Massive Data Analysis: What is under the hood?

S. (Muthu) Muthukrishnan

Google

mysliceofpizza

Talk Overview

- Data Analysis in Different Communities
  - Algorithms, Databases and Networking
- Infrastructure View of Data Analysis
  - Example 1: Cellphone Call Traffic
  - Example 2: IP Packet Traffic Streams
  - Example 3: Web Traffic
- Perspectives

Data Analysis in Different Communities

- Networking:
  - Mining anomalies using traffic feature distributions
    A. Lakhina, M. Crovella, C. Diot. SIGCOMM 05.

entropy

- Algorithms:
  - Streaming and sublinear approximation of entropy and information distances.

    S. Guha, A. McGregor, S. Venkatasubramanian. SODA 2006.
- Databases:
  - Holistic UDAFs at streaming speeds.

    G. Cormode, T. Johnson, F. Korn, S. Muthukrishnan, O. Spatscheck, D. Srivastava. SIGMOD 2004.

User defined aggregate function (UDAF), eg., entropy.

Infrastructure View, Example 1:
Cellphone Calls Analysis

A mobile call: Detailed view of CDRs

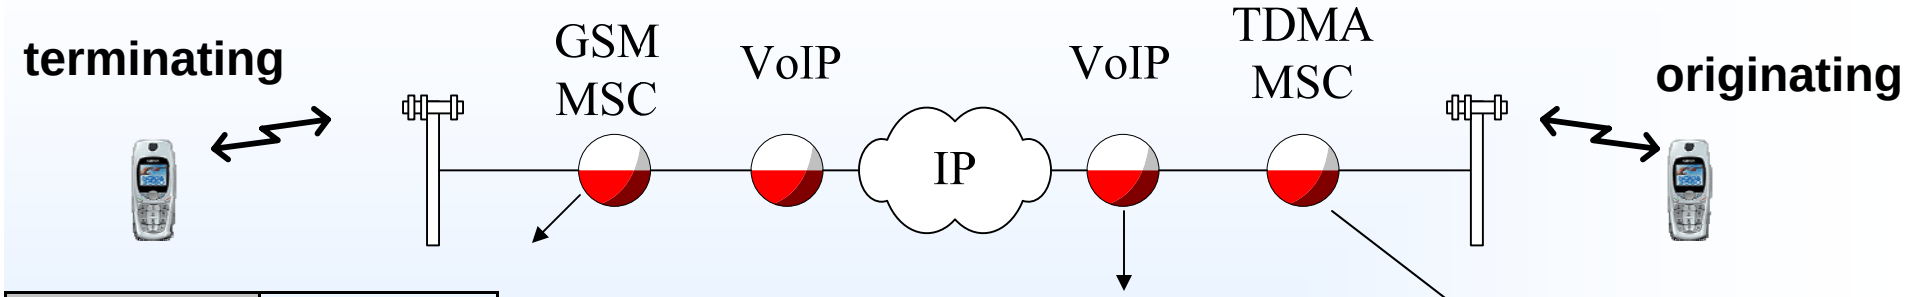

rel	gsm17
scode	3
IMSI	310380049259999
Calling Number	2136109999
Called Number	19493009999
Dialed Digits	
IMEI	352968001799999
Channel Alloc Time	6/26/05 7:28:00
Answer Time	6/26/05 7:28:02
Disconnect Time	6/26/05 7:28:10
Rls Time	6/26/05 7:28:10
Half Rate	0
termcause	004
diag	04127
in adnum	00204
in memkey	00330
out adnum	
out memkey	
in trk seize	6/26/05 7:27:57
out trk seize	
calldur	0000009
BSC in adnum	00520
BSC in memkey	00740
LAC	31038005221
CellID	59165
ChanType	11140
LRN	

Gateway	ANHG2SO
StartTime	6/26/05 7:28:16
Disc_Time	6/26/05 7:28:29
Duration	789
Diag	127
Service	VoIP
ASubNum	2136109999
BSubNum	9516425189 (msrn)
BillNum	9493009999
RouteLabel	RVSDCALBCM5_IM
RouteSelected (Gateway:CLLI)	RVSG5SO:RVSDCALBCM50IMB
LocSIPaddr	155.172.0.9
RemSIPaddr	155.172.0.216
InPSTN_TrkNm	ANHMCACLCM30IMB
InPSTN_CircEnd	1:14:12:7:1079:0x00E37D01:0x00E3C6F2
EgrIP_CircEnd	155.172.0.11:8050/155.172.0.218:8728
PktsOut	620
PktsIn	617
GSX Call Handle	GSX2GSX,0x380D6441
DialedNum	9494661933 (lrn)
GenAddr	9493009999
InCodec	C:1:1
OutCodec	P:1:1
OrigEchCanc	1

Record_type	04
Call_status	2
Call_ID_number	01586580
A_subscriber_number	2136109999
B_subscriber_number	9493009999
Date_for_start_of_charging	6/26/05 7:29:00
Chargeable_duration	7
Time regsz	5
Abnormal_call_release	1
Internal_Cause_and_Location	027B
Outgoing_route	AN2AMGO
Incoming_route	C736CKI

"Transmission fault, incoming" (dropped call)

Analyzing CDRs: Data

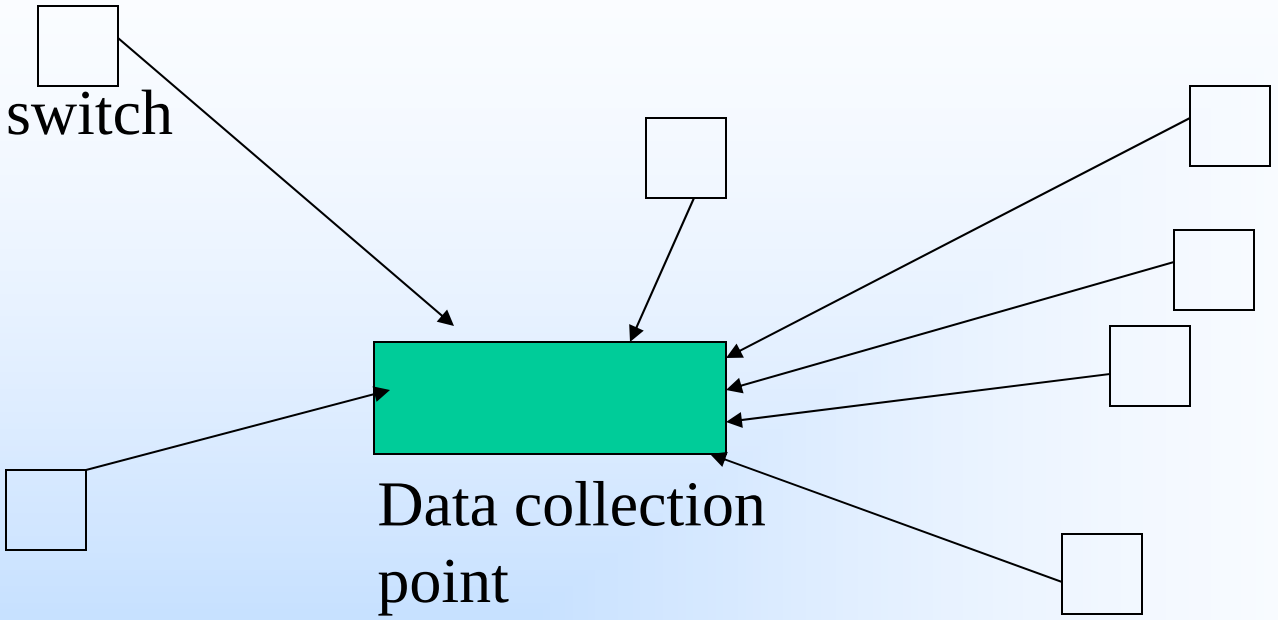

- Data:
  - TDMA: Ericsson, Lucent, and Nortel MSCs; GSM and UMTS: Nortel MSCs; VoIP: Sonus Media Gateways; GPRS: Nortel SGSNs, GGSNs, and MMSCs; SMS logs.
  - 20 - 30 different data formats.
  - Side tables: LERG. Handset info. Trunk info.
  - About 1 Tbyte/month.

