

FINAL: Fast Attributed Network Alignment

Presented by Si Zhang (ASU)



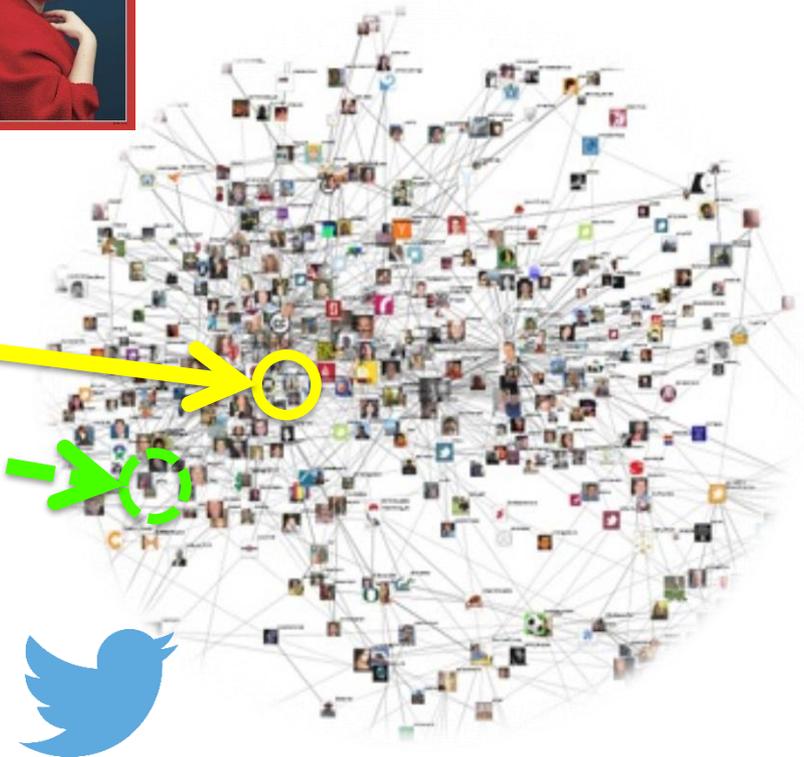
Si Zhang



Hanghang Tong

Why Network Alignment?

- Find Someone Like You

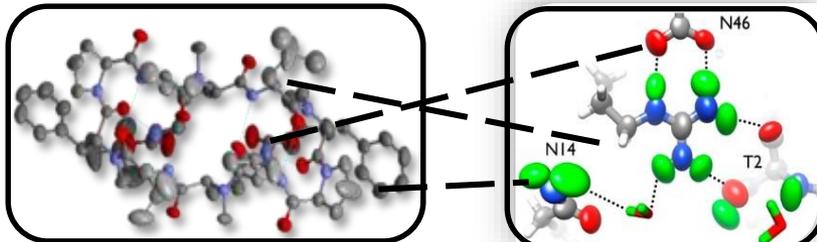


- Q: what if someone lives in a different universe (network)?

More Applications

Identify Species-Specific Pathways [Singh'08]

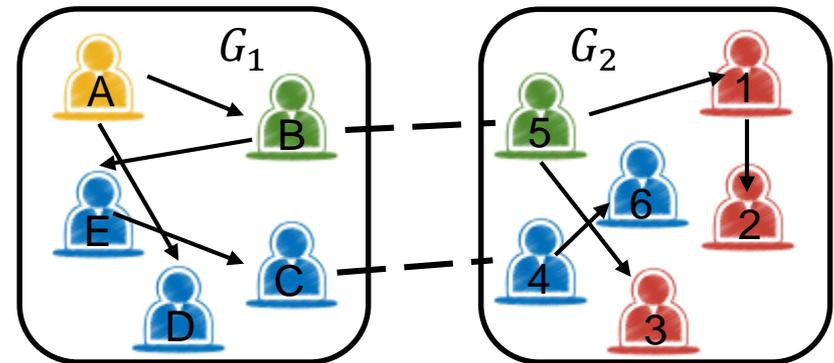
Protein-Protein Interaction (PPI) networks



PPI network 1

PPI network 2

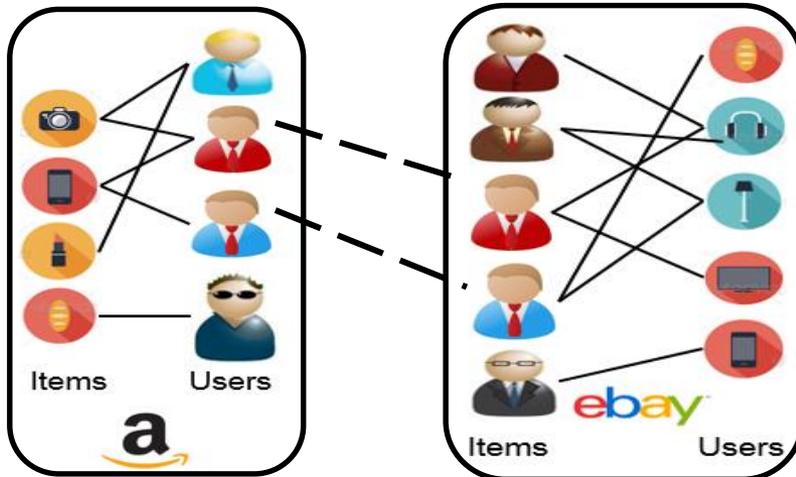
Cross Network Information Diffusion [Zhan'16]



