



# Updating the Stationary Vector

of an

## Irreducible Markov Chain

with an eye on

## Google's PageRank

Amy Langville

Carl Meyer

Department of Mathematics  
North Carolina State University  
Raleigh, NC

SIAM CSE 2/13/05



# Outline

- PageRank Solution Methods
- A Reordering for PageRank
- Updating PageRank



# PageRank

## The Hyperlink Matrix $\mathbf{H}$

$$\mathbf{H}_{ij} = 1/|O_i|$$

## The Stochastic Matrix $\mathbf{S}$

$$\mathbf{S} = \mathbf{H} + \mathbf{a}\mathbf{v}^T$$

( $a_i=1$  for  $i \in D$ , 0, o.w.)

## The Google Matrix $\mathbf{G}$

$$\begin{aligned}\mathbf{G} &= \alpha\mathbf{S} + (1 - \alpha)\mathbf{e}\mathbf{v}^T \\ &= \alpha\mathbf{H} + (\alpha\mathbf{a} + (1 - \alpha)\mathbf{e})\mathbf{v}^T\end{aligned}$$

- $\mathbf{G}$  is irreducible, aperiodic Markov chain.
- Stationary vector of  $\mathbf{G}$  is PageRank vector  $\boldsymbol{\pi}^T$ .

$\pi_i$  is long-run proportion of time that random surfer spends on page  $i$ .



# Computing $\pi^T$

## A Big Problem

$$\text{Solve } \pi^T = \pi^T \mathbf{G}$$

(stationary distribution vector)

$$\pi^T (\mathbf{I} - \mathbf{G}) = \mathbf{0}$$

(too big for direct solves)



# Computing $\pi^T$

## A Big Problem

Solve  $\pi^T = \pi^T \mathbf{G}$  (stationary distribution vector)

$\pi^T (\mathbf{I} - \mathbf{G}) = \mathbf{0}$  (too big for direct solves)

Start with  $\pi_0^T = \mathbf{e}/n$  and iterate  $\pi_{j+1}^T = \pi_j^T \mathbf{G}$  (power method)



# Power Method to compute PageRank

$$\pi_0^T = \mathbf{e}^T / n$$

until convergence, do

$$\pi_{j+1}^T = \pi_j^T \mathbf{G}$$

(dense computation)

end



# Power Method to compute PageRank

$$\pi_0^T = \mathbf{e}^T / n$$

until convergence, do

**X**  $\pi_{j+1}^T = \pi_j^T \mathbf{G}$  (dense computation)

**•**  $\pi_{j+1}^T = \alpha \pi_j^T \mathbf{S} + (1 - \alpha) \pi_j^T \mathbf{e} \mathbf{v}^T$  (sparser computation)

end



# Power Method to compute PageRank

$$\boldsymbol{\pi}_0^T = \mathbf{e}^T / n$$

until convergence, do

**X**  $\boldsymbol{\pi}_{j+1}^T = \boldsymbol{\pi}_j^T \mathbf{G}$  (dense computation)

**X**  $\boldsymbol{\pi}_{j+1}^T = \alpha \boldsymbol{\pi}_j^T \mathbf{S} + (1 - \alpha) \boldsymbol{\pi}_j^T \mathbf{e} \mathbf{v}^T$  (sparser computation)

•  $\boldsymbol{\pi}_{j+1}^T = \alpha \boldsymbol{\pi}_j^T \mathbf{H} + (\alpha \boldsymbol{\pi}_j^T \mathbf{a} + (1 - \alpha)) \mathbf{v}^T$  (even less computation)

end

- $\mathbf{H}$  is very, very sparse with about 3-10 nonzeros per row.
- $\Rightarrow$  one vector-matrix mult. is  $O(nnz(\mathbf{P})) \approx O(n)$ .





