In-core computation of distance distributions and geometric centralities with HyperBall: A hundred billion nodes and beyond

> Paolo Boldi, **Sebastiano Vigna** Laboratory for Web Algorithmics Università degli Studi di Milano, Italy

You have a very large graph (social, web)

- You have a very large graph (social, web)
- You want to understand something of its global structure (not triangles/degree distribution/etc.)

- You have a very large graph (social, web)
- You want to understand something of its global structure (not triangles/degree distribution/etc.)
- First candidate: *distance distribution* (and, in the directed case, the number of *reachable pairs*)

- You have a very large graph (social, web)
- You want to understand something of its global structure (not triangles/degree distribution/etc.)
- First candidate: *distance distribution* (and, in the directed case, the number of *reachable pairs*)
- You want to understand which nodes are important in some sense

First paper at WWW 2011 (with Marco Rosa)

- First paper at WWW 2011 (with Marco Rosa)
- Open-source software part of the WebGraph framework

- First paper at WWW 2011 (with Marco Rosa)
- Open-source software part of the WebGraph framework
- Run on Facebook (whole graph) using just

- First paper at WWW 2011 (with Marco Rosa)
- Open-source software part of the WebGraph framework
- Run on Facebook (whole graph) using just

HOME PA	HOME PAGE TODAY'S PAPER		VIDEO	VIDEO MOST POPULAR		TIMES TOPICS			
Technology									
WORLD	U.S.	.S. N.Y. / REGION		NESS	TECHNOLOGY		SCIENCE	HEALTH	SF
Separating You and Me? 4.74 Degrees By JOHN MARKOFF and SOMINI SENGUPTA Published: November 21, 2011									
The world is even smaller than you thought.								IMEND ER	Mie stuc er 5

Closeness (Bavelas 1946):

+ Closeness (Bavelas 1946): $\frac{1}{\sum_{y} d(y, x)}$

Closeness (Bavelas 1946): 1
The summation is over all *y* such that d(*y*,*x*)

- + Closeness (Bavelas 1946): $\frac{1}{\sum_{y} d(y, x)}$ + The summation is over all y such that
 - The summation is over all y such that $d(y,x) < \infty$
- Harmonic centrality:

- + Closeness (Bavelas 1946): $\frac{1}{\sum_{y} d(y, x)}$ + The summation is over all *y* such that
 - $d(y,x) < \infty$

+ Harmonic centrality: $\sum_{y \neq x} \frac{1}{d(y, x)}$



Why?

 Using HyperBall, we were able to evaluate geometric centrality in an IR setting

Why?

- Using HyperBall, we were able to evaluate geometric centrality in an IR setting
- The (preliminary) results show that harmonic centrality has a very good signal (in fact, better NDCG@10/P@10 than anything we tried)

Why?

- Using HyperBall, we were able to evaluate geometric centrality in an IR setting
- The (preliminary) results show that harmonic centrality has a very good signal (in fact, better NDCG@10/P@10 than anything we tried)
- In general, HyperBall makes it possible to use harmonic centrality on very large graphs

Hollywood: PageRank

Ron Jeremy

Adolf Hitler

Lloyd Kaufman



George W. Bush





Ronald Reagan



Bill Clinton



Martin Sheen



Debbie Rochon

Hollywood: PageRank

Ron Jeremy

Adolf Hitler

Lloyd Kaufman



George W. Bush





Ronald Reagan



Bill Clinton



Martin Sheen



Debbie Rochon

Hollywood: Harmonic

George Clooney



Samuel Jackson



Martin Sheen

Dennis Hopper

Sharon Stone

Antonio Banderas

Tom Hanks





Madonna

 For each node, we compute in sequence the number of nodes at distance exactly t

- For each node, we compute in sequence the number of nodes at distance exactly t
- Adding up over all nodes, we get the distance distribution (modulo normalization)

- For each node, we compute in sequence the number of nodes at distance exactly t
- Adding up over all nodes, we get the distance distribution (modulo normalization)
- Centralities can be rewritten, e.g., harmonic:

- For each node, we compute in sequence the number of nodes at distance exactly t
- Adding up over all nodes, we get the distance distribution (modulo normalization)
- Centralities can be rewritten, e.g., harmonic:

$$\sum_{t>0} \frac{1}{t} |\{y \mid d(y, x) = t\}|$$

 Many many breadth-first visits: O(mn), needs direct access

- Many many breadth-first visits: O(mn), needs direct access
- Sampling: a fraction of breadth-first visits, very unreliable results on graphs that are not strongly connected, needs direct access

- Many many breadth-first visits: O(mn), needs direct access
- Sampling: a fraction of breadth-first visits, very unreliable results on graphs that are not strongly connected, needs direct access
- Edith Cohen's [JCSS 1997] size estimation framework: very powerful but does not scale or parallelize really well, needs direct