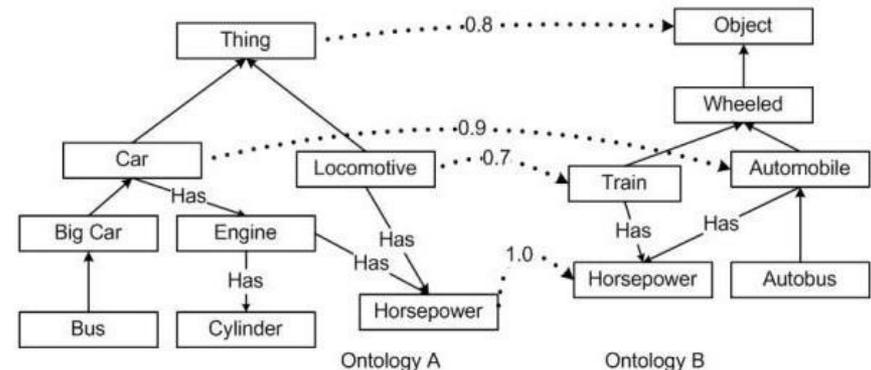
social network 1

social network 2

Cross-Site Recommendation [Zhang'14]



Ontology Matching on Semantic Web [Doan'02]



Why Network Alignment: How to

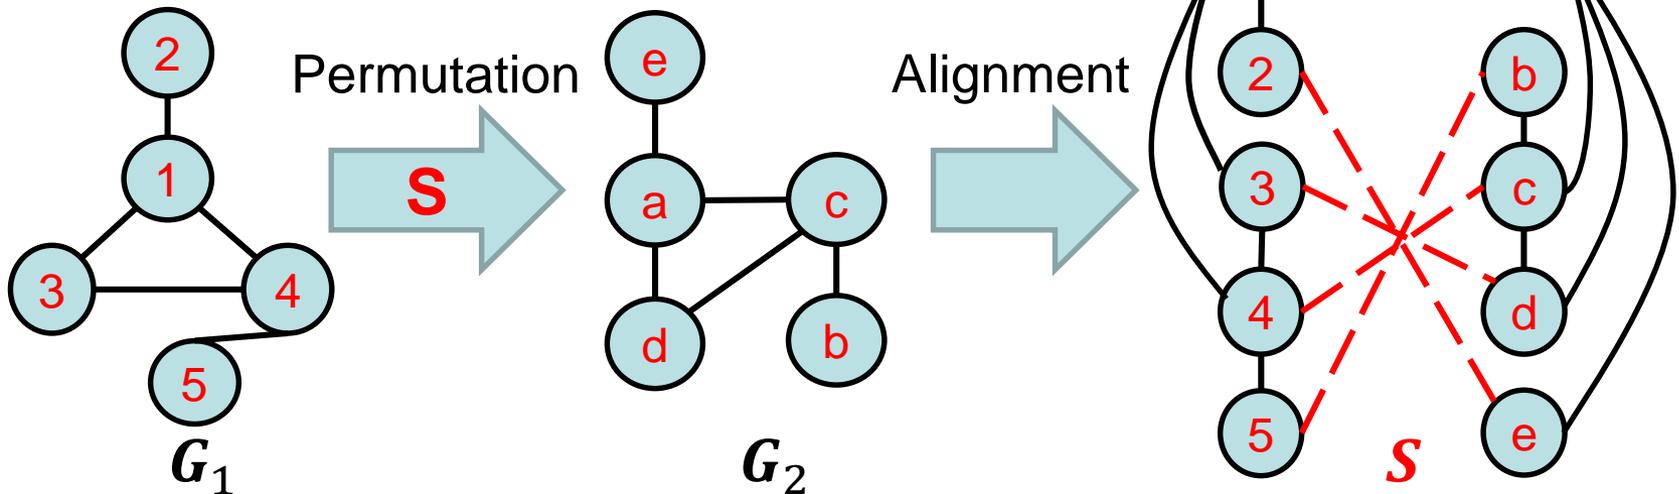
- Existing Methods

- IsoRank [Singh'08], NetAlign [Bayati'09], BigAlign [Koutra'13], UMA [Zhang'15]

- Key Idea: topological consistency

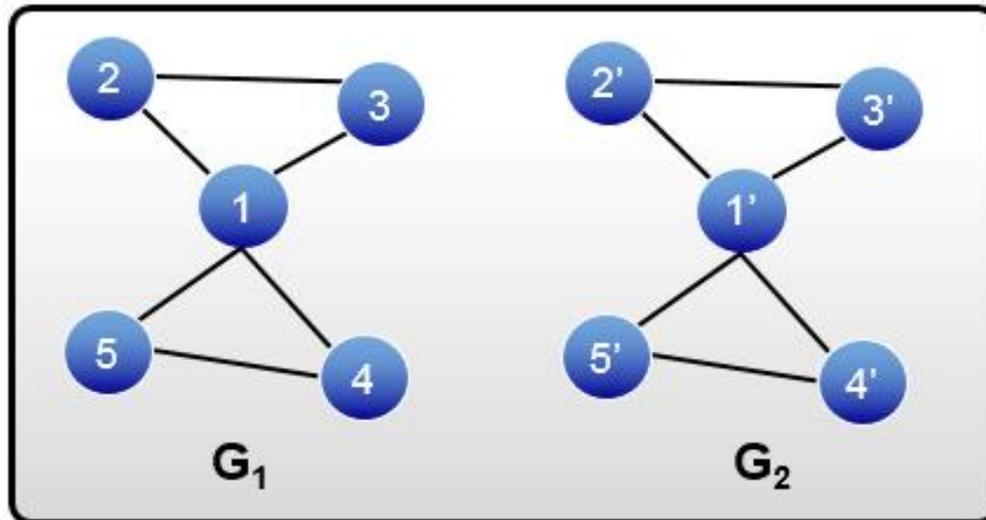
- Network G_2 is a noisy permutation of network G_1

