



Local clustering with graph diffusions and spectral solution paths




Joint with
David F. Gleich,
(Purdue), supported by
NSF CAREER
1149756-CCF

PURDUE
UNIVERSITY

Kyle Kloster
Purdue University

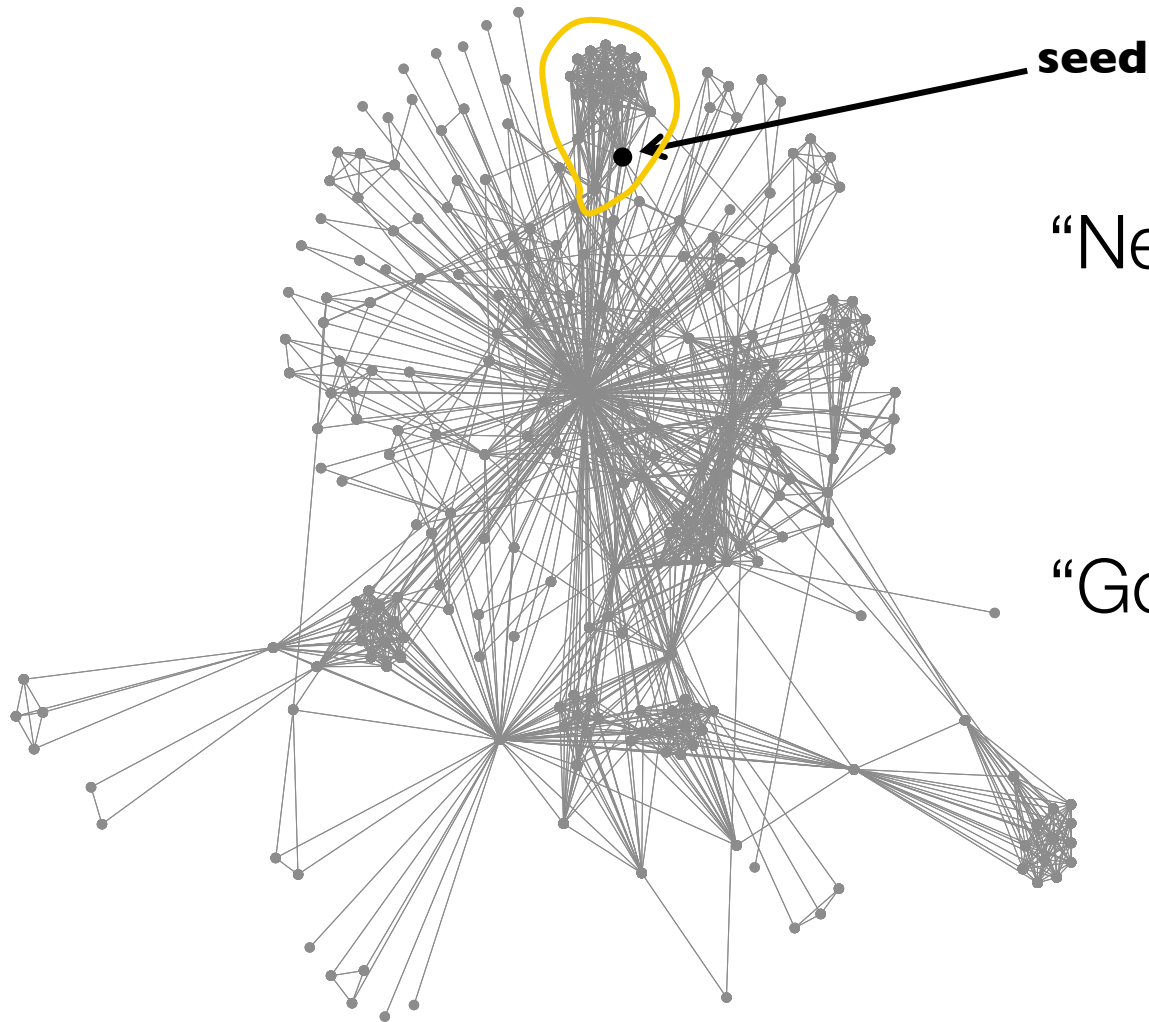
Local Clustering

Given seed(s) S in G , find a good cluster *near* S



Local Clustering

Given seed(s) S in G , find a good cluster *near* S




“Near”? ->
local, small
containing S

“Good”? ->
low conductance

Low-conductance sets are clusters

$$\text{conductance}(T) = \frac{\# \text{ edges leaving } T}{\# \text{ edge endpoints in } T}$$

(for small sets T ,
i.e. $\text{vol}(T) < \text{vol}(G)/2$)



= “chance a random edge
that touches T exits T ”


Low-conductance sets are clusters

$$\text{conductance}(T) = \frac{\# \text{ edges leaving } T}{\# \text{ edge endpoints in } T}$$

(for small sets T ,
i.e. $\text{vol}(T) < \text{vol}(G)/2$)

For a global cluster,
could use Fiedler...

