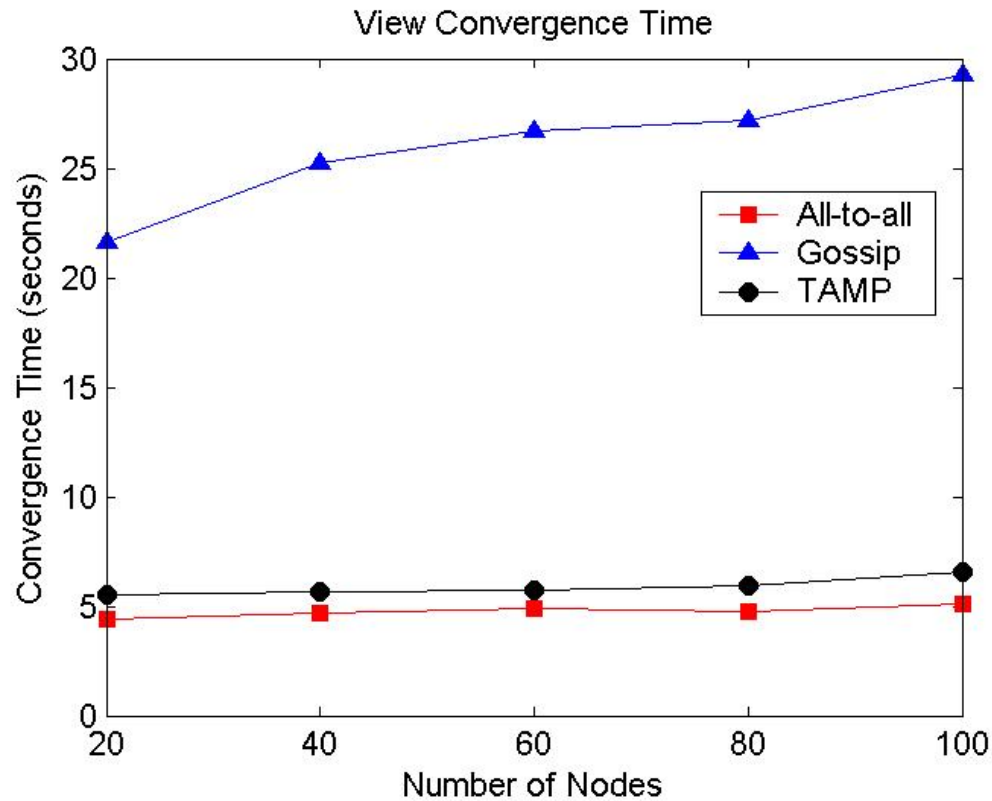
- **All-to-All & Gossip: quadratic increase**
- **TAMP: close to linear**

Failure Detection Time



- **Gossip: $\log(N)$ increase**
- **All-to-All & TAMP: constant**

View Convergence Time



- **Gossip: $\log(N)$ increase**
- **All-to-All & TAMP: constant**

References

- T. Yang, W. Wang, A. Gerasoulis, Relevancy-Based Database Retrieval and Display Techniques, Ask Jeeves/Teoma, 2002. US Patent 7028026.
- K. Shen, H. Tang, T. Yang, and L. Chu, Integrated Resource Management for Cluster-based Internet Services. In *Proc. of Fifth USENIX Sym. on Operating Systems Design and Implementation (OSDI '02)* , pp 225-238, Boston, 2002.
- L. Chu, T. Yang, J. Zhou, Topology-Centric Resource Management for Large Scale Service Clusters, 2005 (Pending patent application).
- L. Chu, K. Shen, H.Tang, T. Yang, and J. Zhou. Dependency Isolation for Thread-based Multi-tier Internet Services. In *Proc. of IEEE INFOCOM 2005*, Miami FL, March, 2005

Concluding Remarks

- Ask.com is focused on leading-edge technology for Internet search.
- Many open/challenging problems for information retrieval, mining, and system scalability.
- Interested in joining Ask.com?
recruiting@ask.com