- $G_2 \approx S^T G_1 S$



Topology Consistency: Limitations

- Topological consistency could be easily violated
 - Same nodes may behave differently across different networks
 - Different nodes may have similar connectivity structures



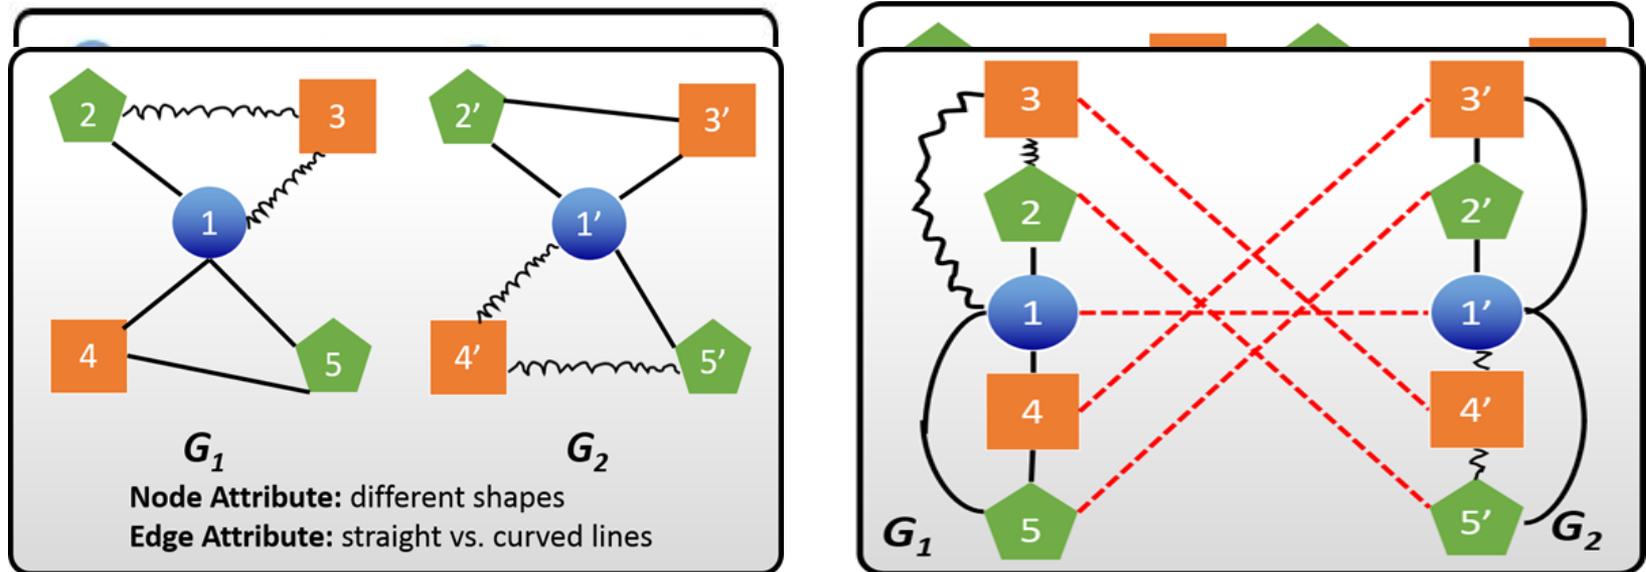
Nodes 2=3=4=5

Nodes 2'=3'=4'=5'

Only topology is not enough!

Topology Consistency: How to Rescue

- Real networks have rich attributes on nodes and/or edges



- Q:** how to calibrate topology-based alignment by leveraging attributes?

Challenges: Attributed Network Alignment

- C1: Formulation
- C2: Optimality
- C3: Scalable Computation

C1. Formulation

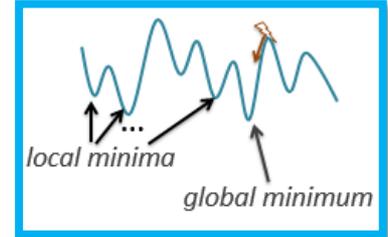
- Typical Network Alignment

NetAlign [Bayati'09]	UMA [Zhang'15]
$\begin{aligned} & \underset{x}{\text{maximize}} \quad \alpha h^T s + \left(\frac{\beta}{2}\right) s^T W s \\ & \text{subject to} \quad A s \leq 1, s_{ii'} \in \{0,1\} \end{aligned}$	$\begin{aligned} & \underset{S}{\text{minimize}} \quad \ S^T A S - B\ _F^2 \\ & \text{s.t.} \quad S \mathbf{1}^{n_2 \times 1} \leq \mathbf{1}^{n_1 \times 1} \\ & \quad \quad S^T \mathbf{1}^{n_1 \times 1} \leq \mathbf{1}^{n_2 \times 1} \end{aligned}$

- **Obs:** only encode topological information
- **Q:** what are their attributed counterparts?