**But we want
a local cluster**



Fiedler

Compute Fiedler vector, \mathbf{v} :

$$L\mathbf{v} = \lambda_2 D\mathbf{v}$$

“Sweep” over \mathbf{v} :

1. sort:

$$v(1) \geq v(2) \geq \dots$$

2. for each set $S_k = (1, \dots, k)$
compute conductance

$$\phi(S_k)$$

3. output best S_k

Fiedler

Compute Fiedler vector, \mathbf{v} :

$$L\mathbf{v} = \lambda_2 D\mathbf{v}$$

“Sweep” over \mathbf{v} :

1. sort:

$$v(1) \geq v(2) \geq \dots$$

2. for each set $S_k = (1, \dots, k)$
compute conductance

$$\phi(S_k)$$

3. output best S_k

Cheeger Inequality:
Fiedler finds a cluster “not too much worse” than global optimal

But we want **local...**

Local Fiedler and diffusions

[Mahoney Orecchia Vishnoi 12]

“A local spectral method...”

$$L\mathbf{v} = D\mathbf{v}[\lambda] \quad \text{Fiedler}$$

$$L\mathbf{v} = D\mathbf{v}[\lambda] + \mathbf{s} \quad \text{with local bias} \quad (\text{MOV})$$

(normalized seed vector \mathbf{s})

THM: MOV is a scaling of personalized PageRank*!

Local Fiedler and diffusions

Intuition: why MOV ~ PageRank

$$L\mathbf{v} = D\mathbf{v}[\lambda] \quad \text{Fiedler}$$

$$L\mathbf{v} = D\mathbf{v}[\lambda] + \text{“s”} \quad \text{with local bias}$$

$$(I - D^{-1/2}AD^{-1/2})\hat{\mathbf{v}} = \hat{\mathbf{v}}[\lambda] + \text{“s”}$$

$$AD^{-1}\hat{\mathbf{v}} = \hat{\mathbf{v}}[1 - \lambda] + \text{“s”}$$

$$(I - \alpha P)\hat{\mathbf{v}} = \text{“s”} \quad \text{PageRank vector, a diffusion}$$

PageRank and other diffusions

“Personalized” PageRank (PPR)

[Andersen, Chung, Lang 06]: **local** Cheeger inequality and fast algorithm, “Push” procedure

Standard setting

$$(I - \alpha \mathbf{P}) \mathbf{x} = \hat{\mathbf{s}} \longrightarrow$$

Diffusion perspective

$$\mathbf{x} = \sum_{k=0} \alpha^k \mathbf{P}^k \hat{\mathbf{s}}$$

PageRank and other diffusions


$$\mathbf{x} = \sum_{k=0} \alpha^k \mathbf{P}^k \hat{\mathbf{s}}$$

“Personalized” PageRank (PPR)

[Andersen, Chung, Lang 06]: local Cheeger inequality and fast algorithm, “Push” procedure

Heat Kernel diffusion (HK)
(many more!)

$$\mathbf{f} = \sum_{k=0} \frac{t^k}{k!} \mathbf{P}^k \hat{\mathbf{s}}$$



Various diffusions explore different aspects of graphs.

Diffusions, theory & practice

good
conductance

fast
algorithm

PR

Local Cheeger Inequality

[Andersen Chung Lang 06]
“PPR-push” is $O(1/(\epsilon(1-\alpha)))$

HK

Local Cheeger Inequality
[Chung 07]

[K., Gleich 2014]
“HK-push” is $O(e^t C/\epsilon)$

TDPR

Open question

[Avron, Horesh 2015]

Gen
Diff

Open question

This talk

Diffusions, theory & practice

good
conductance

fast
algorithm

PR

Local Cheeger Inequality

[Andersen Chung Lang 06]
“PPR-push” is $O(1/(\epsilon(1-\alpha)))$

HK

Local Cheeger Inequality
[Chung 07]

[K., Gleich 2014]
“HK-push” is $O(e^t C/\epsilon)$

TDPR

Open question

[Avron, Horesh 2015]