# Convergence

Can prove  $\lambda_2(\mathbf{G}) \leq \alpha$

( $\Rightarrow$  asymptotic rate of convergence of PageRank method is rate at which  $\alpha^k \rightarrow 0$ )

## Google

- uses  $\alpha = .85$  (5/6, 1/6 interpretation)
- report 50-100 iterations til convergence
- still takes days to converge



# Enhancements to the PR power method

- Kamvar et al. Extrapolation
- Kamvar et al. Adaptive PageRank
- Kamvar et al. BlockRank
- Lee et al. Lumpability of Dangling Nodes
- Langville/Meyer: Updating PageRank
- Ipsen/Kirkland: more theory for Langville/Meyer



# Langville/Meyer Updating

## Motivation

- Updating PR is huge problem. Currently done monthly, but web changes hourly.
- Chien et al. use aggregation to focus on pages whose PR is most likely to change.

## Idea

- Use iterative aggregation to extend Chien idea.
- Focus on bad states, aggregate good, fast-converging states into one superstate.
- $\Rightarrow$  only work on much smaller aggregated chain.

## Results

- speedup by factor of 5-10 on some datasets.

## Issue


- Partitioning into good and bad states is hard, and IAD is very sensitive to partition.



# Idea behind Aggregation

Best for NCD systems

(Simon and Ando (1960s), Courtois (1970s))



Pro

exploits structure to reduce work

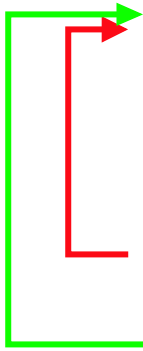
Con

produces an approximation, quality is dependent on degree of coupling



# Iterative Aggregation

- Problem: repeated aggregation leads to fixed point.
- Solution: Do a power step to move off fixed point.
- Do this iteratively. Approximations improve and approach exact solution.
- Success with NCD systems, not in general.



Input: approximation to  $\Pi^T$   
get censored distributions  $\Pi^T \quad \Pi^T \quad \Pi^T$   
get coupling constants  $\xi_i$

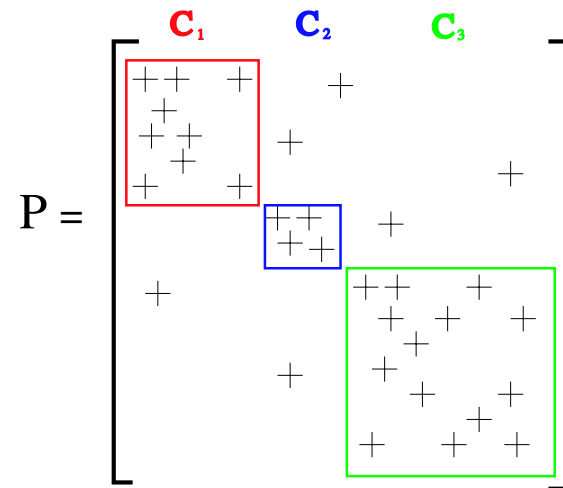
Output: get approximate global stationary distribution  $\Pi^T = \left[ \xi_1 \Pi^T \quad \xi_2 \Pi^T \quad \xi_3 \Pi^T \right]$

Output: move off fixed point with power step



# Exact Aggregation

(Meyer 1989)



$s_i^T =$  censored (stat.) dist. of stochastic complement  $S_i$

$$S_i = P_{ii} + P_{i*} (I - P_*)^{-1} P_{*i}$$

For 2-level partition,

$$S_1 = P_{11} + P_{12} (I - P_{22})^{-1} P_{21}$$

$$A = \begin{matrix} c_1 & c_2 & c_3 \\ c_1 & c_2 & c_3 \\ c_2 & & \\ c_3 & & \end{matrix} \begin{bmatrix} & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \\ & & \end{bmatrix} \quad a_{ij} = s_i^T P_{ij} e \quad \xi^T = \begin{bmatrix} \xi_1 & \xi_2 & \xi_3 \end{bmatrix}$$

$$\Pi^T = \begin{bmatrix} \xi_1 s_1^T & \xi_2 s_2^T & \xi_3 s_3^T \end{bmatrix}$$