C2. Optimality

- **Obs #1:** many topology-based approaches are non-convex or even NP-hard → find the local minima



- **Obs #2:** attribute may complicate the optimization problem

$$\min_{\mathbf{S}} J(\mathbf{S}) = \sum_{a,b,x,y} \left[\frac{\mathbf{S}(x,a)}{\sqrt{f(x,a)}} - \frac{\mathbf{S}(y,b)}{\sqrt{f(y,b)}} \right]^2 \times \mathbf{A}_1(a,b)\mathbf{A}_2(x,y)$$

without attributes

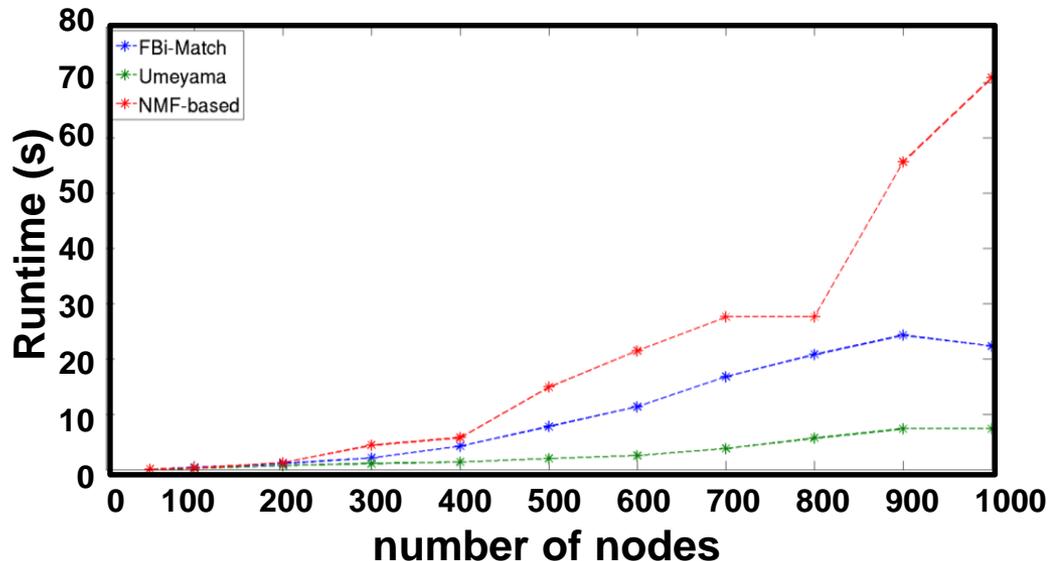
$$\min_{\mathbf{S}} J(\mathbf{S}) = \sum_{a,b,x,y} \left[\frac{\mathbf{S}(x,a)}{\sqrt{f(x,a)}} - \frac{\mathbf{S}(y,b)}{\sqrt{f(y,b)}} \right]^2 \mathbf{A}_1(a,b)\mathbf{A}_2(x,y) \times \mathbb{I}(\mathbf{N}_1(a,a) = \mathbf{N}_2(x,x))\mathbb{I}(\mathbf{N}_1(b,b) = \mathbf{N}_2(y,y)) \times \mathbb{I}(\mathbf{E}_1(a,b) = \mathbf{E}_2(x,y))$$

with attributes

- **Q #1:** what is the exact optimality of the attributed network alignment?
- **Q #2:** how to get the optimal solution, with a comparable complexity (as the topology-alone alignment)?

C3. Scalable Computation

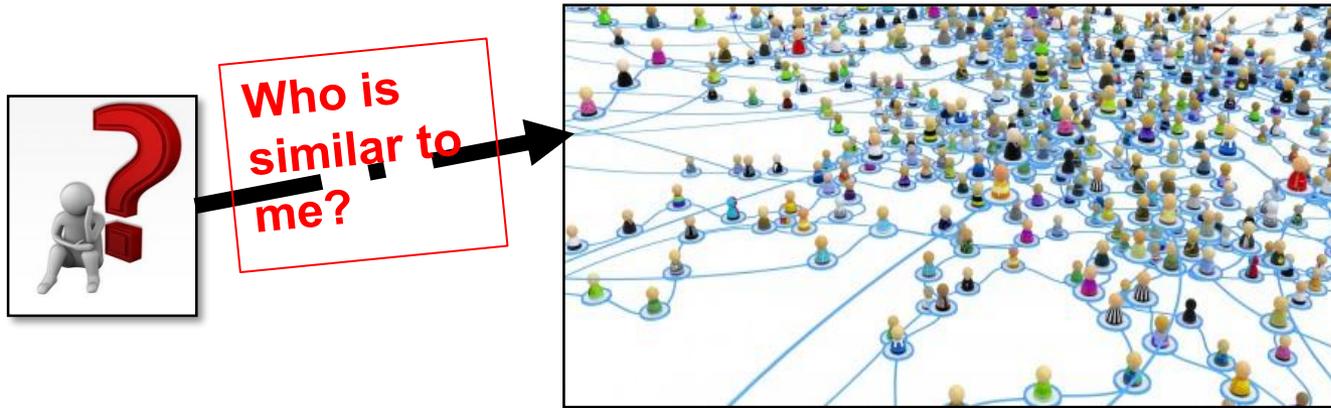
- **Obs #1:** most existing methods have an $O(mn)$ complexity [Singh'08].
- **Obs #2:** best empirical scalability is near-linear [Koutra'13].