Gen
Diff

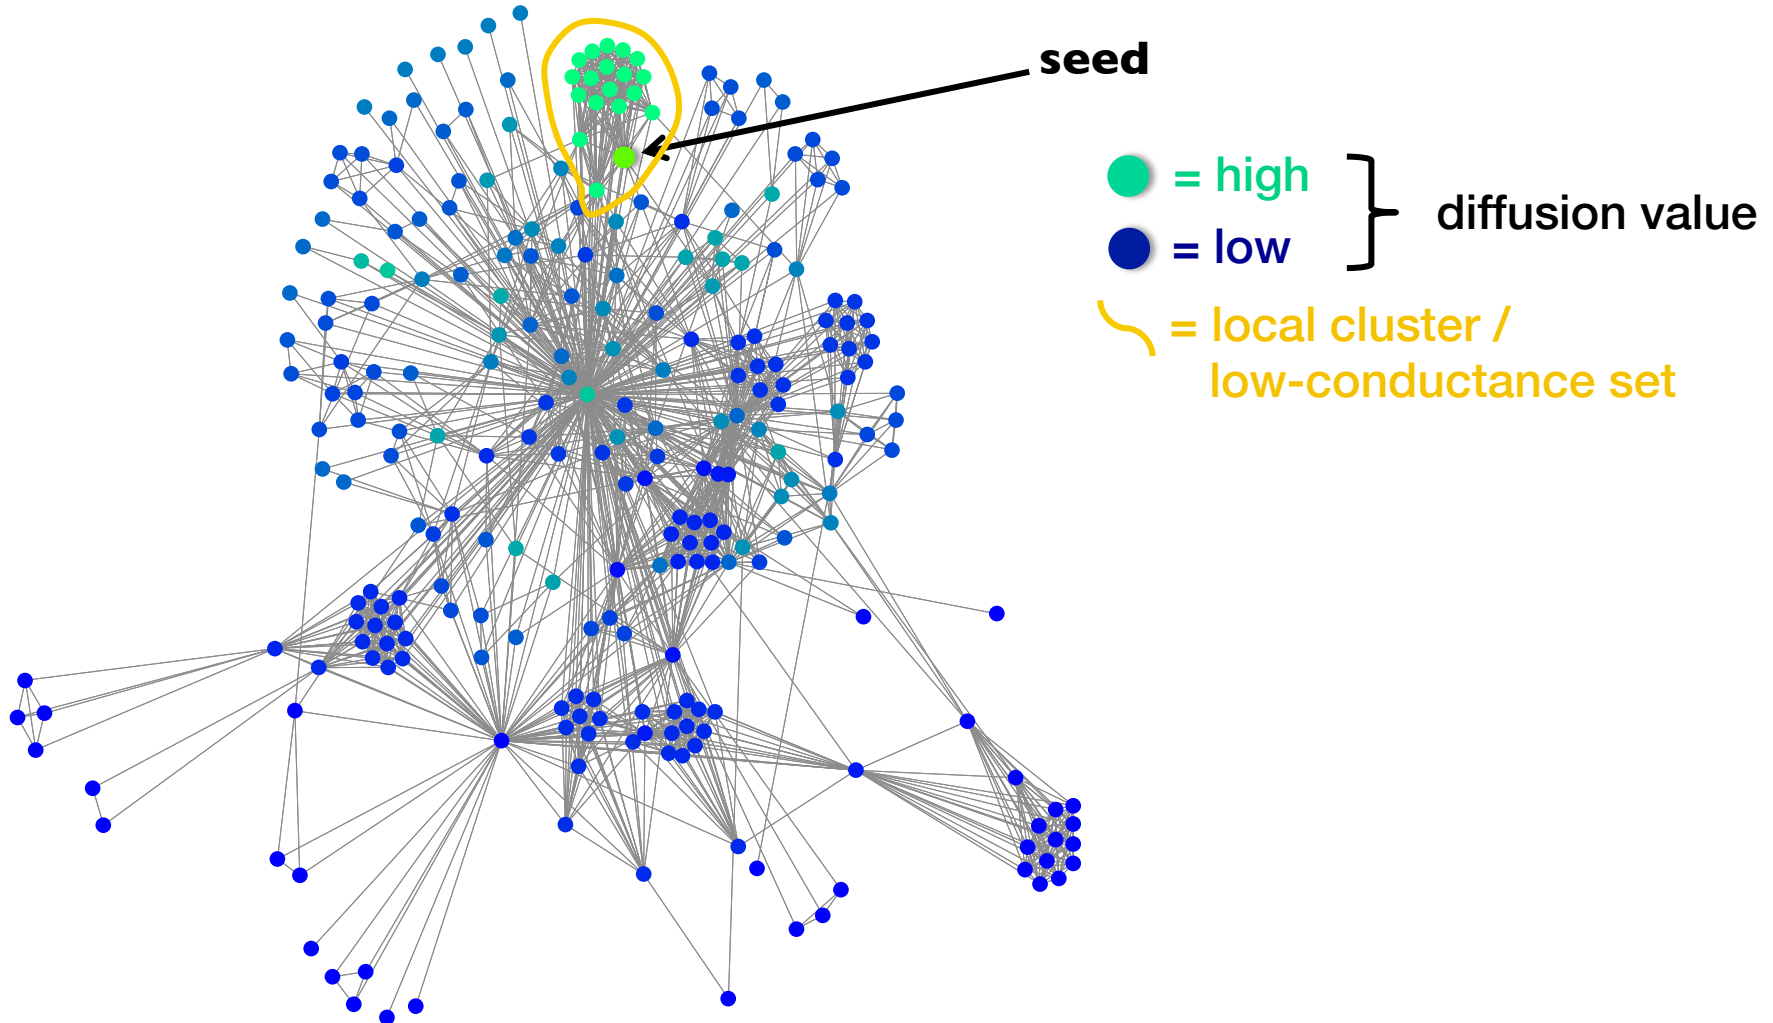
Open question

This talk

David Gleich and I are working with Olivia Simpson
(a student of Fan Chung's)