Alternative: Diffusion

Alternative: Diffusion

+ Basic idea: Palmer et. al, KDD '02

Alternative: Diffusion

+ Basic idea: Palmer et. al, KDD '02

Let B_t(x) be the ball of radius t around x
 (nodes at distance at most t from x)
Alternative: Diffusion

- * Basic idea: Palmer et. al, KDD '02
- Let B_t(x) be the ball of radius t around x
 (nodes at distance at most t from x)

+ Clearly $B_0(x) = \{x\}$

Alternative: Diffusion

- * Basic idea: Palmer et. al, KDD '02
- Let B_t(x) be the ball of radius t around x
 (nodes at distance at most t from x)
- + Clearly $B_0(x) = \{x\}$
- + But also $B_{t+1}(x) = \bigcup_{x \to y} B_t(y) \bigcup \{x\}$

Alternative: Diffusion

- + Basic idea: Palmer et. al, KDD '02
- Let B_t(x) be the ball of radius t around x
 (nodes at distance at most t from x)
- + Clearly $B_0(x) = \{x\}$
- + But also $B_{t+1}(x) = \bigcup_{x \to y} B_t(y) \bigcup \{x\}$
- * So we can compute balls by enumerating the arcs $x \rightarrow y$ and performing set unions



















































Each set uses linear space; overall quadratic

+ Each set uses linear space; overall quadratic

+ Impossible!

- Each set uses linear space; overall quadratic
- * Impossible!
- But what if we use approximate sets?

- Each set uses linear space; overall quadratic
- * Impossible!
- But what if we use approximate sets?
- Idea: use *probabilistic counters*, which represent sets but answer just to "size?" questions

- Each set uses linear space; overall quadratic
- * Impossible!
- But what if we use approximate sets?
- Idea: use *probabilistic counters*, which represent sets but answer just to "size?" questions
- + Very small!

Main trick

Main trick

 Choose an approximate set such that unions can be computed quickly
- Choose an approximate set such that unions can be computed quickly
- ANF [Palmer *et al.*, KDD '02] uses Martin– Flajolet (MF) counters (log *n* + *c* space)

- Choose an approximate set such that unions can be computed quickly
- ANF [Palmer *et al.*, KDD '02] uses Martin– Flajolet (MF) counters (log *n* + *c* space)
- We use HyperLogLog counters [Flajolet *et al.*,2007] (log log *n* space)

- Choose an approximate set such that unions can be computed quickly
- ANF [Palmer *et al.*, KDD '02] uses Martin– Flajolet (MF) counters (log *n* + *c* space)
- We use HyperLogLog counters [Flajolet *et al.*,2007] (log log *n* space)
- MF counters can be combined with an OR

- Choose an approximate set such that unions can be computed quickly
- ANF [Palmer *et al.*, KDD '02] uses Martin– Flajolet (MF) counters (log *n* + *c* space)
- We use HyperLogLog counters [Flajolet *et al.*,2007] (log log *n* space)
- MF counters can be combined with an OR
- We use broadword programming to combine HyperLogLog counters quickly!

 Instead of actually counting, we *observe* a statistical feature of a set (think stream) of elements

- Instead of actually counting, we *observe* a statistical feature of a set (think stream) of elements
- The feature: the number of trailing zeroes of the value of a very good hash function

- Instead of actually counting, we *observe* a statistical feature of a set (think stream) of elements
- The feature: the number of trailing zeroes of the value of a very good hash function
- We keep track of the maximum *m* (log log *n* bits!)

- Instead of actually counting, we *observe* a statistical feature of a set (think stream) of elements
- The feature: the number of trailing zeroes of the value of a very good hash function
- We keep track of the maximum *m* (log log *n* bits!)
- + The number of distinct elements $\propto 2^m$

- Instead of actually counting, we *observe* a statistical feature of a set (think stream) of elements
- The feature: the number of trailing zeroes of the value of a very good hash function
- We keep track of the maximum *m* (log log *n* bits!)
- + The number of distinct elements $\propto 2^m$
- Important: the counter of stream *AB* is simply

To increase confidence, we need *several* counters (usually 2^b, b≥4) and take their harmonic mean

- To increase confidence, we need *several* counters (usually 2^b, b≥4) and take their harmonic mean
- Thus each set is represented by a list of small (typically 5-bit) counters (unlikely >6 bits!)

