Alcatel·Lucent

A Simple and Efficient Estimation Method for Stream Expression Cardinalities

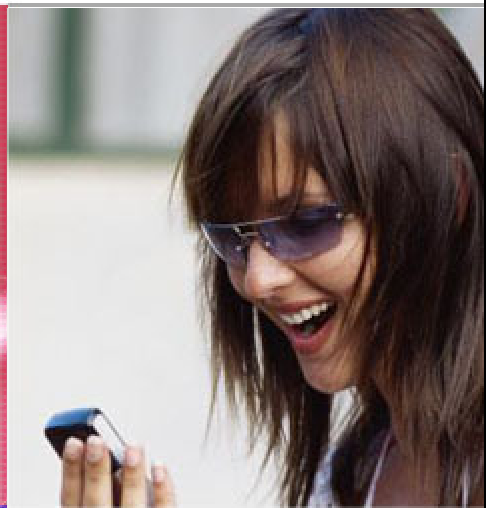

Aiyou Chen, Jin Cao, and Tian Bu (Bell Labs, Alcatel-Lucent)

Contact: aychen@research.bell-labs..com

Acknowledgement: Minos Garofalakis (Presenter)

Sept 25, 2007 VLDB'07, Vienna, Austria

Outline

1. Motivation of stream expression cardinalities
2. Literature review
3. A statistical model, algorithms and analysis
4. Simulation studies and applications for Traffic Matrix

Motivation of Stream Expression Cardinalities

Nature of network data streams

- Massive and fast
- Distributed at remote nodes

Statistical challenges

- What statistical sketches would be most efficient?
- Limited memory and computation

Traffic matrix problem: Each remote node sees one packet stream, how many flows or packets do each pair of streams share ?

The cardinality problem over multiple streams

- Two streams:

  Stream A = (a1, a2, …)

  Stream B = (b1, b2, …)

  (The same item may appear multiple times.)

  How many distinct items are in $A \cap B$ (set operation)?

Example: A can be a stream of packets: each item is the flow ID of a packet. $|A \cap B|$ is the number of common flows in both A and B.

- For more streams:

$$A \cap B \cap C$$

$$(A \cap B) - C$$

…….

Literature review on stream expression cardinalities

Key references:

- Use 1-stream cardinality algorithm:
  - Bitmap: Estan, Varghese & Fisk, IMC'03
  - LogLogN-Bitmap: Cai, Pan, Kwok & Hwang, SIGCOMM'05

  Basic idea: use set operation |AnB|=|A|+|B|-|AuB|. Not efficient $Error: O(|A \cup B| / |A \cap B|)$

  Our approach: use correlation information efficiently $Error: O(\sqrt{|A \cup B| / |A \cap B|})$

- Two-level-hash: Ganguly, Garofalakis & Rastogi, VLDBJ'04 ; Ganguly, ISAAC'05

  Basic idea: stores correlation information directly, space not efficient O((logN)^2)

  Our approach: Use continuous Flajolet-Martin sketches, space O(LogLogN)

For one stream cardinality problem, many works exist:

- Flajolet & Martin, FOCS'83, Alon, Matias & Szegedy, STOC'96,
- Durand & Flajolet, ESA'03 (LoglogN)
- And so on

Our solution: Generalize the Flajolet-Martin sketches

We propose a continuous version of Flajolet-Martin sketch:

Static sketch:

1) Associate each item with a pseudo-random number (universal hash)

2) For each stream, record the minimal value of the random numbers (FM sketch)

Let Y1 be the record for stream A and Y2 for stream B. Then Y1 and Y2 are correlated.

Partition all items into three groups:

$$A \cap B, A - B, B - A$$

Note: The famous FM sketch uses geometric random numbers. Here we use a continuous random number (truncated in the decimals in applications). We handle insert-only streams.

The Statistical Model

Diagram

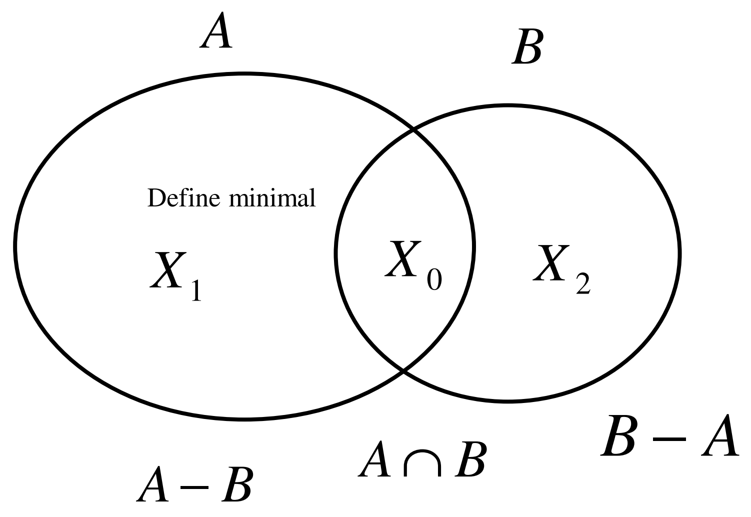

Relationship: Y1=min(X0,X1) & Y2=min(X0,X2).