Analyzing CDRs: Analyses

- Analyses:
  - 100's of reports a month.
- Example Analyses:
  - Dropped calls per handset type
  - Glare detection
  - 2A or 2B connections.
  - Fraudulent transit calls
  - Cell adjacency graph

Example Analysis:
Distant Tower Problem

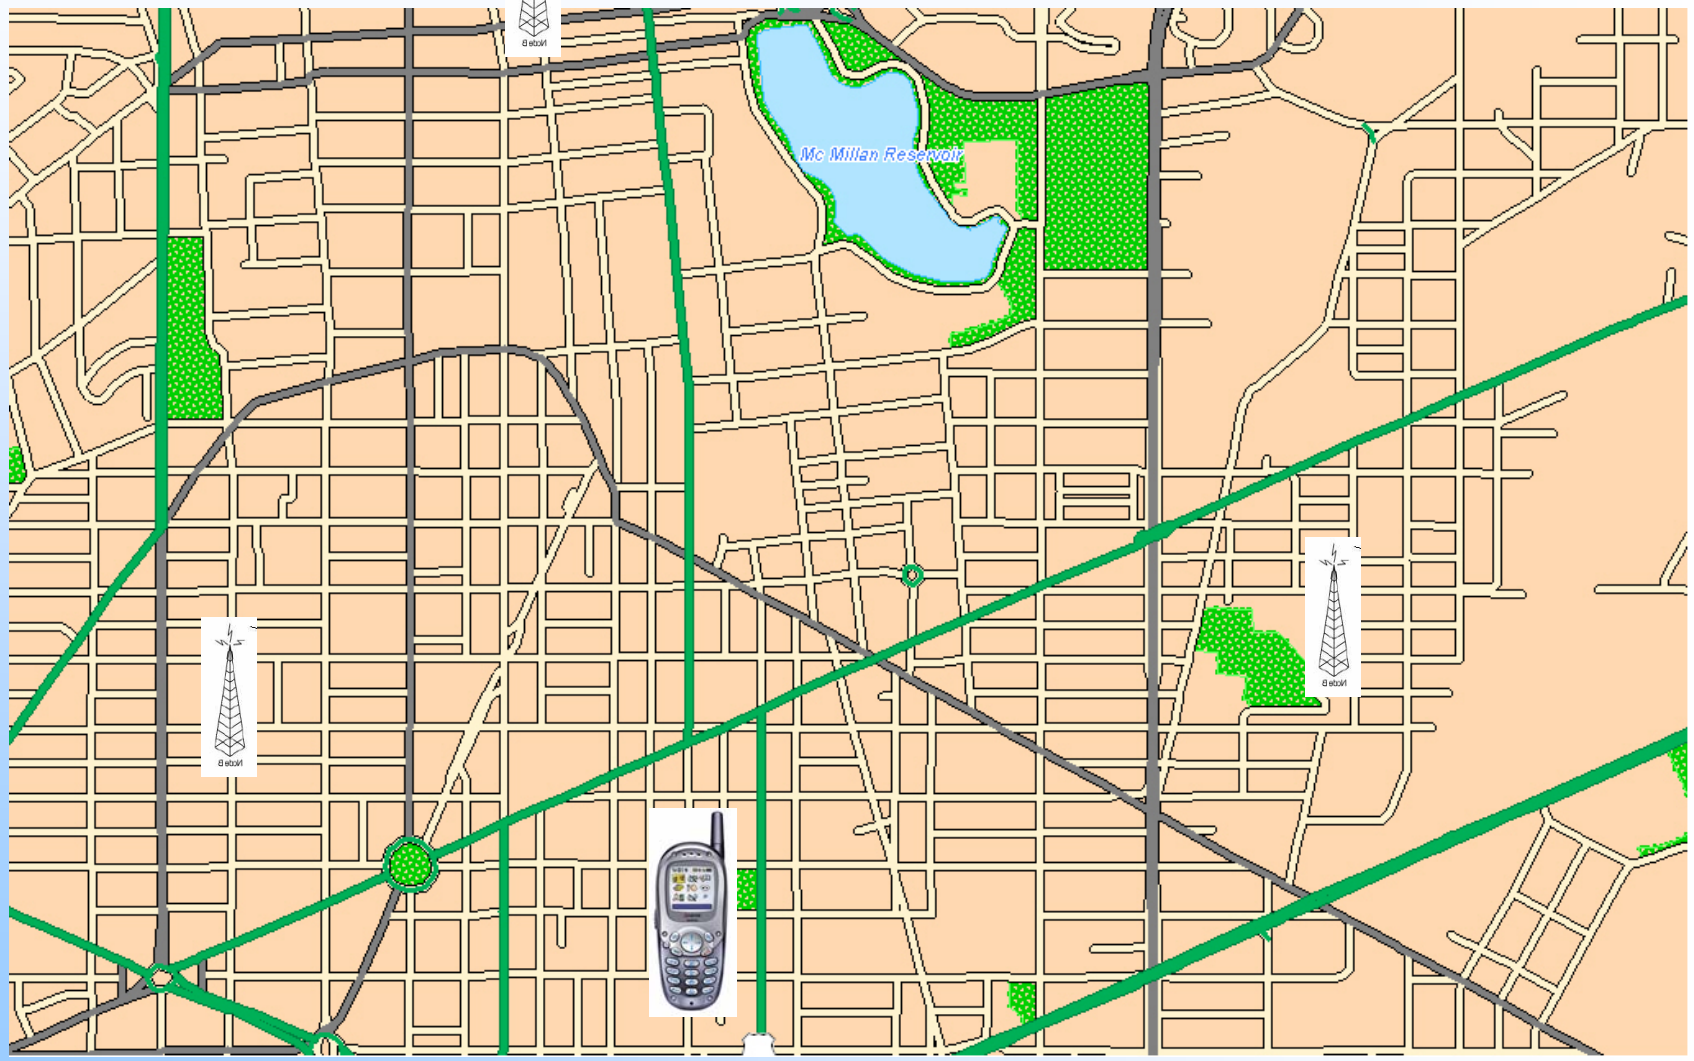

Distant Tower Problem

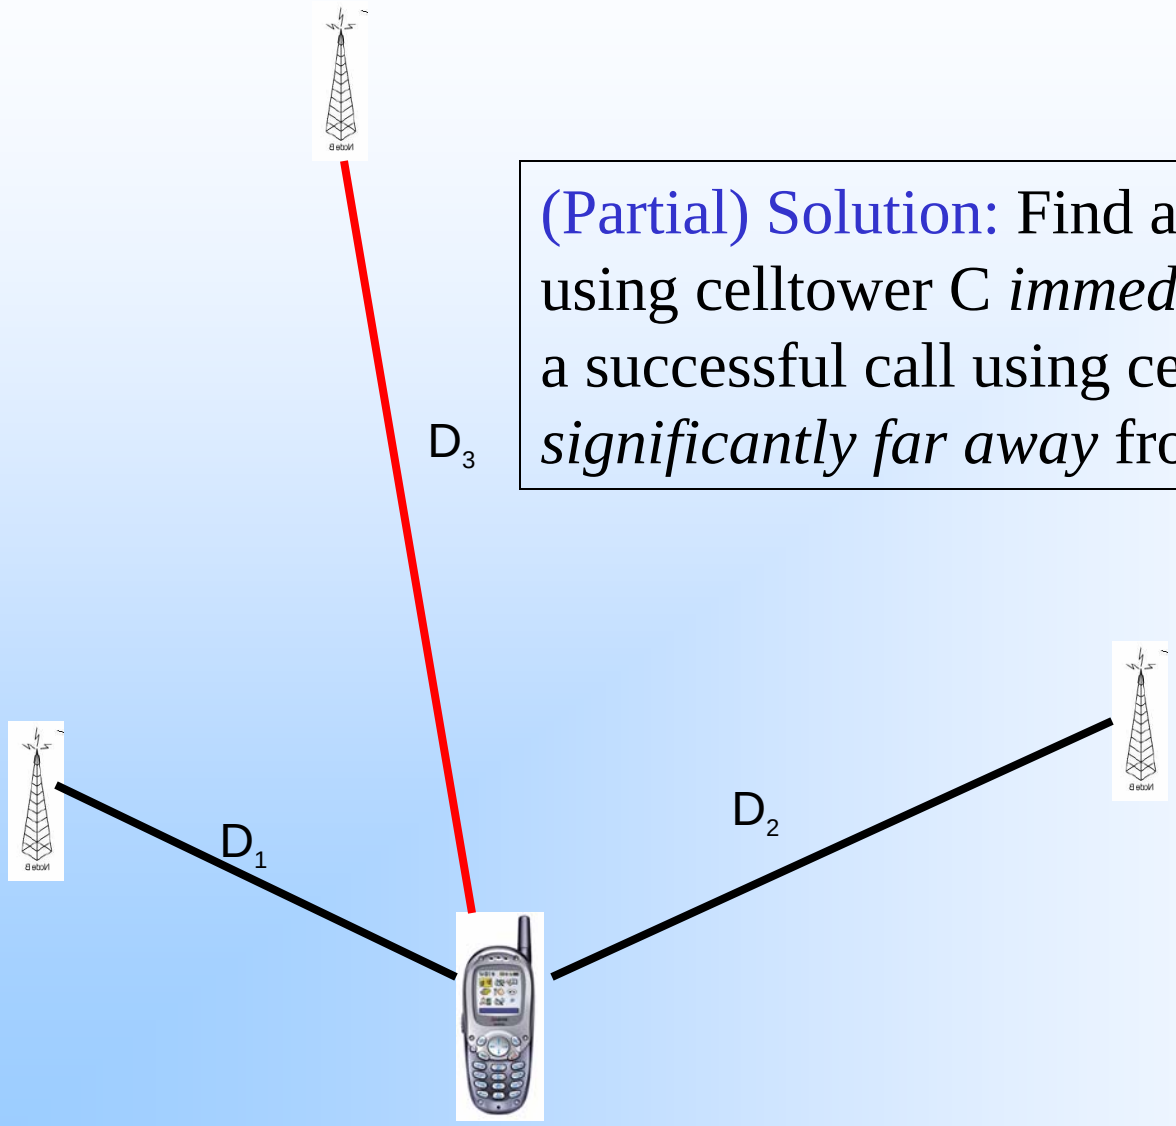

(Partial) Solution: Find a dropped call using celltower C *immediately preceding* a successful call using celltower D *significantly far away* from C.

Analyzing CDRs: Infrastructure

- Challenge is not the size of the data.
  - understanding the data, translating a business problem down to CDR analysis.
- Turnaround time: Days or weeks.
- Small team of analysts responsible.

Infrastructure:

- Large disks.
- Multiple CPU machines.
- Scripting languages, standard file system.

Talk Overview

- Data Analysis in Different Communities
  - Algorithms, Databases and Networking
- Infrastructure View of Data Analysis
  - Example 1: Cellphone Call Traffic
  - Example 2: IP Packet Traffic Streams 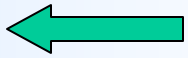
  - Example 3: Web Traffic
- Perspectives

Infrastructure View:
IP Traffic Analysis

Analyzing IP Traffic (ISP View): Data

- SNMP, IP flows, packet header logs, packet contents, routing tables, BGP updates, fault alarms.
- OC48, 192, 768: xTbytes/hour. 6M -- 96M pkts/sec.
- Real time, router speed analysis.
- Example:
  - Reporting, SLA mediation.
  - Anomaly/Attack detection.
  - Lawful intercept
  - Monitoring failures.
  - Traffic classification.

Gigascope Architecture

- Gigascope is an SQL-based operational IP traffic analysis tool at AT&T.
- Has two level arch.
  - Low-level queries perform initial fast selection and aggregation on high speed stream.
  - Complex aggregation on high level, at monitor server
- Depending on the capabilities of the NIC, can push operators and low-level queries into it.