- To increase confidence, we need *several* counters (usually 2^b, b≥4) and take their harmonic mean
- Thus each set is represented by a list of small (typically 5-bit) counters (unlikely >6 bits!)
- To compute the union of two sets these must be maximized one-by-one

- To increase confidence, we need *several* counters (usually 2^b, b≥4) and take their harmonic mean
- Thus each set is represented by a list of small (typically 5-bit) counters (unlikely >6 bits!)
- To compute the union of two sets these must be maximized one-by-one
- Extracting by shifts, maximizing and putting back by shifts is unbearably slow

- To increase confidence, we need *several* counters (usually 2^b, b≥4) and take their harmonic mean
- Thus each set is represented by a list of small (typically 5-bit) counters (unlikely >6 bits!)
- To compute the union of two sets these must be maximized one-by-one
- Extracting by shifts, maximizing and putting back by shifts is unbearably slow





























 We keep track of modifications: we do not maximize with unmodified counters

- We keep track of modifications: we do not maximize with unmodified counters
- Systolic computation: each modified set signals back to predecessors that something is going to happen (much fewer updates O(m log n) in expectation! [Cohen])

- We keep track of modifications: we do not maximize with unmodified counters
- Systolic computation: each modified set signals back to predecessors that something is going to happen (much fewer updates O(m log n) in expectation! [Cohen])
- Multicore exploitation by decomposition: a task is updating just a batch of counters whose overall outdegree is predicted using an Elias-Fano representation of the cumulative outdegree distribution (almost

Footprint

Footprint

+ Scalability: a minimum of 20 bytes per node
Footprint

Scalability: a minimum of 20 bytes per node
On a 2TiB machine, 100 billion nodes

Footprint

- Scalability: a minimum of 20 bytes per node
- + On a 2TiB machine, 100 billion nodes
- Graph structure is accessed by memorymapping in a compressed form (WebGraph)

Footprint

- Scalability: a minimum of 20 bytes per node
- + On a 2TiB machine, 100 billion nodes
- Graph structure is accessed by memorymapping in a compressed form (WebGraph)
- Pointer to the graph are store using quasisuccinct lists (Elias-Fano representation)

On a 177K nodes / 2B arcs graph, RSD ~14%:

- On a 177K nodes / 2B arcs graph, RSD ~14%:
- Hadoop: 2875s per iteration [Kang, Papadimitriou, Sun and H. Tong, 2011]

- On a 177K nodes / 2B arcs graph, RSD ~14%:
- Hadoop: 2875s per iteration [Kang, Papadimitriou, Sun and H. Tong, 2011]
- HyperBall on this laptop: 70s per iteration

- On a 177K nodes / 2B arcs graph, RSD ~14%:
- Hadoop: 2875s per iteration [Kang, Papadimitriou, Sun and H. Tong, 2011]
- HyperBall on this laptop: 70s per iteration
- On a 32-core workstation: 23s per iteration

- On a 177K nodes / 2B arcs graph, RSD ~14%:
- Hadoop: 2875s per iteration [Kang, Papadimitriou, Sun and H. Tong, 2011]
- HyperBall on this laptop: 70s per iteration
- + On a 32-core workstation: 23s per iteration
- On ClueWeb09 (4.8G nodes, 8G arcs) on a 40-core workstation: 141m (avg. 40s per iteration)

Convergence Harmonic centrality



 Cohen's estimation framework provides error bounds for the relative error of the probability mass function (and centralities)

- Cohen's estimation framework provides error bounds for the relative error of the probability mass function (and centralities)
- ANF/HyperANF give only pointwise guarantees, but provide error for the *absolute* error of the probability mass function (and centralities)

- Cohen's estimation framework provides error bounds for the relative error of the probability mass function (and centralities)
- ANF/HyperANF give only pointwise guarantees, but provide error for the *absolute* error of the probability mass function (and centralities)
- Sampling provides only the latter and only for strongly connected graphs

- Cohen's estimation framework provides error bounds for the relative error of the probability mass function (and centralities)
- ANF/HyperANF give only pointwise guarantees, but provide error for the *absolute* error of the probability mass function (and centralities)
- Sampling provides only the latter and only for strongly connected graphs
- ...but we can retrofit Cohen's estimators on HyperANF, obtaining an extremely efficient version of Cohen's framework!

 Perfect and natural fit for distributed computation (GraphLab, Pregel, etc.)

- Perfect and natural fit for distributed computation (GraphLab, Pregel, etc.)
- Apply the same computational framework to other size estimators

- Perfect and natural fit for distributed computation (GraphLab, Pregel, etc.)
- Apply the same computational framework to other size estimators
- Edith Cohen new HIP estimators for HyperLogLog counters might work

- Perfect and natural fit for distributed computation (GraphLab, Pregel, etc.)
- Apply the same computational framework to other size estimators
- Edith Cohen new HIP estimators for HyperLogLog counters might work
- * http://webgraph.di.unimi.it/ *** software

- Perfect and natural fit for distributed computation (GraphLab, Pregel, etc.)
- Apply the same computational framework to other size estimators
- Edith Cohen new HIP estimators for HyperLogLog counters might work
- * <u>http://webgraph.di.unimi.it/</u>
 software
- http://law.di.unimi.it/ datasets