Pro




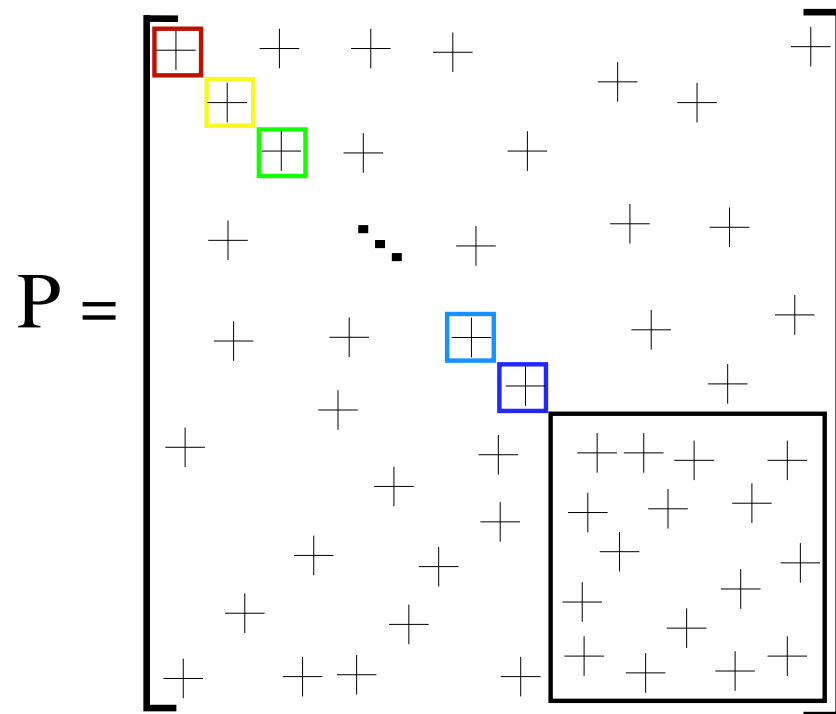
Con

only one step needed to produce exact global vector

SC matrices  $S_i$  are very expensive to compute



# Back to Updating . . .





# Aggregation

## Partitioned Matrix

$$\mathbf{P}_{n \times n} = \begin{matrix} G & \overline{G} \\ G & \overline{G} \end{matrix} \begin{pmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \\ \mathbf{P}_{21} & \mathbf{P}_{22} \end{pmatrix} = \left[ \begin{array}{c|c|c|c} p_{11} & \cdots & p_{1g} & \mathbf{r}_1^T \\ \hline \vdots & \ddots & \vdots & \vdots \\ \hline p_{g1} & \cdots & p_{gg} & \mathbf{r}_g^T \\ \hline \mathbf{c}_1 & \cdots & \mathbf{c}_g & \mathbf{P}_{22} \end{array} \right]$$

$$\boldsymbol{\pi}^T = (\pi_1, \dots, \pi_g \mid \pi_{g+1}, \dots, \pi_n)$$

## Advantages of this Partition

$p_{11} \cdots p_{gg}$  are  $1 \times 1 \implies$  Stochastic complements = 1

$\implies$  censored distributions = 1

Only one significant complement  $\mathbf{S}_2 = \mathbf{P}_{22} + \mathbf{P}_{21}(\mathbf{I} - \mathbf{P}_{11})^{-1}\mathbf{P}_{12}$

Only one significant censored dist  $\mathbf{s}_2^T \mathbf{S}_2 = \mathbf{s}_2^T$

A/D Theorem  $\implies \mathbf{s}_2^T = (\pi_{g+1}, \dots, \pi_n) / \sum_{i=g+1}^n \pi_i$