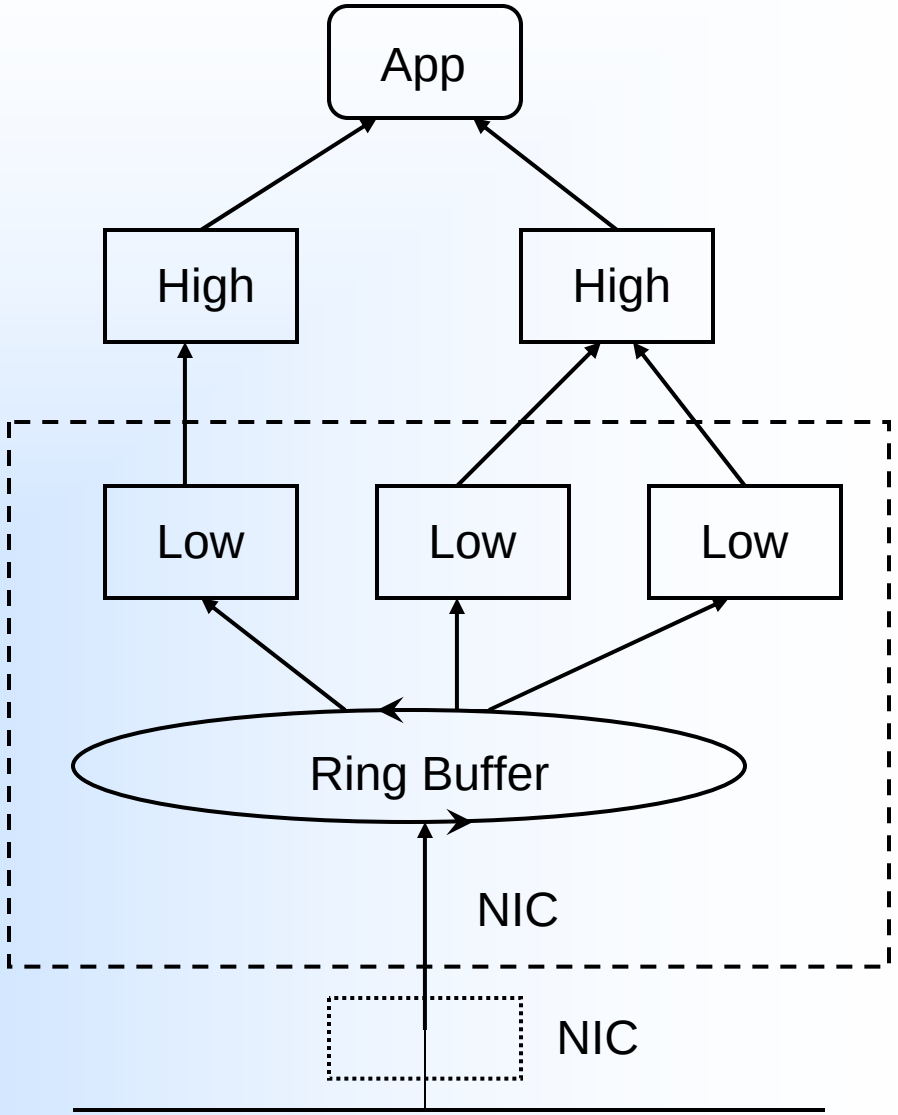

GSQL Query Splitting

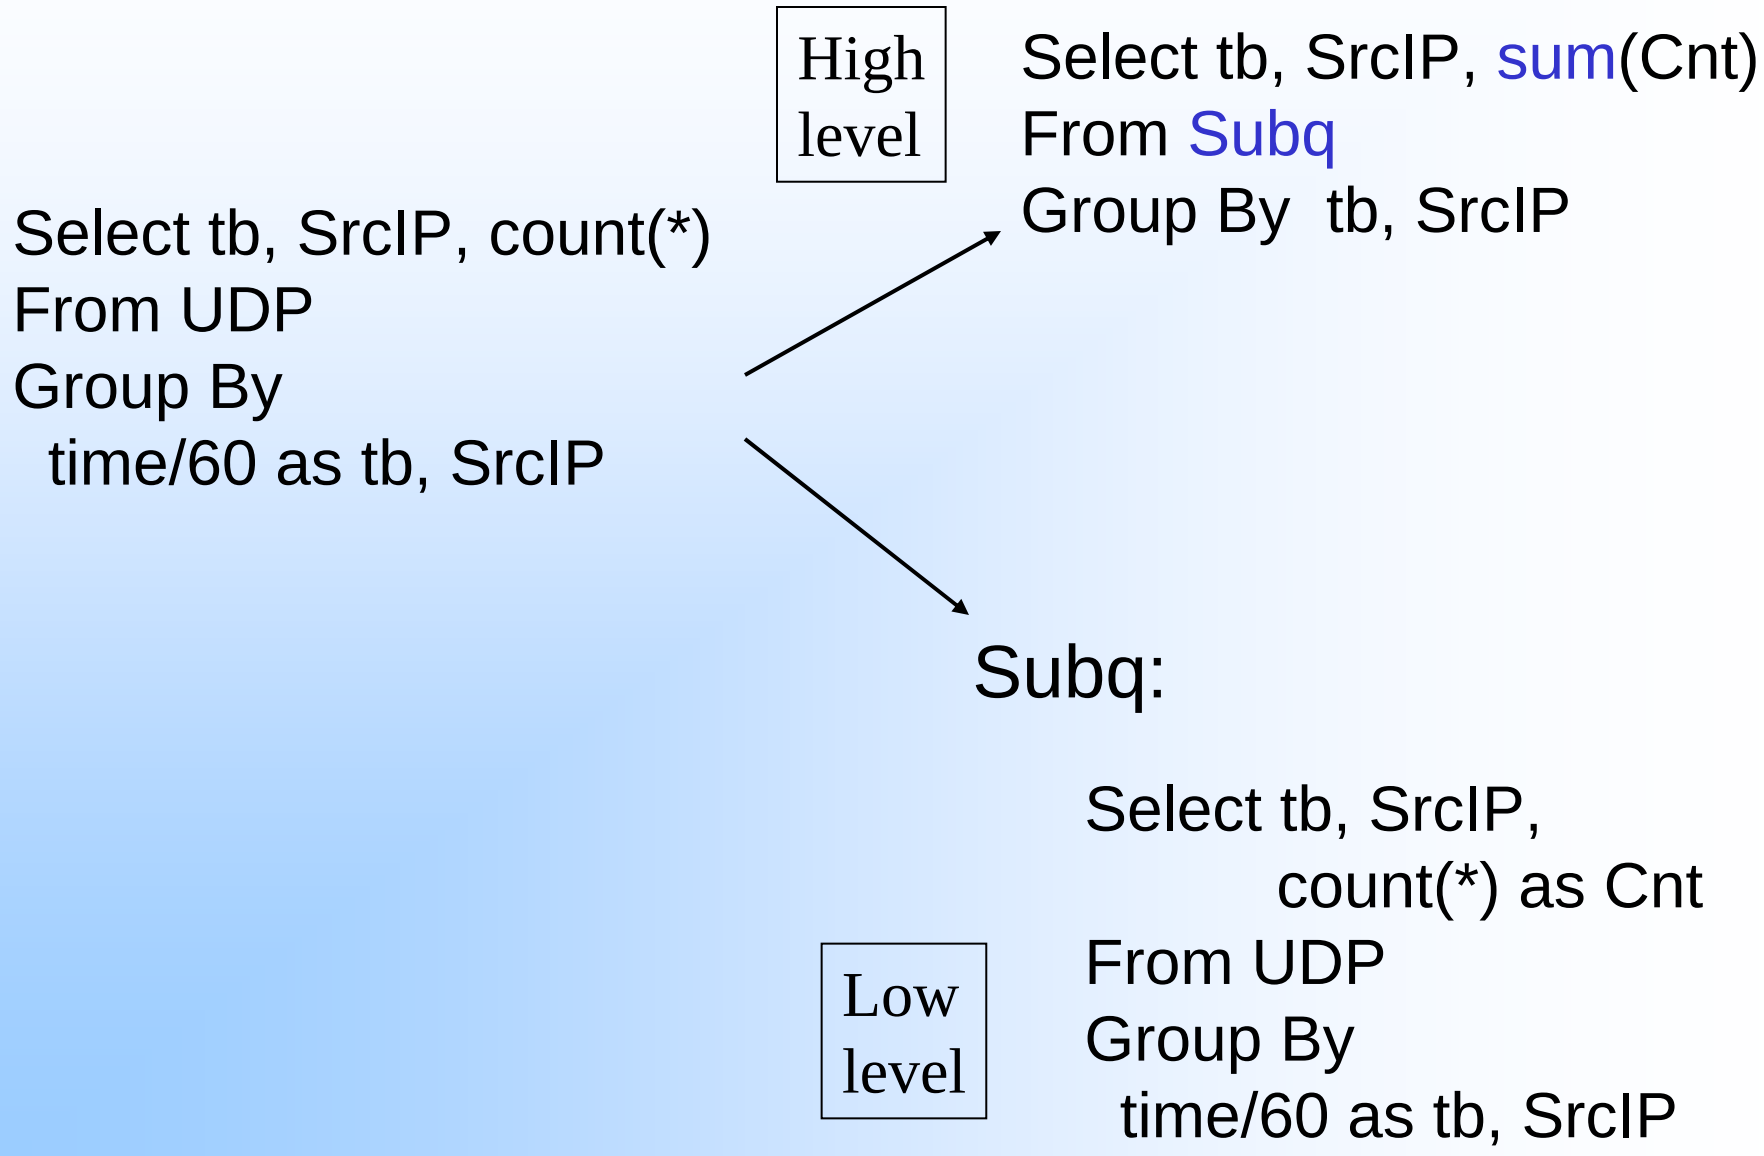

Gigascope, Status

Currently supports:

- GSQL, UDAFs.
  - stream aggregate queries.
- Sampling.
  - Operator can be specialized to most stream sampling methods.
  - Most complex queries can be executed with semantic sampling to provide correct output.
- Regex matcher for flows.
  - Match contents across packets in presence of duplicates, out-of-order or overlapping packets.
- Heartbeats.
  - Prelim distributed implementation.
- Query-aware query partitioning.
- Deployed

Ted Johnson S. Muthukrishnan
Irina Rozenbaum Vlad Shkapenyuk
Oliver Spatscheck.

Sampling Operator

- Many sampling algorithms known for IP traffic streams.
  - Uniform random sampling
  - Priority sampling
  - Value sampling
  - Distinct, inverse, minwise sampling.
- Observation:
  - Most sampling algorithms have a overall common execution structure.
- Our approach:
  - Define and optimize a single sampling operator.

Stream Sampling Operator

- Operator:

```
Select <select expression list>.
From <stream>.
Where <predicate>.