By using exponential random numbers, X0~exp(1/|AnB|), similarly X1, X2 are exp.

Statistical principle: Maximum Likelihood Estimation –the likelihood is a function of the three cardinality parameters and estimate them by maximizing the likelihood function -- *MLE is Asymptotically unbiased and achieve the Cramer-Rao lower bound, but for more than 2 streams, MLE becomes complicated. Parameter numbers grow as 2^d-1 for expression over d streams.*

Alcatel·Lucent

A Proportional Union (PU) Estimator: Simple and Efficient

A nice property due to randomness: Because each item has the equal chance $\frac{1}{|A \cup B|}$ to be the global minimal, i.e. min(Y1,Y2), we have in probability 1:

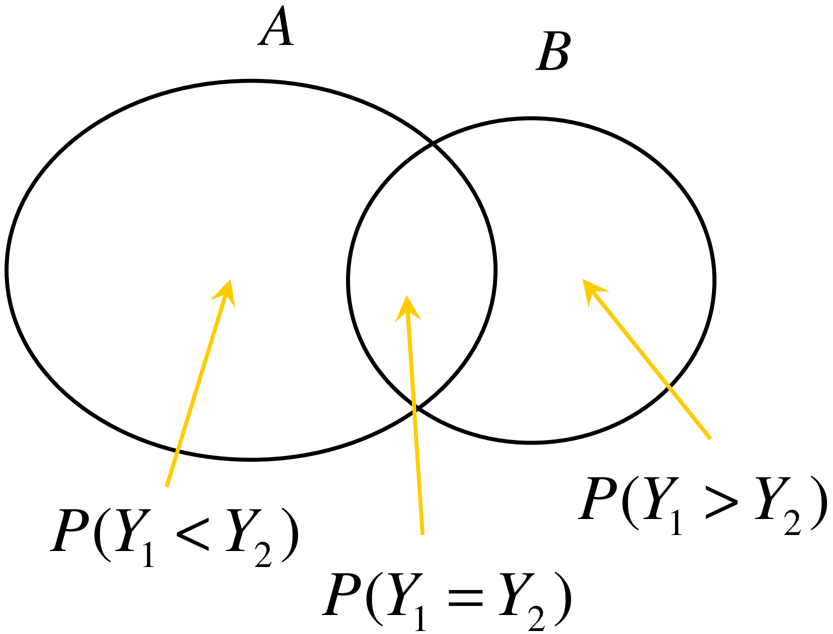

Y1=Y2 ←→ X0 is global minimal

Y1<Y2 ←→ X1 is global minimal

Y1>Y2 ←→ X2 is global minimal

Proportional property:

$$P(Y_1 = Y_2) = \frac{|A \cap B|}{|A \cup B|}$$

$$P(Y_1 < Y_2) = \frac{|A - B|}{|A \cup B|}$$

$$P(Y_1 > Y_2) = \frac{|B - A|}{|A \cup B|}$$

Stochastic sketches: To avoid multiple hashing, partition the item space randomly into **m** buckets and each bucket records a FM sketch

Estimation methods and performance

Estimation approaches:

- Maximum likelihood estimate (MLE).
- PU method: $|A \cap B| \propto P(Y_1 = Y_2)$

Note: min(Y1,Y2) ~ exp(1/|AuB|) => estimate |AuB|

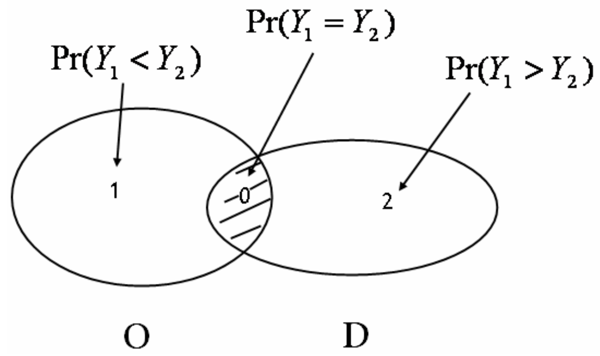

Cardinality and Proportions

Theorem (Relative Error).

$$\sqrt{m}\left(\frac{\hat{n}_{PU}}{n} - 1\right) \approx \sqrt{\frac{|A \cup B|}{|A \cap B|}} Normal(0,1),$$

where $n \equiv |A \cap B|$.