# Aggregation Matrix

$$\mathbf{A} = \begin{bmatrix} p_{11} & \cdots & p_{1g} & \mathbf{r}_1^T \mathbf{e} \\ \vdots & \ddots & \vdots & \vdots \\ p_{g1} & \cdots & p_{gg} & \mathbf{r}_g^T \mathbf{e} \\ \mathbf{s}_2^T \mathbf{c}_1 & \cdots & \mathbf{s}_2^T \mathbf{c}_g & \mathbf{s}_2^T \mathbf{P}_{22} \mathbf{e} \end{bmatrix}_{(g+1) \times (g+1)} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \mathbf{e} \\ \mathbf{s}_2^T \mathbf{P}_{21} & \mathbf{1} - \mathbf{s}_2^T \mathbf{P}_{21} \mathbf{e} \end{bmatrix}$$

## The Aggregation/Disaggregation Theorem

If  $\alpha^T = (\alpha_1, \dots, \alpha_g, \alpha_{g+1}) =$  stationary dist for  $\mathbf{A}$

Then  $\pi^T = (\alpha_1, \dots, \alpha_g \mid \alpha_{g+1} \mathbf{s}_2^T) =$  stationary dist for  $\mathbf{P}$

## Trouble! Always A Big Problem

$G$  small  $\Rightarrow \bar{G}$  big  $\Rightarrow \mathbf{S}_2 = \mathbf{P}_{22} + \mathbf{P}_{21}(\mathbf{I} - \mathbf{P}_{11})^{-1} \mathbf{P}_{12}$  large

$G$  big  $\Rightarrow \mathbf{A}$  large



# Approximate Aggregation

## Assumption

Updating involves relatively few states

$$G \text{ small} \Rightarrow \mathbf{A} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12}\mathbf{e} \\ \mathbf{s}_2^T \mathbf{P}_{21} & \mathbf{1} - \mathbf{s}_2^T \mathbf{P}_{21}\mathbf{e} \end{bmatrix}_{(g+1) \times (g+1)} \text{ small}$$

**Approximation**  $(\pi_{g+1}, \dots, \pi_n) \approx (\phi_{g+1}, \dots, \phi_n)$ ,  
 where  $\phi^T$  is old PageRank vector and  $\pi^T$  is new, updated PageRank

$$\mathbf{s}_2^T = \frac{(\pi_{g+1}, \dots, \pi_n)}{\sum_{i=g+1}^n \pi_i} \approx \frac{(\phi_{g+1}, \dots, \phi_n)}{\sum_{i=g+1}^n \phi_i} = \tilde{\mathbf{s}}_2^T$$

(avoids computing  $\tilde{\mathbf{s}}_2^T$  for large  $\mathbf{S}_2$ )

$$\mathbf{A} \approx \tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12}\mathbf{e} \\ \tilde{\mathbf{s}}_2^T \mathbf{P}_{21} & \mathbf{1} - \tilde{\mathbf{s}}_2^T \mathbf{P}_{21}\mathbf{e} \end{bmatrix}$$

$$\alpha^T \approx \tilde{\alpha}^T = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_g, \tilde{\alpha}_{g+1})$$

$$\pi^T \approx \tilde{\pi}^T = (\tilde{\alpha}_1, \dots, \tilde{\alpha}_g \mid \tilde{\alpha}_{g+1} \tilde{\mathbf{s}}_2^T)$$

(not bad)



# Iterative Aggregation

Improve By Successive Aggregation / Disaggregation?

NO

Can't do A/D twice — a fixed point emerges

## Solution

Perturb A/D output to move off of fixed point

Move it in direction of solution

$$\tilde{\pi}^T = \tilde{\pi}^T \mathbf{P}$$

(a smoothing step)

## The Iterative A/D Updating Algorithm

Determine the “ $G$ -set” partition  $\mathcal{S} = G \cup \overline{G}$

Approximate A/D step generates approximation  $\tilde{\pi}^T$

Smooth the result  $\tilde{\pi}^T = \tilde{\pi}^T \mathbf{P}$

Use  $\tilde{\pi}^T$  as input to another approximate aggregation step

⋮



# How to Partition for Updating Problem?

## Intuition

- There are some bad states ( $G$ ) and some good states ( $\overline{G}$ ).
- Give more attention to bad states. Each state in  $G$  forms a partitioning level. Much progress toward correct PageRank is made during aggregation step.
- Lump good states in  $\overline{G}$  into 1 superstate. Progress toward correct PageRank is made during smoothing step (power iteration).