Group by <group-by variables definition list>.
Cleaning when <predicate>.
Cleaning by <predicate>.
[Having <predicate>].
```

  - **`Cleaning when`** – condition for triggering a cleaning phase.
  - **`Cleaning by`** – condition for sample reduction.

- Can be specialized for wide variety of stream sampling algorithms.
- Encourages experimentation and development of new sampling algorithms.

T. Johnson, S. Muthukrishnan and I. Rozenbaum, SIGMOD 2002.

Sampling Operator

War story:

- During SYN flooding and DDOS attacks, Cisco Netflow generator is overwhelmed and produces useless output.
- Packet sampling does not provide accurate flow samples.
- By combining flow sampling and flow generation logic using the sampling operator, Gigascope produces meaningful, valuable flow samples even at peak rates of flows such as in attacks.

Example Analysis

- Heavy hitter q-gram in packet contents.
- Design sampling+sketching method to skip over vast number of packets.

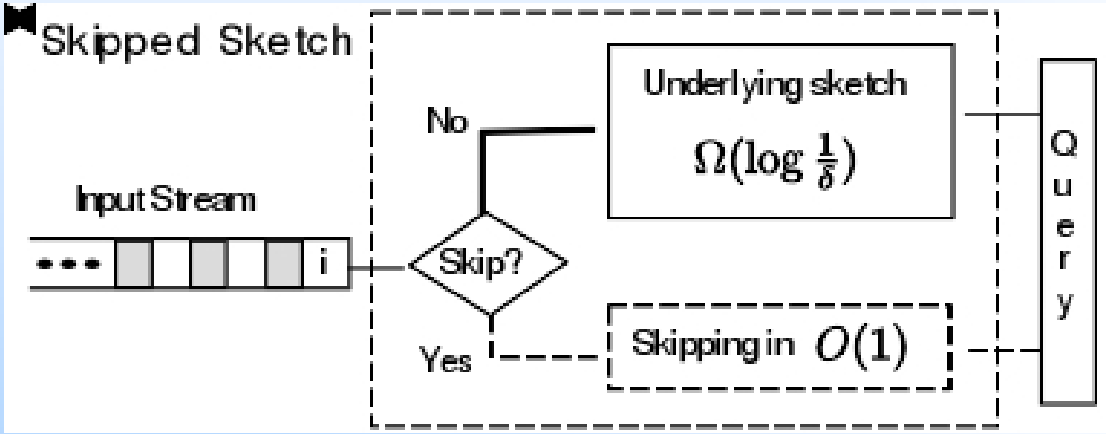

- Orders of magnitude improvement over prior work in networking, skipping fraction of packets.