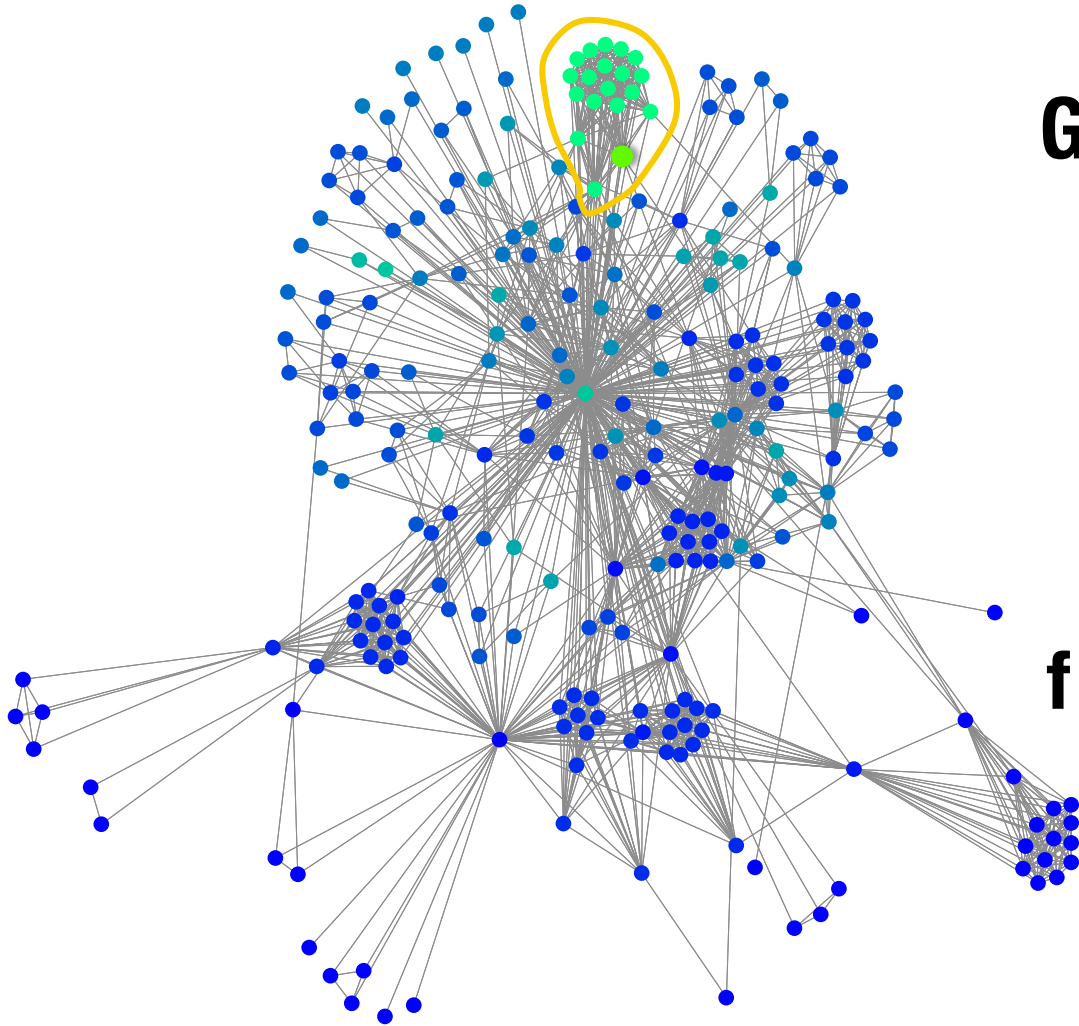
General diffusions: intuition

A diffusion propagates “rank” from a seed across a graph.



General diffusions

A diffusion propagates “rank” from a seed across a graph.



General diffusion vector

$$\mathbf{f} = \sum_{k=0} c_k \mathbf{P}^k \hat{\mathbf{s}}$$

$$\mathbf{f} = c_0 \mathbf{p}_0 + c_1 \mathbf{p}_1 + c_2 \mathbf{p}_2 + c_3 \mathbf{p}_3 + \dots$$

Sweep over \mathbf{f} !

General algorithm


1. Approximate \mathbf{f} so $\|\mathbf{D}^{-1}(\mathbf{f} - \hat{\mathbf{f}})\|_{\infty} \leq \epsilon$
2. Scale, $\mathbf{D}^{-1}\hat{\mathbf{f}}$
3. Then sweep!

How to do this efficiently?


Algorithm Intuition

From parameters c_k , ε , seed \mathbf{s} ...

Starting from here...




How to end up here?