- **Q:** how to scale up attributed network alignment?

C3. Scalable Computation

- **Obs:** cross-network search – to find similar users in one network for a given user in another network.



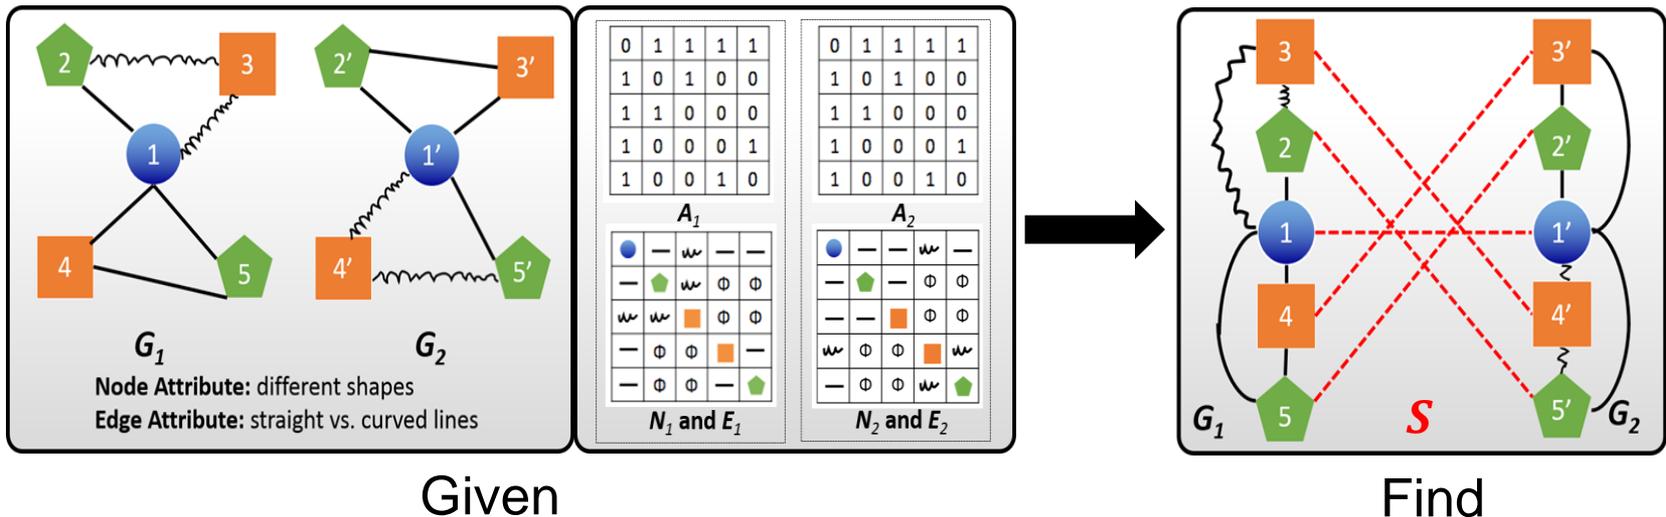
- **Q:** how to speed up on-query network alignment, without solving the full alignment problem?

Outline

- Motivations ✓
- Q1: FINAL Formulation
- Q2: FINAL Algorithms
- Q3: FINAL Speed-up Computation
- Experimental Results
- Conclusions

Prob. Def: Attributed Network Alignment

- **Given:**
 - (1) two attributed networks $G_1 = \{A_1, N_1, E_1\}$ and $G_2 = \{A_2, N_2, E_2\}$;
 - (2 – optional) a prior alignment preference H .
- **Find:** alignment/similarity matrix S
- An Illustrative Example