S. Bhattacharyya, A. Maderia, S. Muthukrishnan and T. Ye. Sprint ATL Technical Report, 2006.

IP Traffic Analysis: Infrastructure

- Challenge:
  - Size, rate of data. Analyses: Simple.
  - Turnaround time: Minutes, days.
  - Moderate sized team of analysts.
- Special infrastructure:
  - Optical splitters, NIC
  - Multiple CPU machines
  - Data stream management systems (DSMSs): different architectures.

Talk Overview

- Data Analysis in Different Communities
  - Algorithms, Databases and Networking
- Infrastructure View of Data Analysis
  - Example 1: Cellphone Call Traffic
  - Example 2: IP Packet Traffic Streams
  - Example 3: Web Traffic 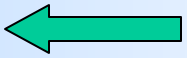
- Perspectives

Infrastructure View: Web Traffic Analysis

Google

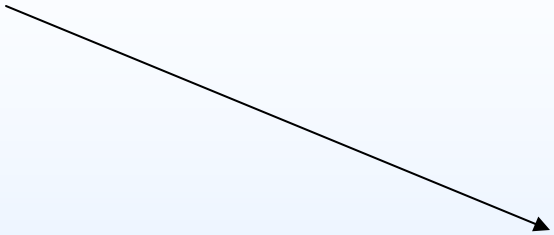

Search

Web
Image
Video
News
Usenet Groups
Blogs

Google: Calculator Co.