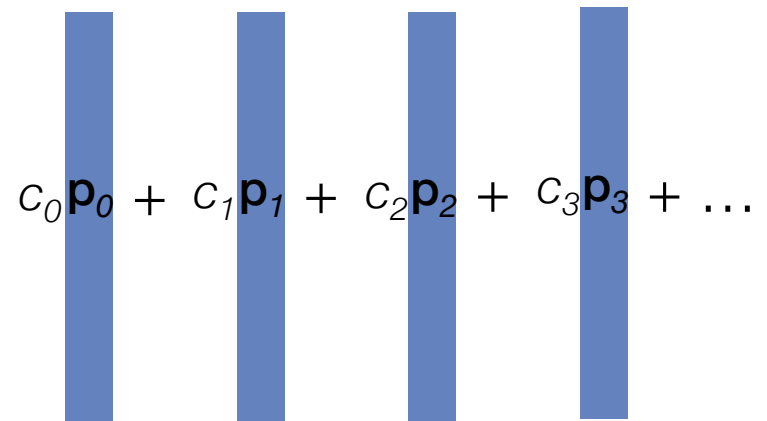
Algorithm Intuition

Begin with mass at seed(s)
in a “residual” staging area, r_0

The residuals r_k hold mass that
is unprocessed – it’s like *error*




Idea: “push” any entry
 $r_k(j) / d_j > (\text{some threshold})$



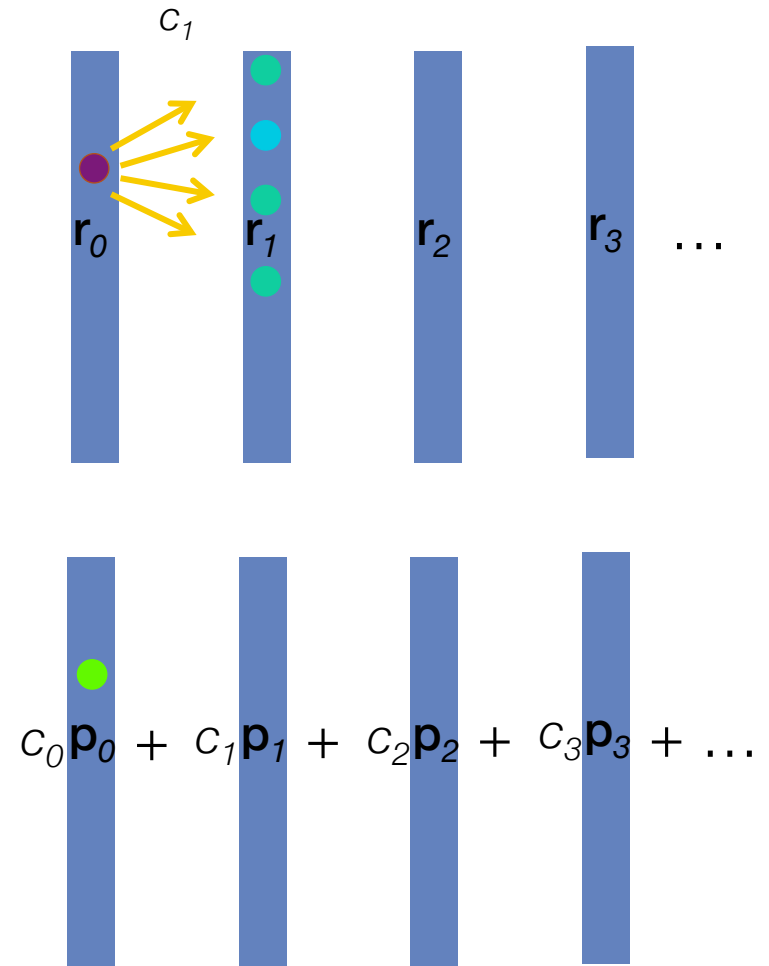
Push Operation

push – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{f} ,



Push Operation


- push** – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{f} ,
(3) then scale and spread to neighbors in next \mathbf{r}



Push Operation

push – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{f} ,
(3) then scale and
spread to neighbors
in next \mathbf{r}



(repeat)



Push Operation

push – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{f} ,
(3) then scale and
spread to neighbors
in next \mathbf{r}



(repeat)



Push Operation

push – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{f} ,
(3) then scale and
spread to neighbors
in next \mathbf{r}

(repeat)




Thresholds

ERROR equals weighted sum of entries left in \mathbf{r}_k

→ Set threshold so “leftovers” sum to $< \varepsilon$

entries $<$ threshold



$$c_0 \mathbf{p}_0 + c_1 \mathbf{p}_1 + c_2 \mathbf{p}_2 + c_3 \mathbf{p}_3 + \dots$$

The diagram shows the weighted sum of vectors $\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \dots$ represented as vertical blue bars. The coefficients $c_0, c_1, c_2, c_3, \dots$ are shown to the left of each bar. The vectors \mathbf{p}_k contain colored dots (green, cyan, blue) representing the entries.