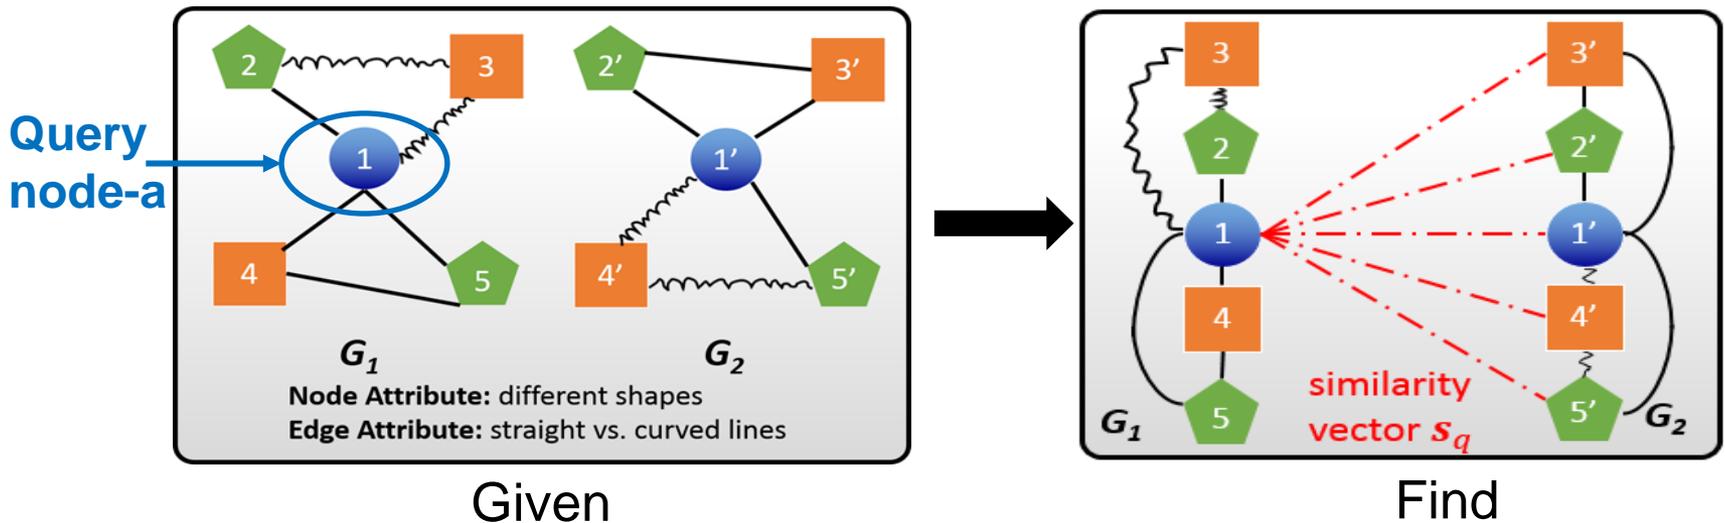


Prob. Def: On-query Attributed Network Alignment

■ Given:

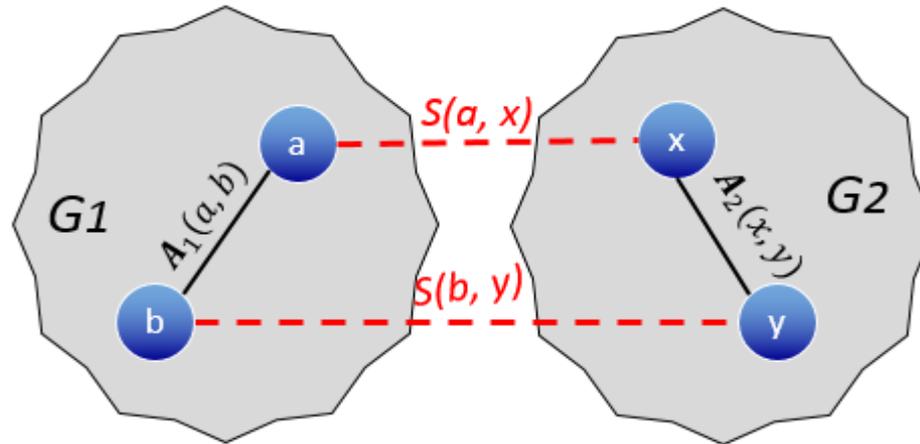
- (1) two attributed networks $G_1 = \{A_1, N_1, E_1\}$ and $G_2 = \{A_2, N_2, E_2\}$;
- (2 – optional) a prior alignment preference H ;
- (3) a query node- a in G_1

■ Find: a vector s_q (similarities of node- a vs. all nodes in G_2)



FINAL Formulation #1: Topological Consistency

- **Intuition:** similar node-pairs tend to have similar neighboring node-pairs



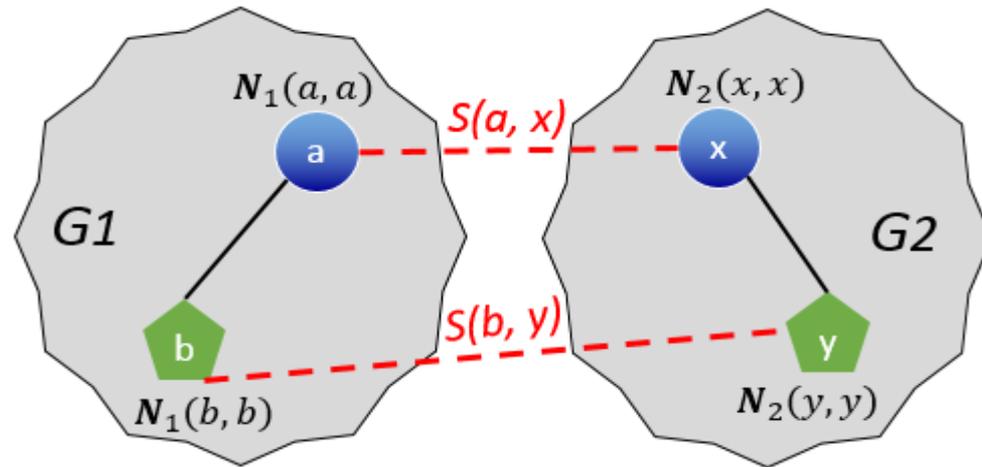
- **Example:**

- large $S(a, x)$
 - large $A_1(a, b)$ and $A_2(x, y)$
- } \longrightarrow large $S(b, y)$