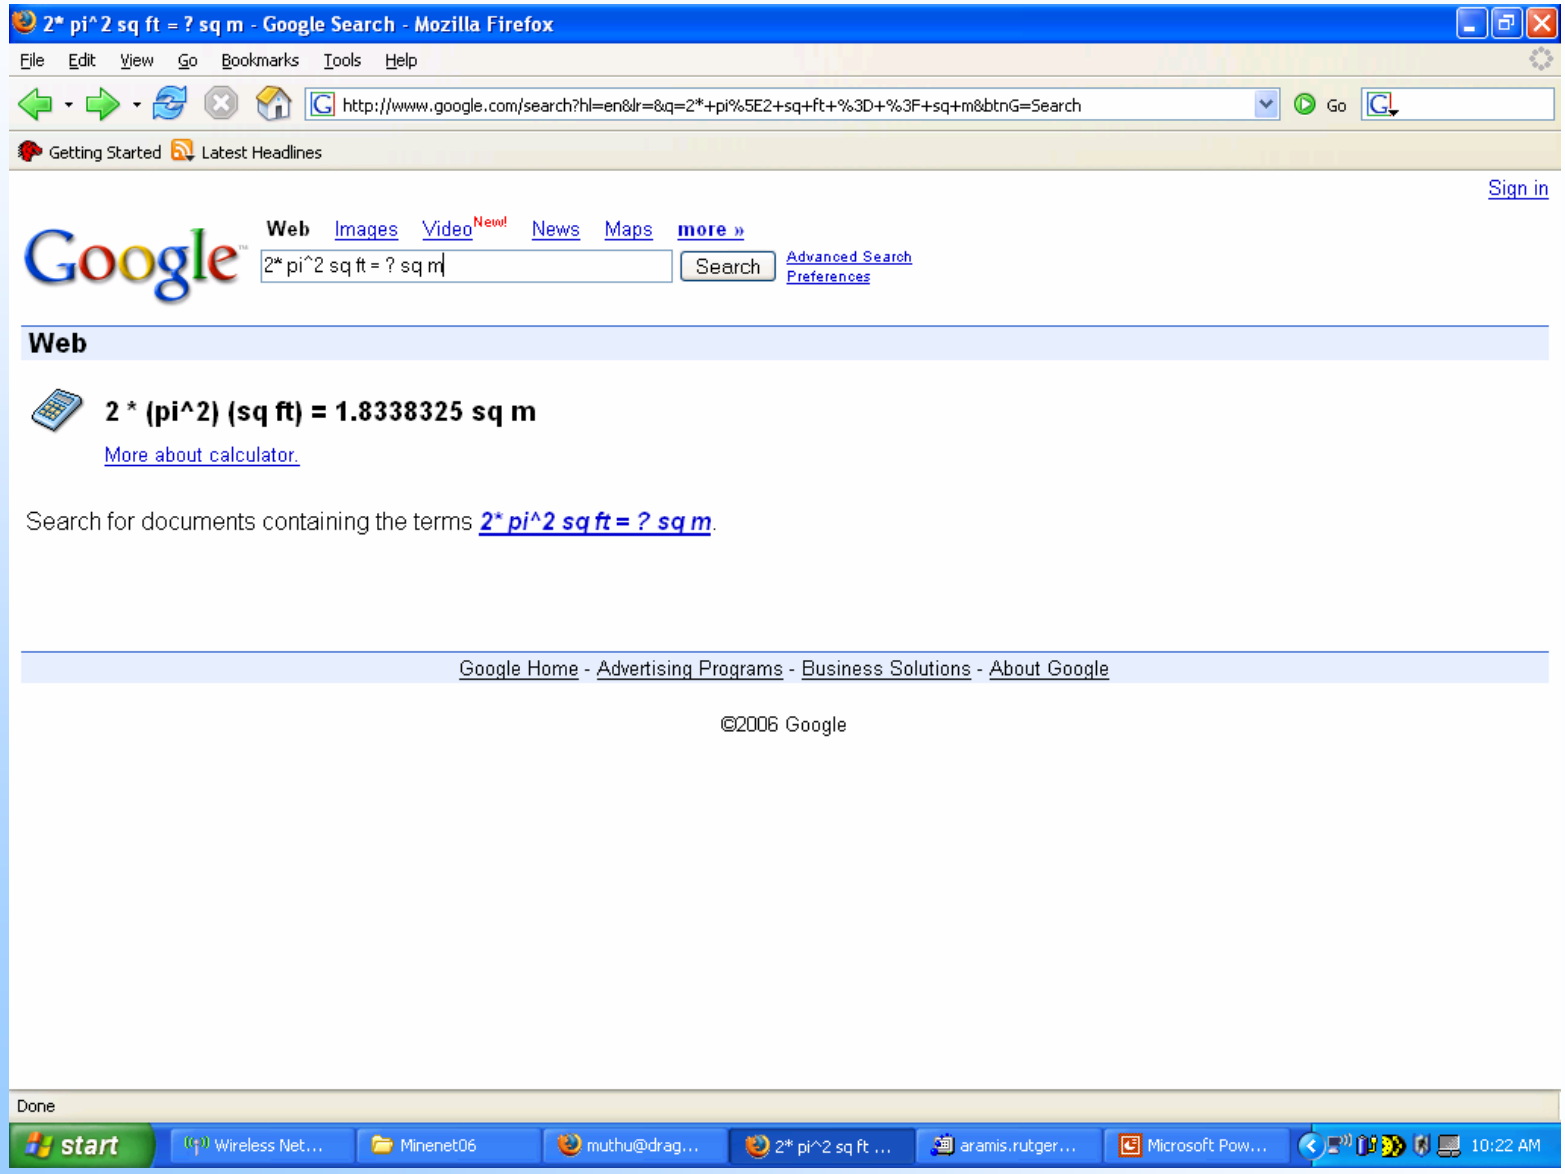

Google: Advertising

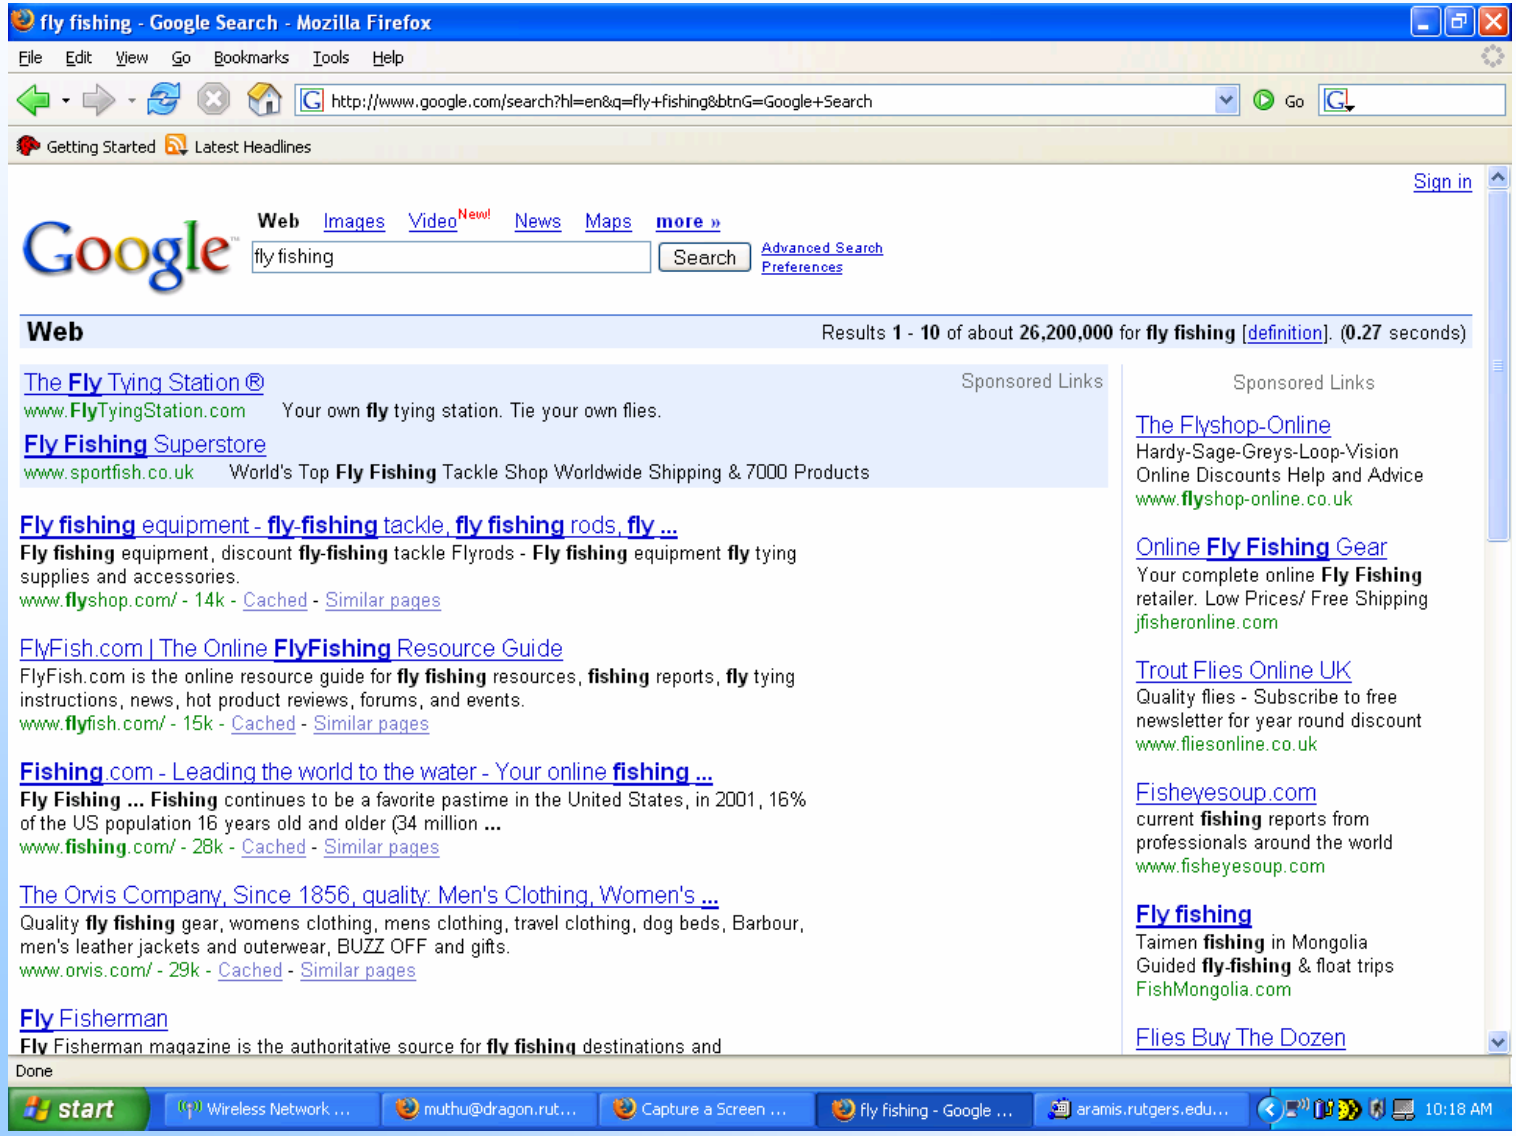

Google

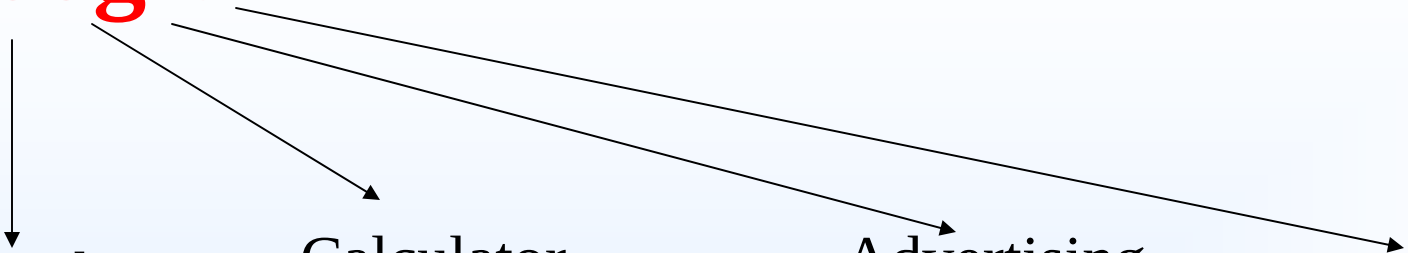

Search

Web
Image
Video
News
Usenet Groups
Blogs

Calculator Co.

Convert units,
Calculate.

Advertising

AdWords
AdSense
Partner sites
Coupons

Earth
Map
Finance
Trends
Writely
Personalize
Froogle
….

Example: Sponsored Search

- Advertisers want to place ads in response to user queries.
- Search companies place ads by running an auction in response to user queries.
- Have to figure out what queries are interesting, how much to bid on each query, what is the budget,…