# Definitions for “Good” and “Bad”

1. Good = states least likely to have  $\pi_i$  change  
Bad = states most likely to have  $\pi_i$  change
2. Good = states with smallest  $\pi_i$  after  $k$  transient steps  
Bad = states “nearby”, with largest  $\pi_i$  after  $k$  transient steps
3. Good = smallest  $\pi_i$  from old PageRank vector  
Bad = largest  $\pi_i$  from old PageRank vector
4. Good = **fast**–converging states  
Bad = **slow**–converging states



# Determining “Fast” and “Slow”

## Consider power method and its rate of convergence

$$\pi_{k+1}^T = \pi_k^T \mathbf{P} = \pi_k^T \mathbf{e} \pi^T + \lambda_2^k \pi_k^T \mathbf{x}_2 \mathbf{y}_2^T + \lambda_3^k \pi_k^T \mathbf{x}_3 \mathbf{y}_3^T + \cdots + \lambda_n^k \pi_k^T \mathbf{x}_n \mathbf{y}_n^T$$

Asymptotic rate of convergence is rate at which  $\lambda_2^k \rightarrow 0$

## Consider convergence of elements

Some states converge to stationary value faster than  $\lambda_2$ -rate, due to LH e-vector  $\mathbf{y}_2^T$ .

## Partitioning Rule

Put states with largest  $|\mathbf{y}_2^T|_i$  values in bad group  $G$ , where they receive more individual attention in aggregation method.

## Practicality

$\mathbf{y}_2^T$  expensive, but for PageRank problem, Kamvar et al. show states with large  $\pi_i$  are slow-converging.  $\Rightarrow$  inexpensive soln = use old  $\pi^T$  to determine  $G$ .  
(adaptively approximate  $\mathbf{y}_2^T$ )



# Implications of Web's scale-free nature

## Facts:

(1)  $\pi^T$  follows power law since WWW is scale-free

(experimental and theoretical justification)

(2) not all pages converge to their PageRanks at same rate

(3) pages with high PR are slow-converging


⇒ very few pages are slow-converging, but these are the pages that cause power method to drag on



# Power law for PageRank

Scale-free Model of Web network creates power laws

(Kamvar, Barabasi, Raghavan)







# Convergence

## Theorem

Always converges to stationary dist  $\pi^T$  for  $\mathbf{P}$

Converges for all partitions  $\mathcal{S} = G \cup \overline{G}$

Rate of convergence is rate at which  $\mathbf{S}_2^n$  converges

$$\mathbf{S}_2 = \mathbf{P}_{22} + \mathbf{P}_{21}(\mathbf{I} - \mathbf{P}_{11})^{-1}\mathbf{P}_{12}$$

Dictated by Jordan structure of  $\lambda_2(\mathbf{S}_2)$

$\lambda_2(\mathbf{S}_2)$  simple  $\implies \pi_k^T \rightarrow \pi^T$  at the rate at which  $\lambda_2^n \rightarrow 0$

## The Game

Goal now is to find a relatively small  $G$  that minimizes  $\lambda_2(\mathbf{S}_2)$



# Ipsen/Kirkland Updating Theory

## Motivation

- L/M prove updating method converges at rate  $(\lambda_2(\mathbf{S}_2))^k \rightarrow 0$ .
- Ipsen/Kirkland wonder: can  $\lambda_2(\mathbf{S}_2) > \alpha$  ?