Thresholds

ERROR equals weighted sum of entries left in \mathbf{r}_k



→ Set threshold so “leftovers” sum to $< \epsilon$

Threshold for stage \mathbf{r}_k is

$$\epsilon / \left(\sum_{j=k+1}^{\infty} c_j \right)$$

Then $\| \mathbf{D}^{-1} (\mathbf{f} - \hat{\mathbf{f}}) \|_{\infty} \leq \epsilon$

entries $<$ threshold



Another perspective

$$L\mathbf{v} = D\mathbf{v}[\lambda] \quad \text{Fiedler}$$

$$L\mathbf{v} = D\mathbf{v}[\lambda] + \text{“S”} \quad \text{with local bias}$$

$$(I - D^{-1/2}AD^{-1/2})\hat{\mathbf{v}} = \hat{\mathbf{v}}[\lambda] + \text{“S”}$$

$$AD^{-1}\hat{\mathbf{v}} = \hat{\mathbf{v}}[1 - \lambda] + \text{“S”}$$

$$(I - \alpha P)\hat{\mathbf{v}} = \text{“S”} \quad \text{PageRank vector, a diffusion}$$

Another perspective

$$LV_k = DV_k \Lambda_k \quad \text{Fiedler}$$

$$LV_k = DV_k \Lambda_k + \mathbf{S} \quad \text{with local bias}$$

$$(I - D^{-1/2} A D^{-1/2}) \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k \Lambda_k + \hat{\mathbf{S}}$$

$$A D^{-1} \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k (I - \Lambda_k) + \hat{\mathbf{S}}$$

Another perspective

$$LV_k = DV_k \Lambda_k \quad \text{Fiedler}$$

$$LV_k = DV_k \Lambda_k + \mathbf{S} \quad \text{with local bias}$$

$$(I - D^{-1/2} A D^{-1/2}) \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k \Lambda_k + \hat{\mathbf{S}}$$

$$A D^{-1} \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k (I - \Lambda_k) + \hat{\mathbf{S}}$$

$$P \hat{\mathbf{V}}_k \Gamma = \hat{\mathbf{V}}_k + \hat{\mathbf{S}}$$

Mix-product property
For Kronecker product

Another perspective

$$\mathbf{LV}_k = \mathbf{DV}_k \Lambda_k \quad \text{Fiedler}$$

$$\mathbf{LV}_k = \mathbf{DV}_k \Lambda_k + \mathbf{S} \quad \text{with local bias}$$

$$(\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}) \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k \Lambda_k + \hat{\mathbf{S}}$$

$$\mathbf{A} \mathbf{D}^{-1} \hat{\mathbf{V}}_k = \hat{\mathbf{V}}_k (\mathbf{I} - \Lambda_k) + \hat{\mathbf{S}}$$

$$\mathbf{P} \hat{\mathbf{V}}_k \Gamma = \hat{\mathbf{V}}_k + \tilde{\mathbf{S}}$$

Mix-product property
For Kronecker product

$$(\mathbf{I} - \Gamma^T \otimes \mathbf{P}) \text{vec}(\hat{\mathbf{V}}_k) = \text{vec}(\tilde{\mathbf{S}})$$

Another perspective

$$\left. \begin{aligned} (\mathbf{I} - \Gamma^T \otimes \mathbf{P}) \text{vec}(\hat{\mathbf{V}}_k) &= \text{vec}(\tilde{\mathbf{S}}) \\ (\mathbf{I} - \alpha \mathbf{P}) \hat{\mathbf{v}} &= \tilde{\mathbf{s}} \end{aligned} \right\}$$

- generalizes PageRank to “matrix teleportation parameter”

Standard spectral approach: $\Gamma = (\mathbf{I} - \mathbf{\Lambda}_k)^{-1}$

Another perspective

$$\left. \begin{aligned} (\mathbf{I} - \Gamma^T \otimes \mathbf{P}) \text{vec}(\hat{\mathbf{V}}_k) &= \text{vec}(\tilde{\mathbf{S}}) \\ (\mathbf{I} - \alpha \mathbf{P}) \hat{\mathbf{v}} &= \tilde{\mathbf{s}} \end{aligned} \right\}$$

- generalizes PageRank to “matrix teleportation parameter”

Our framework
is equivalent to:

$$\Gamma = \begin{bmatrix} 0 & \tilde{c}_0 & & \\ & 0 & \ddots & \\ & & \ddots & \\ & & & \tilde{c}_N \\ & & & & 0 \end{bmatrix}$$

(Details in [K., Gleich KDD 14])

General diffusions: conclusion

THM: For diffusion coefficients $c_k \geq 0$ satisfying

$$\left. \begin{array}{l} \sum_{k=0}^{\infty} c_k = 1 \quad \text{and} \quad \sum_{k=0}^N c_k \leq \epsilon/2 \end{array} \right\} \text{“rate of decay”}$$

“generalized push” approximates the diffusion \mathbf{f}

on a symmetric graph so that $\|\mathbf{D}^{-1}(\mathbf{f} - \hat{\mathbf{f}})\|_{\infty} \leq \epsilon$

in work bounded by $O(2N^2/\epsilon)$

Constant for any inputs!
(If diffusion decays fast)