- **Mathematical Details:** $\min_S [S(a, x) - S(b, y)]^2 A_1(a, x) A_2(b, y)$

FINAL Formulation #2: Node Attribute Consistency

- **Intuition:** similar node-pairs share same node attributes



- **Example:**

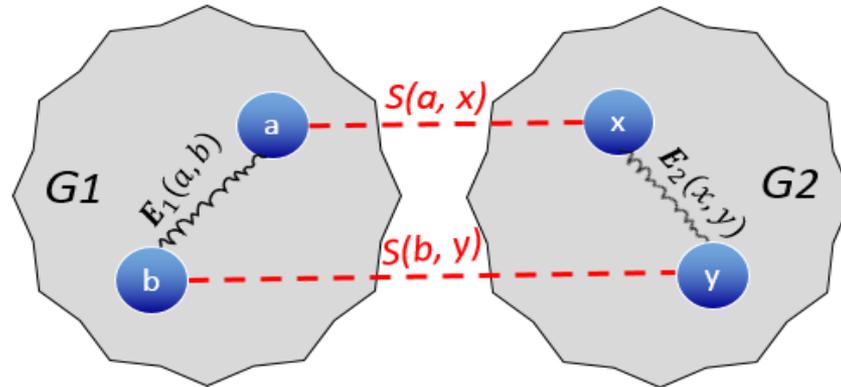
- large $S(a, x)$ \longrightarrow node- a and node- x share same node attribute

- **Mathematical Details:** if $N_1(a, a) = N_2(x, x)$ and $N_1(b, b) = N_2(y, y)$,

$$\min_S [S(a, x) - S(b, y)]^2 A_1(a, x) A_2(b, y)$$

FINAL Formulation #3: Edge Attribute Consistency

- **Intuition:** similar node-pairs connect to their neighbor-pairs via same edge attributes

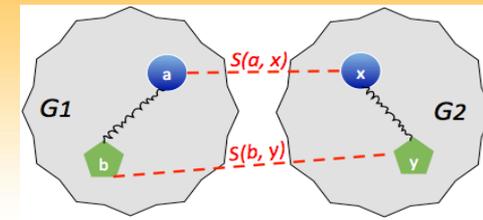


- **Example:**
 - large $S(a,x)$
 - large $S(b,y)$ } \longrightarrow edge (a,b) and (x,y) share same attribute

- **Mathematical Details:** if $E_1(a,b) = E_2(x,y)$,

$$\min_{\mathcal{S}} [\mathcal{S}(a,x) - \mathcal{S}(b,y)]^2 A_1(a,x) A_2(b,y)$$

Putting everything together



■ Objective Function:

$$\min_{\mathbf{S}} J(\mathbf{S}) = \sum_{a,b,x,y} \left[\frac{\mathbf{S}(x,a)}{\sqrt{f(x,a)}} - \frac{\mathbf{S}(y,b)}{\sqrt{f(y,b)}} \right]^2 \underbrace{A_1(a,b)A_2(x,y)}_{\text{\#1. Topology Consistency}}$$

$$\times \underbrace{\mathbb{I}(N_1(a,a) = N_2(x,x))\mathbb{I}(N_1(b,b) = N_2(y,y))}_{\text{\#2. Node Attribute Consistency}}$$

$$\times \underbrace{\mathbb{I}(E_1(a,b) = E_2(x,y))}_{\text{\#3. Edge Attribute Consistency}}$$

■ $f(\mathbf{x}, \mathbf{a})$:

- ‘joint degree’ of node- a and node- x
- normalization to make the optimization problem convex

■ Generalization:

- replacing $\mathbb{I}(\cdot)$ by an attribute similarity function
- can handle numerical attributes on nodes and/or edges.

FINAL Formulation: Matrix Form

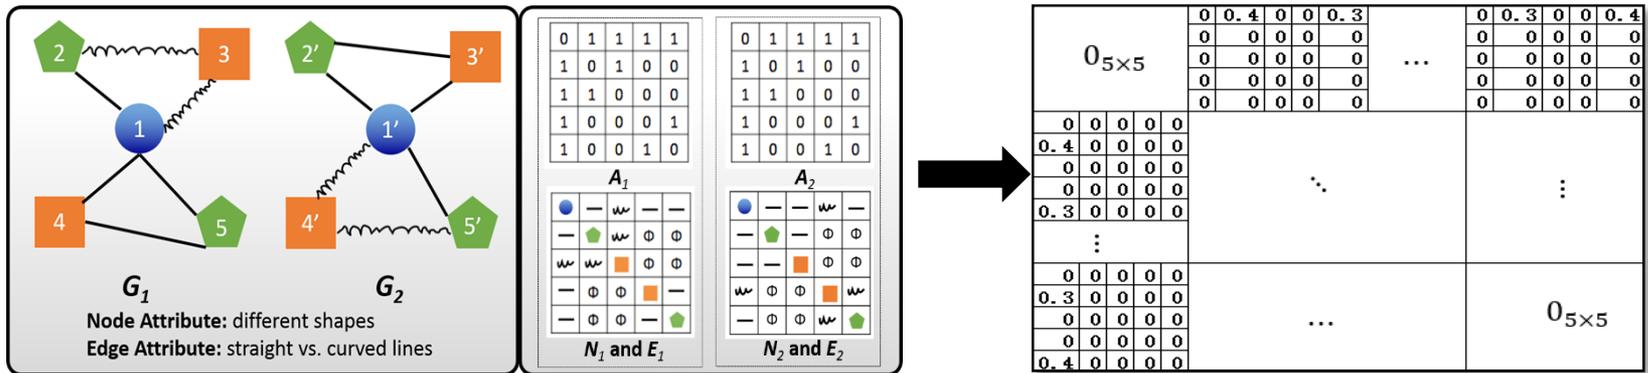
- Matrix-Form Objective Function

$$\min_{\mathbf{S}} J(\mathbf{S}) = \min_{\mathbf{s}} \sum_{v,w} \left[\frac{\mathbf{s}(v)}{\sqrt{\mathbf{D}(v,v)}} - \frac{\mathbf{s}(w)}{\sqrt{\mathbf{D}(w,w)}} \right]^2 \mathbf{W}(v,w)$$