## Results

- $\lambda_2(\mathbf{S}_2) \leq \alpha$  for all partitions.
- $\lambda_2(\mathbf{S}_2) < \alpha$  under two trivial assumptions on  $\mathbf{P}$ .  
( $\mathbf{P}$  is reducible, and at least one page in each essential class does not self-link)



# Ipsen/Kirkland Updating Theory

## Motivation

- L/M prove updating method converges at rate  $(\lambda_2(\mathbf{S}_2))^k \rightarrow 0$ .
- Ipsen/Kirkland wonder: can  $\lambda_2(\mathbf{S}_2) > \alpha$  ?

## Results

- $\lambda_2(\mathbf{S}_2) \leq \alpha$  for all partitions.
- $\lambda_2(\mathbf{S}_2) < \alpha$  under two trivial assumptions on  $\mathbf{P}$ .  
( $\mathbf{P}$  is reducible, and at least one page in each essential class does not self-link)

But ... how do we find partition so that  $\lambda_2(\mathbf{S}_2) \ll \alpha$  ?



# Experiments

## Test Networks From Crawl Of Web

NCState

(NCSU internal crawl)

10,000 nodes    101,118 links

California

(Sites concerning “california” query)

9,664 nodes    16,150 links



# Parameters

**Number Of Nodes (States) Added**

50

**Number Of Nodes (States) Removed**

30

**Number Of Links Added**

(Different values have little effect on results)

300

**Number Of Links Removed**

200

**Stopping Criterion**

1-norm of residual  $< 10^{-10}$



# NC State

## Power Method

<u>Iterations</u>	<u>Time</u>
162	9.79

## Iterative Aggregation

<u><math> G </math></u>	<u>Iterations</u>	<u>Time</u>
500	160	10.18
1000	51	3.92
1500	33	2.82
2500	16	2.15
3000	13	1.99
5000	7	1.77

*nodes = 10,000    links = 101,118*



# NC State

## Power Method

Iterations	Time
162	9.79

## Iterative Aggregation

$ G $	Iterations	Time
500	160	10.18
1000	51	3.92
1500	33	2.82
<b>2000</b>	<b>21</b>	<b>2.22</b>
2500	16	2.15
3000	13	1.99
5000	7	1.77

*nodes = 10,000 links = 101,118*



# California

## Power Method

Iterations	Time
176	5.85

## Iterative Aggregation

$ G $	Iterations	Time
500	19	1.12
1000	15	.92
1250	20	1.04
2000	13	1.17
5000	6	1.25

*nodes = 9,664 links = 16,150*





# California

## Power Method

Iterations	Time
176	5.85

## Iterative Aggregation

$ G $	Iterations	Time
500	19	1.12
1000	15	.92
1250	20	1.04
<b>1500</b>	<b>14</b>	<b>.90</b>
2000	13	1.17
5000	6	1.25

*nodes = 9,664 links = 16,150*



# Advantage

— updating algorithm can be combined with other PR acceleration methods.


**Power**      **Power+Quad(10)**      **Iter. Agg.**      **Iter.Agg.+Quad(10)**

Iter.	Time	Iter.	Time	$ G $	Iter.	Time	Iter.	Time
162	9.69	81	5.93	500	160	10.18	57	5.25
				1000	51	3.92	31	2.87
				1500	33	2.82	23	2.38
				<b>2000</b>	<b>21</b>	<b>2.22</b>	<b>16</b>	<b>1.85</b>
				2500	16	2.15	12	1.88
				3000	13	1.99	11	1.91
				5000	7	1.77	6	1.86

*nodes = 10,000*    *links = 101,118*



# Residual Plot for NC State





# Large-Scale Implementation

## Partitioning

- need more theoretical work on good partitioning.

## IAD's Aggregated System Solve

- direct vs. sparse methods

## Simulating updates to Web

- how to do this accurately, and keep scale-free properties of web
- need collections of the web over time.



# Conclusions

- Aggregation methods reduce PageRank computation for the updating problem. However, partitioning is a difficult, unresolved issue.
- Many of the acceleration methods can be combined to achieve even greater speedups.
- We are moving closer to lofty goal of computing real-time personalized PageRank.