Proof sketch

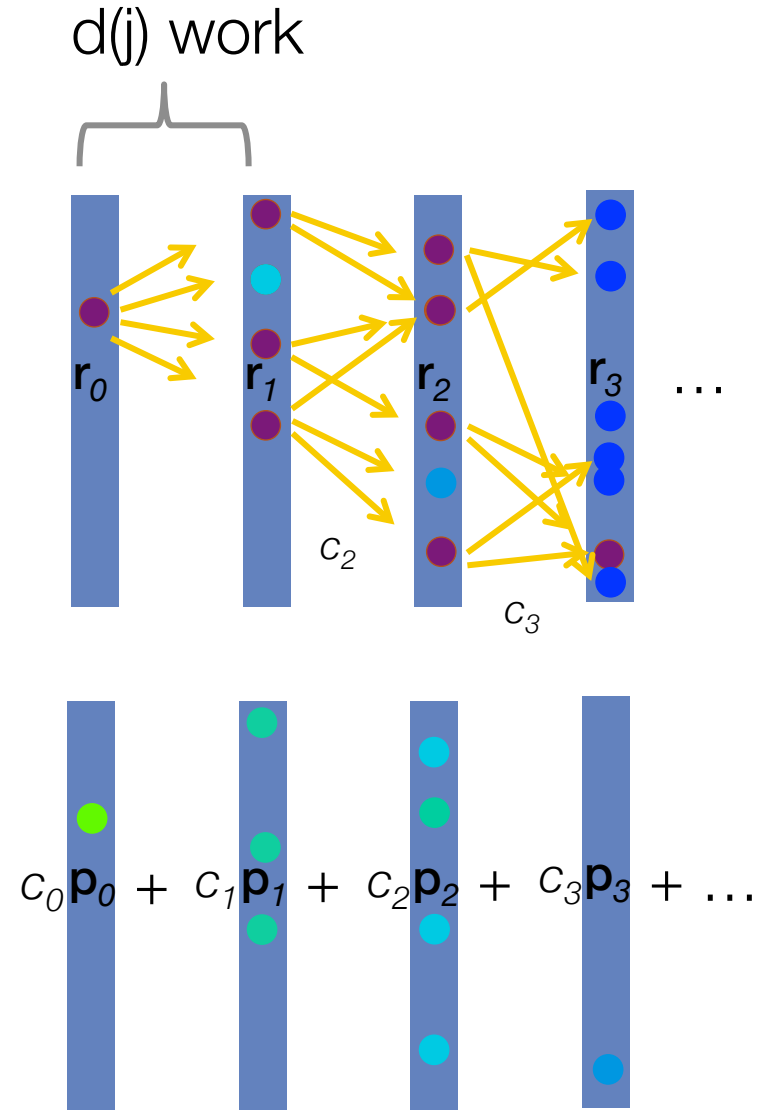
1. Stop pushing after N terms. $\sum_{k=0}^N c_k \leq \epsilon/2$

2. Push residual entries in first N terms if $r_k(j) \geq d(j)\epsilon/(2N)$

3. Total work is # pushes: $\sum_{k=0}^{N-1} \sum_{t=1}^{m_k} d(j_t)$

Push Recap

- push** – (1) remove entry in \mathbf{r}_k ,
(2) put in \mathbf{p} ,
(3) then scale and spread to neighbors in next \mathbf{r}



Proof sketch

1. Stop pushing after N terms. $\sum_{k=0}^N c_k \leq \epsilon/2$

2. Push residual entries in first N terms if $r_k(j) \geq d(j)\epsilon/(2N)$

3. Total work is # pushes: $\sum_{k=0}^{N-1} \sum_{t=1}^{m_k} d(j_t)$

Proof sketch

1. Stop pushing after N terms. $\sum_{k=0}^N c_k \leq \epsilon/2$

2. Push residual entries in first N terms if $r_k(j) \geq d(j)\epsilon/(2N)$

3. Total work is # pushes: $\sum_{k=0}^{N-1} \sum_{t=1}^{m_k} d(j_t) \leq \sum_{k=0}^{N-1} \sum_{t=1}^{m_k} r_k(j_t)(2N)/\epsilon$

Proof sketch

1. Stop pushing after N terms. $\sum_{k=0}^N c_k \leq \epsilon/2$

2. Push residual entries in first N terms if $r_k(j) \geq d(j)\epsilon/(2N)$

3. Total work is # pushes: $\sum_{k=0}^{N-1} \sum_{t=1}^{m_k} d(j_t) \leq \sum_{k=0}^{N-1} \sum_{t=1}^{m_k} r_k(j_t)(2N)/\epsilon$

4. Each r_k sums to ≤ 1
(each push is added to \mathbf{f} , which sums to 1) $\sum_{t=1}^{m_k} r_k(j_t) \leq 1$

$$O(2N^2/\epsilon)$$

Solutions Paths

Benefit of these “push” diffusions?

A direct decomposition is a black box:
Feed in input, get output.