Better than State-of-the-Art algorithms

- Relative error is **independent of |AuB|** (Scale-Free).
- The relative error grows as **square root of noise-signal-ratio** ($\sqrt{\frac{|A \cup B|}{|A \cap B|}}$), better than expected. *Almost as efficient as MLE when the noise-signal-ratio is big*.
- E.g. If AnB has proportion 10%, then the standard relative error of PU is about 5% for m=4,000, no matter how big AuB is; MLE works slightly better.

Simulations and Comparison with State-of-the-Art

Compare the PU algorithm with 2-level-hash (Ganguly'05), LogLog-Bitmap (CPKH'05) for estimating |(AnB)-C|. Proportion is defined as |(AnB)-C|/|AuBuC|.

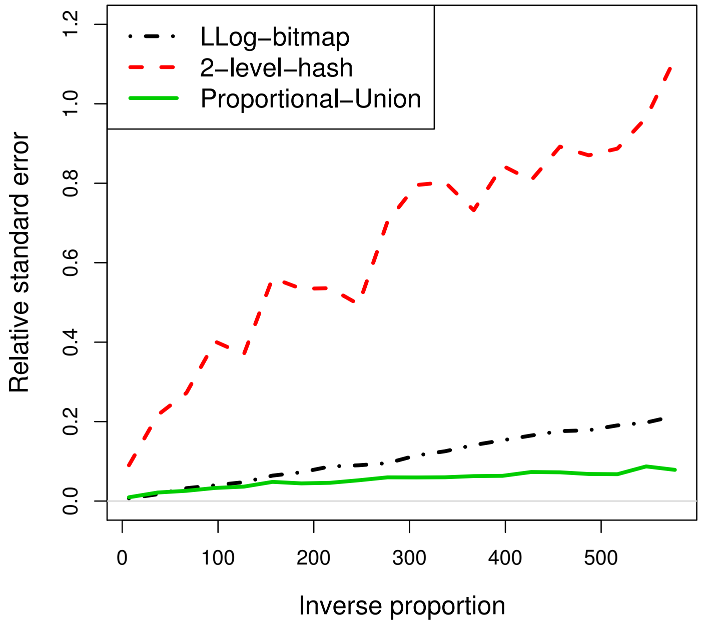

Relative std error of three methods

where N=3e7 and fix memory 188Kbytes.

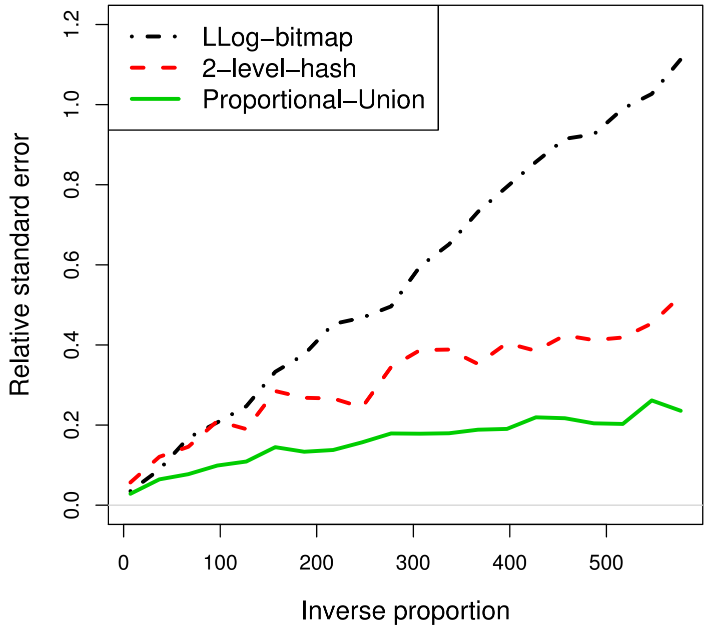

Relative std error of three methods

where N=3e7 and fix m=1e4.

Conclusion: PU works best, especially for small proportion.

Alcatel·Lucent

Experimental validation – Traffic matrix estimation

Traffic matrix (defined as numbers of common flows, i.e. OD flows, for each pair of network nodes) in 5 minutes network traffic based on a major network in US with 1800 OD pairs. Use m=60,000 at each node.