Google Sponsored Search Auction

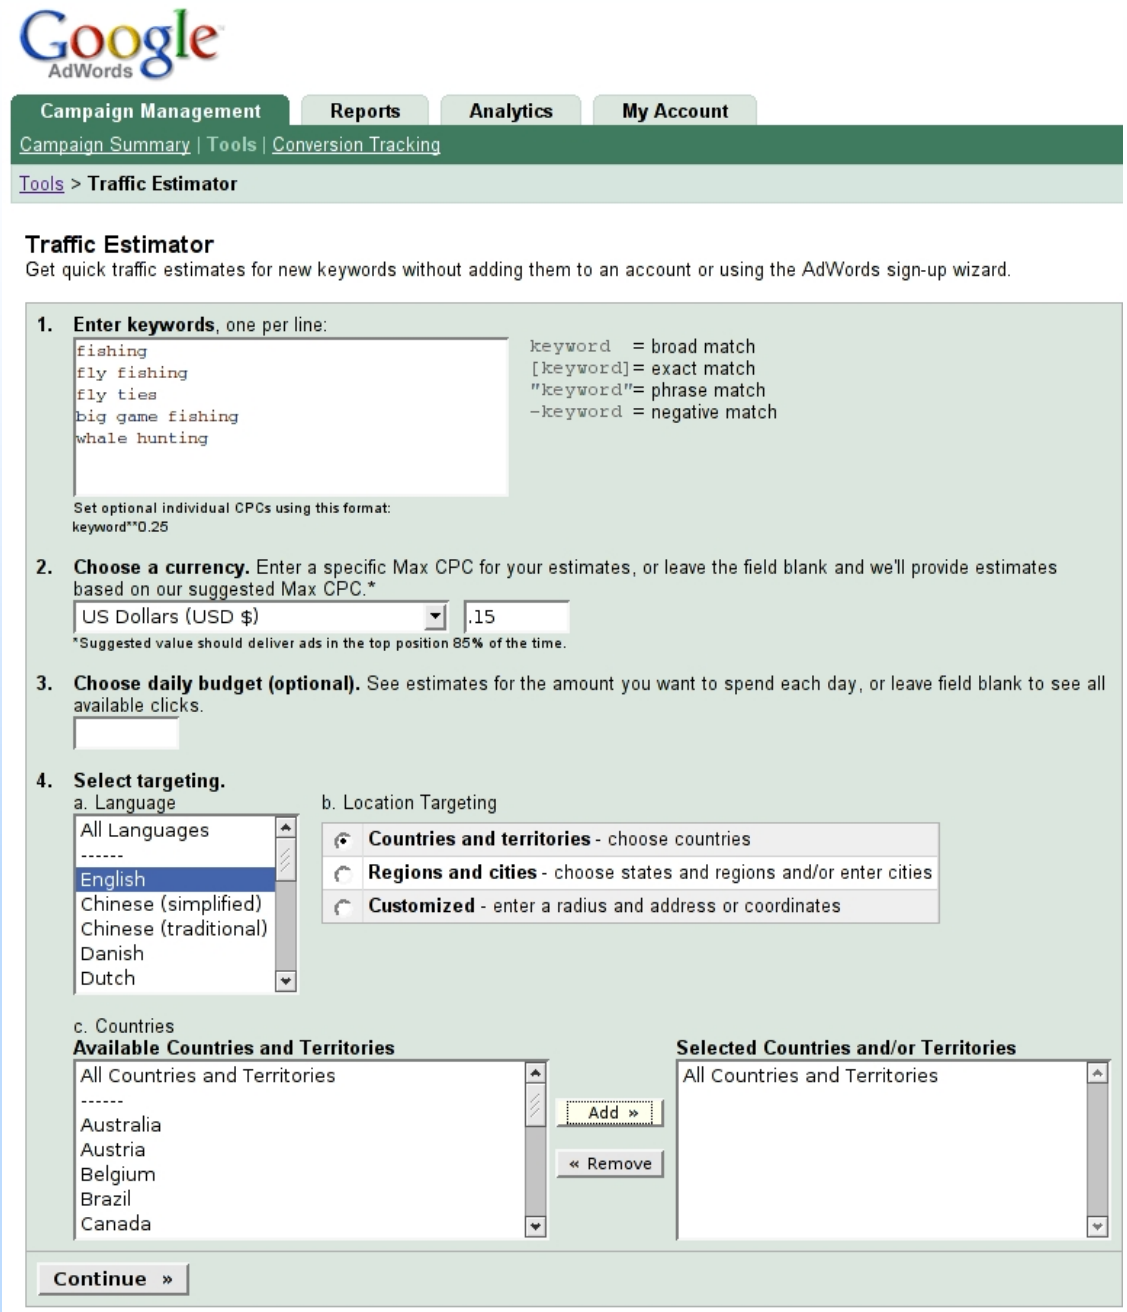

Traffic Estimation for Sponsored Search

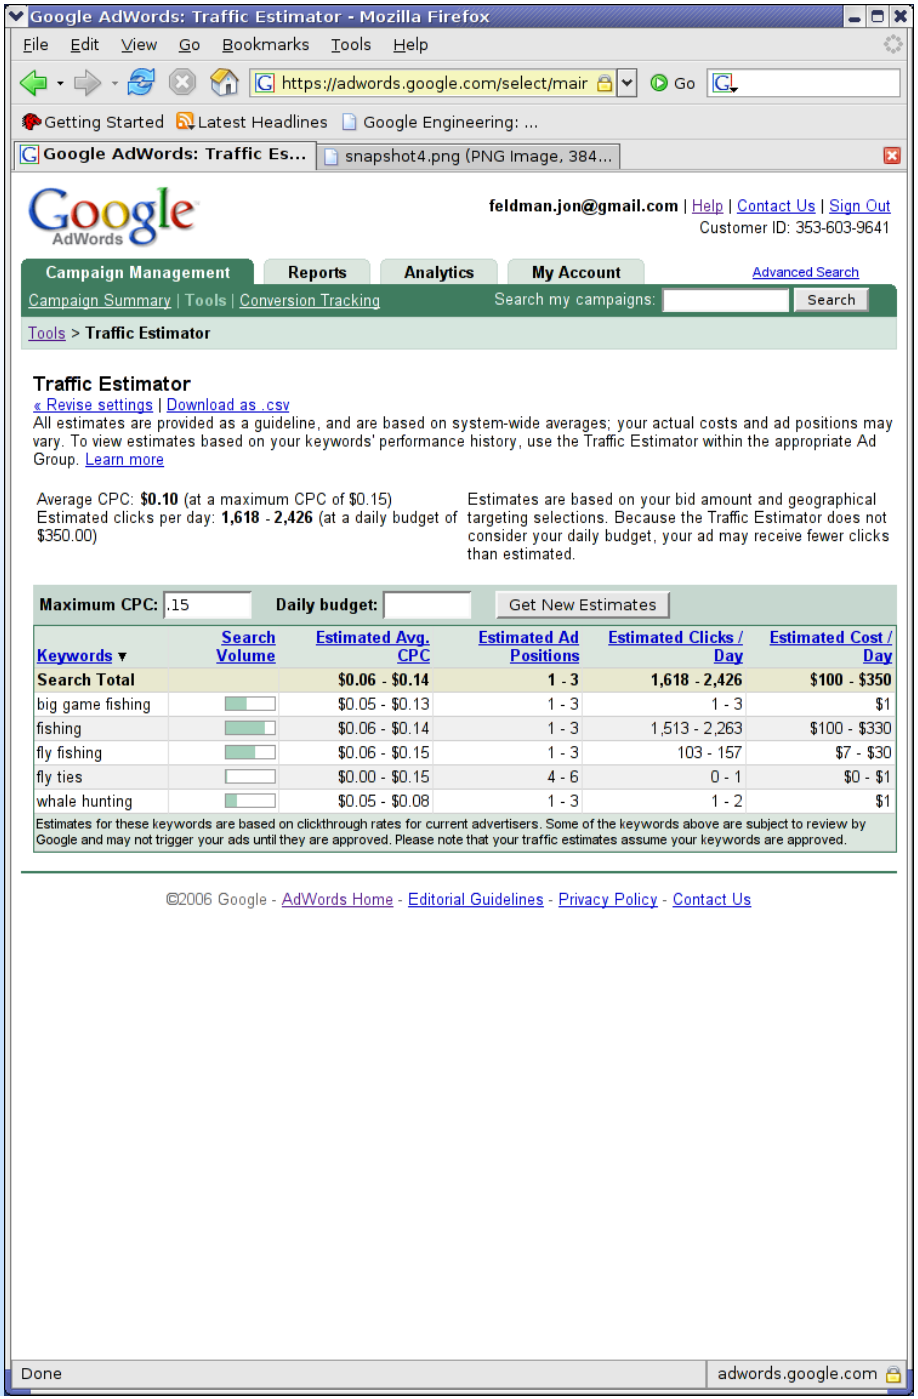

Example Analysis: Traffic Estimation

- Problem: Given a set of queries and a potential bid, output the distribution of
  - Number of clicks expected
  - Expected position on the ad list
  - Expected price.
- Input: queries, ads shown, bids, price, etc.Terabytes of data on 1000's of commodity machines.

MapReduce [Dean, Ghemawat OSDI04]

- Parallel programming infrastructure at Google.
- Users specify map and reduce functions.
- Input: set of records.
  - Each record is mapped to a set of (key, value) pairs.
  - All pairs with same key are considered together and a reduce function is applied to the values.
- System automatically takes care of
  - Parallelizing on 100's++ commodity machines.
  - Fault tolerance
  - Scheduling, load balance, locality, inter-machine communication, etc.

Traffic Estimation Using MapReduce

- Logs consist of $(q,b_1,p_1,b_2,p_2,..,c)$.
  - q is the query.
  - b_i is the bid of advertiser in *i*th place and p_i the price.
  - c is the ad clicked on.
- Map to $(q,b_i,p_i,i,1 \text{ if } c=i)$ for all i; q is the key.
- Reduce will have all records with same q. Calculate.
  - number of clicks,
  - average position,
  - average cost per click, etc.
- Run this periodically and index for each q. Lookup when needed.

Web Traffic Analysis: Infrastructure

- Terabytes of data on 1000's of commodity machines.
- 100's of engineers running many analyses simultaneously any day.
- Enormously successful at Google for machine learning, graph computing to index generation.