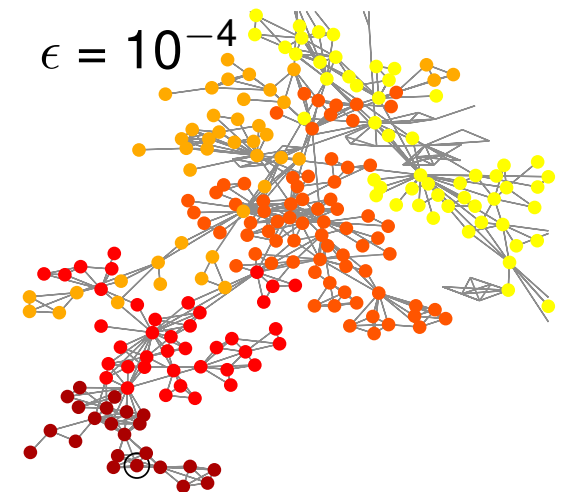


In contrast, the iterative nature of “push” means running the algorithm is essentially “watching” the diffusion process occur.

Solutions Paths

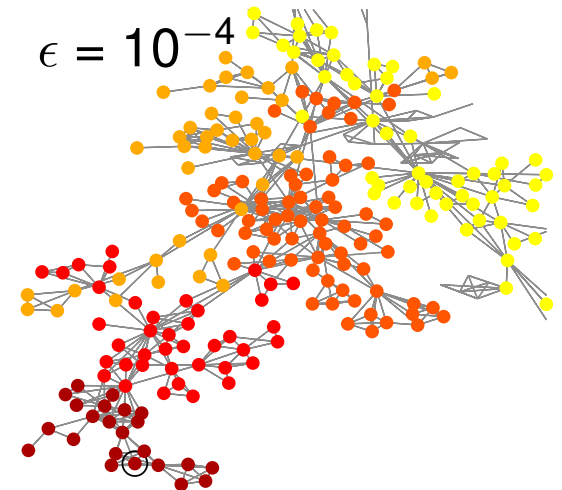

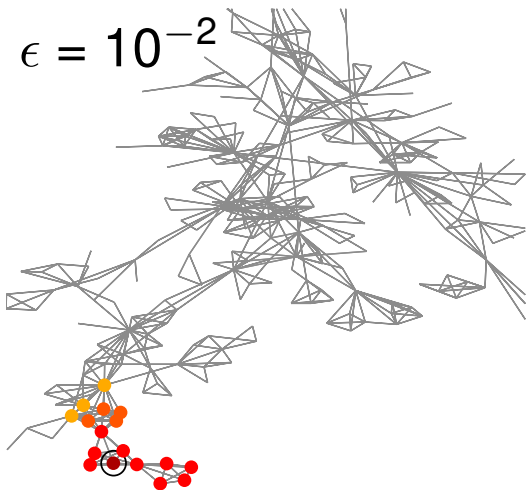

Benefit of these “push” diffusions?

A direct decomposition is a black box:
Feed in input, get output.

In contrast, the iterative nature of “push” means
running the algorithm is essentially “watching” the
diffusion process occur.







Solutions Paths



Solutions Paths

Each curve is a node. Its value increases as ϵ goes to 0.


Thick black line shows set of best conductance.



Solutions Paths


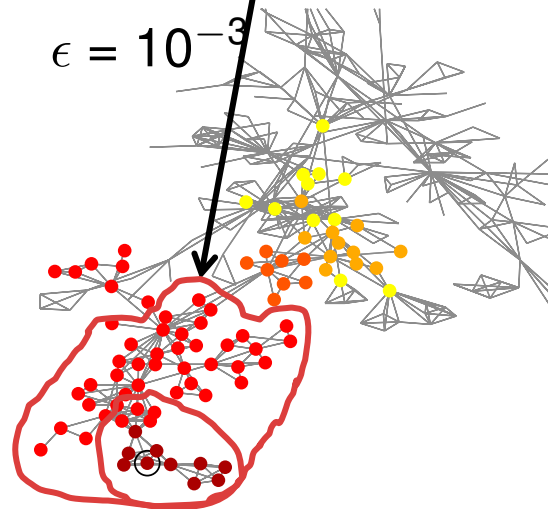

Each curve is a node. Its value increases as ϵ goes to 0.

Thick black line shows set of best conductance.



Bundles of curves are good clusters

Paths identify *nested* clusters



Solutions Paths

Locate nested, good-conductance sets that a single diffusion + sweep could miss.

Can be done efficiently because the constant-time approach to computing diffusions enables efficient storage and analysis of the push process

Total Paths work (for PageRank): $O\left(\frac{1}{\epsilon(1-\alpha)}\right)^2$
Still efficient!

Thank you

Heat kernel code available at

`http://www.cs.purdue.edu/homes/dgleich/codes/hkgrow`

Solution paths: `http://arxiv.org/abs/1503.00322`

(Solution paths, generalized diffusion code soon)

Ongoing work

- Generalized local Cheeger Inequality
for broader class of diffusions

Questions or suggestions? Email Kyle Kloster at `kkloste-at-purdue-dot-edu`