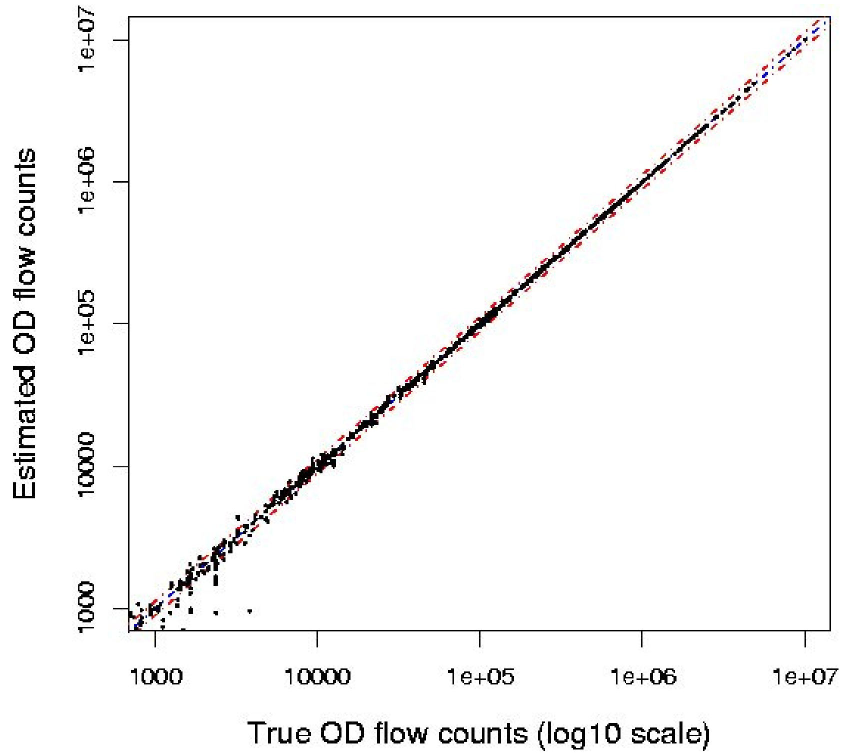

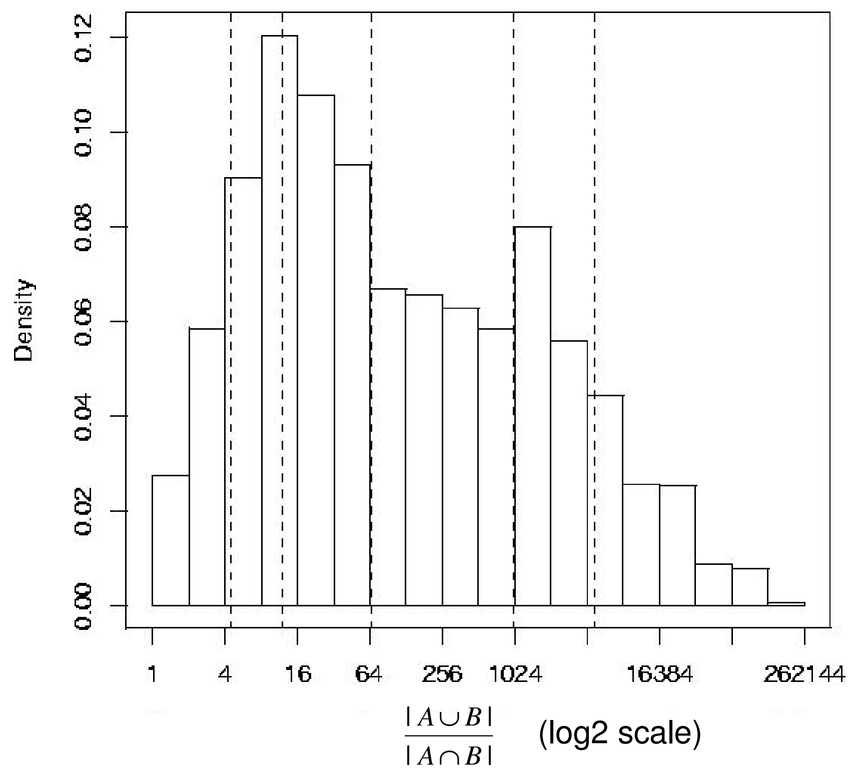

- Left panel shows good results: more than 50% node pairs have relative errors 1% or less and more than 90% node pairs have relative errors 10% or less.
- Right panel shows the Noise-Signal Ratios for node pairs: 50% have traffic proportion 1/64 or less.

Open question: Does there exist even more efficient sketches than the "continuous" FM sketches for estimating stream expression cardinalities?

THANK YOU!