MapReduce was used for 29k jobs, dealt with 3k TB, 300+ programs, 79k machine days, in Aug 04, [OSDI04]

MapReduce

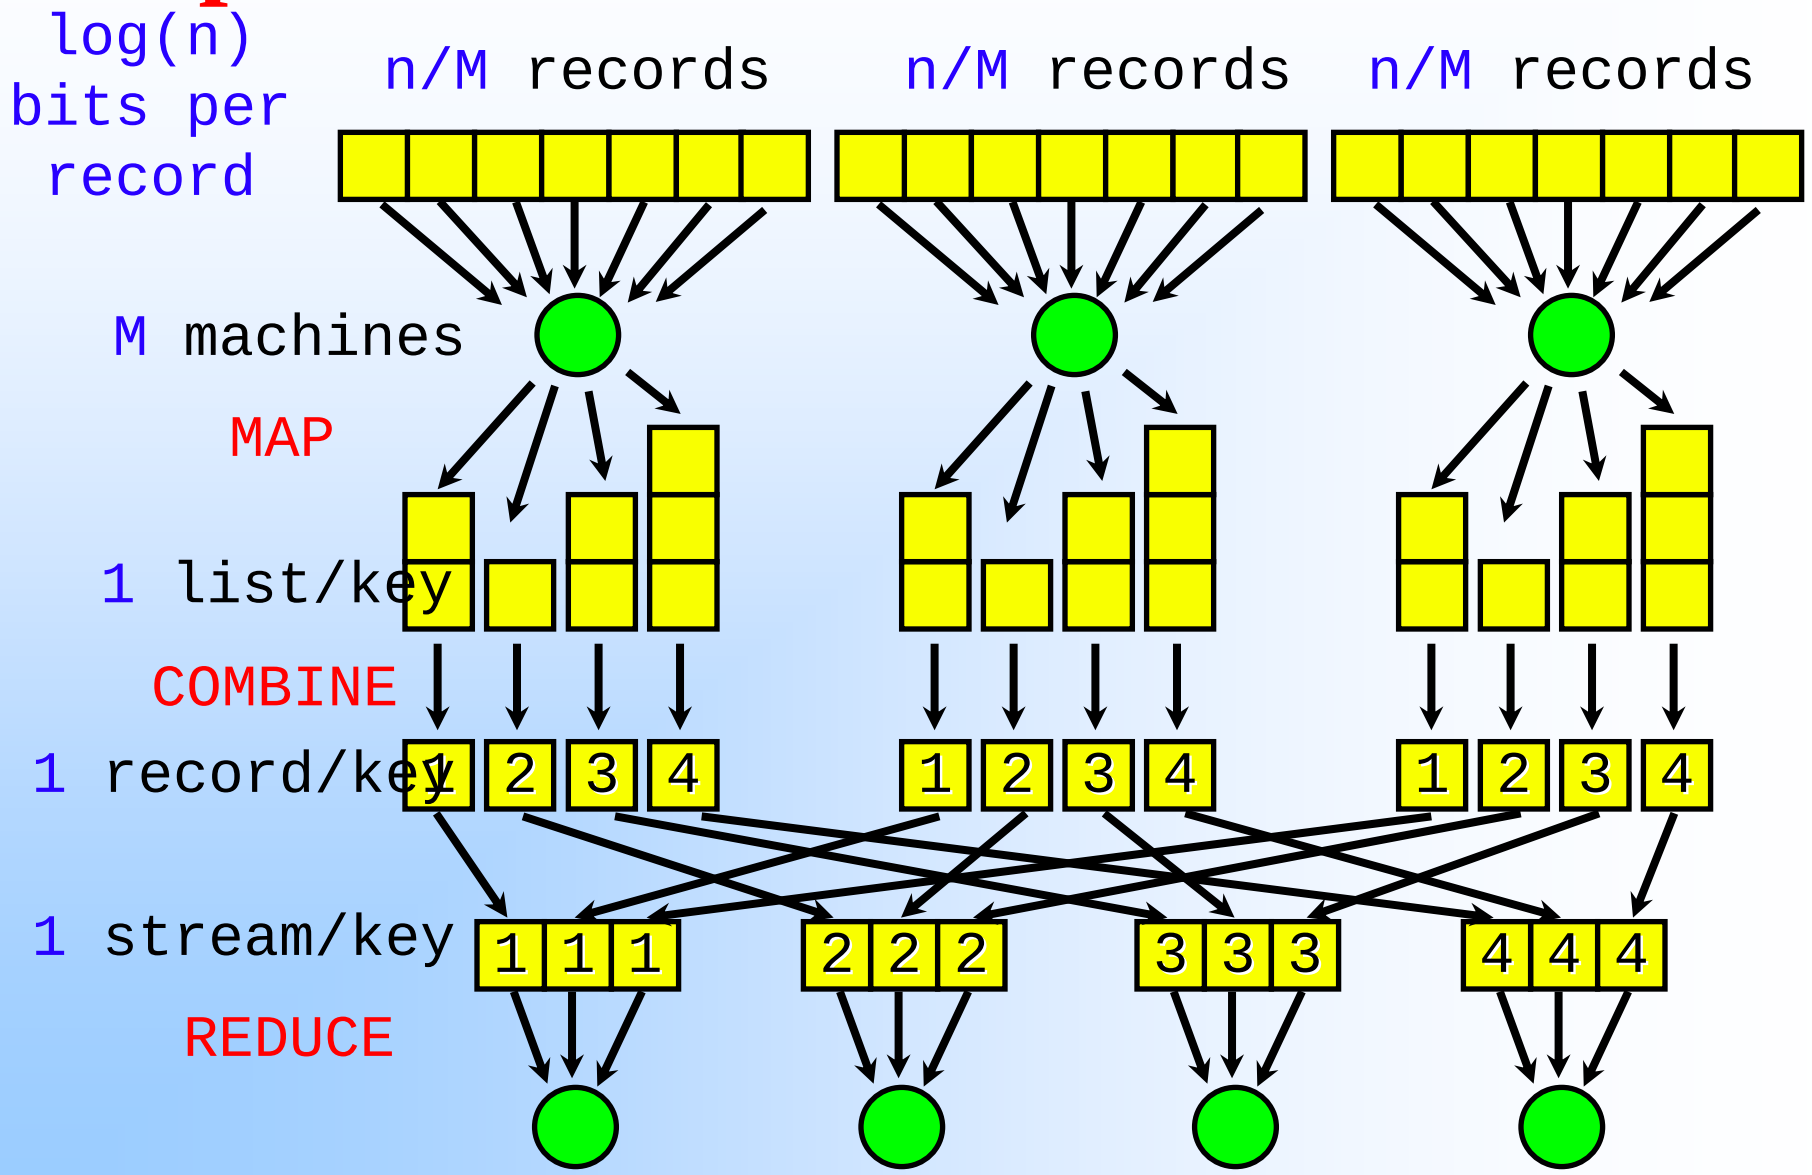
log(n) bits per record
n/M records
n/M records
n/M records
M machines
MAP
1 list/key
COMBINE
1 record/key
1 2 3 4
1 2 3 4
1 2 3 4
1 stream/key
1 1 1
2 2 2
3 3 3
4 4 4
REDUCE

MapReduce: Theoretical Model

- MUD Model: Assume each mapper is a stream, each reducer is a stream, and there is a single key.
  - Looks like distributed streaming.
- How is MUD related to streaming?
- For symmetric, total exact functions: MUD = SS.
- For promise problems and approximate functions, MUD \not= SS.
- With multiple keys, we can simulate PRAM.
- Open Problem: Given k keys and l rounds, can you solve various problems.

J. Feldman, S. Muthukrishnan, T. Sidiropoulos, Z. Svitkina, C. Stein.

Talk Overview

- Data Analysis in Different Communities
  - Algorithms, Databases and Networking
- Infrastructure View of Data Analysis
  - Example 1: Cellphone Call Traffic
  - Example 2: IP Packet Traffic Streams
  - Example 3: Web Traffic
- Perspectives 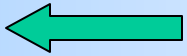

Summary

Cellphone traffic (cellco)	IP Traffic (ISP)	Web Traffic (Search Engine)
TB/month weekly/monthly Reports.	TB/hour min/hours/days Detect attacks, appl.	PB/month hours/days Nearly all services.
Small team of analysts.	Small/Moderate # of researchers	Large number of engineers/analysts
File system, script language, parallel CPUs.	Optical splitters, NICs, stream mgmt engines.	1000's of m/c's, GFS, MapReduce, Bigtable, …
No publications	Alg/DB since 96. Mainly publ.	Mainly systems.

Acknowledgements

- Thanks to Nathan Hamilton for 5+ years of cellular data analysis.
- Thanks to colleagues at Sprint, AT&T, Narus, Google.
- Thanks to students at Rutgers.