$$\mathbf{s} = \text{vec}(\mathbf{S}) = \min_{\mathbf{s}} \mathbf{s}^T (\mathbf{I} - \widetilde{\mathbf{W}}) \mathbf{s}$$

- $\mathbf{W} = \mathbf{N}[\mathbf{E} \odot (\mathbf{A}_1 \otimes \mathbf{A}_2)]\mathbf{N}$, i.e., the attributed Kronecker product
- \mathbf{D} is the degree matrix of \mathbf{W}

$\widetilde{\mathbf{W}} = \mathbf{D}^{-\frac{1}{2}}\mathbf{W}\mathbf{D}^{-\frac{1}{2}}$ is the symmetrically normalization of \mathbf{W}



FINAL Formulation: Matrix Form with Regularization

- Add a regularization term

$$\min_s \alpha s^T (I - \widetilde{W})s + (1 - \alpha) \|s - \mathbf{h}\|_2^2$$

$$\mathbf{h} = \text{vec}(\mathbf{H})$$

- \mathbf{h} is default as a uniform vector
- \mathbf{h} encodes the prior knowledge of alignment preferences
- \mathbf{h} avoids trivial solution, e.g., optimal solution $s = \mathbf{0}$ w/o \mathbf{h}
- teleport vector in PageRank, restart vector in RWR (on the attributed Kronecker product graph)

Relationship with Existing Methods

- FINAL vs. IsoRank [Singh'08]
 - w/o attribute, FINAL = IsoRank (by a scaling factor $D^{\frac{1}{2}}$)
- FINAL vs. Random Walk Graph Kernel (RWGK) [Vishwanathan'10]
 - $k(G_1, G_2) = \sum_i \mathbf{q}(i)\mathbf{s}(i)$, where \mathbf{q} is the stopping probability vector
- FINAL vs. SimRank (Node Proximity) [Jeh'02]
 - $G_1 = G_2$ and w/o attribute, FINAL = SimRank by a scaling factor $D^{\frac{1}{2}}$
- FINAL vs. Random Walk with Restart (RWR) [Tong'06]
 - \mathbf{s} = RWR vector (defined on the attributed Kronecker graph)

Outline

- Motivations ✓
- Q1: FINAL Formulation ✓
- Q2: FINAL Algorithms
- Q3: FINAL Speed-up Computation
- Experimental Results
- Conclusions

FINAL Solutions

$$s = \alpha D^{-\frac{1}{2}} N \text{vec} \left(\sum_{l=1}^L (E_2^l \odot A_2) Q (E_1^l \odot A_1)^T \right) + (1 - \alpha) h$$

$$\min_s \alpha s^T (I - \widetilde{W}) s + (1 - \alpha) \|s - h\|_2^2$$

- **Obs:** a convex optimization problem
- **Benefits:** a fixed-point solution converging to the global optimal solution

$$s = \alpha \widetilde{W} s + (1 - \alpha) h \Rightarrow s = (1 - \alpha) (I - \alpha \widetilde{W})^{-1} h \text{ (closed form)}$$

- **Intuition:** a similarity propagation to neighboring node-pairs, which is additionally filtered by node/edge attributes
- **Challenges:** computationally VERY expensive
 - Iterative solution: $O(m^2 t_{\max})$ (due to Kronecker product)
 - Closed form solution: $O(m^6)$ (due to matrix inversion)
- **Q:** how to scale up and speed up?

Outline

- Motivations ✓
- Q1: FINAL Formulation ✓
- Q2: FINAL Algorithms ✓
- Q3: FINAL Speed-up Computation
- Experimental Results
- Conclusions

FINAL — Speed-up Full Alignment

- **Obs:** FINAL vs. RWGK and RWR
- **Solution:** leverage the existing fast solutions for RWGK and/or RWR [Kang 2012]

- **An Example:** only consider node attributes

$$s = (1 - \alpha) \left(I - \alpha \mathbf{D}_N^{-\frac{1}{2}} \mathbf{N} (\mathbf{A}_1 \otimes \mathbf{A}_2) \mathbf{N} \mathbf{D}_N^{-\frac{1}{2}} \right)^{-1} \mathbf{h}$$

- **Key Idea:** low rank approximation of \mathbf{A}_1 and \mathbf{A}_2

$$\begin{aligned} \mathbf{A}_1 &\approx \mathbf{U}_1 \mathbf{\Lambda}_1 \mathbf{U}_1^T \\ \mathbf{A}_2 &\approx \mathbf{U}_2 \mathbf{\Lambda}_2 \mathbf{U}_2^T \end{aligned}$$

Sherman-
Morrison Lemma

$$\begin{aligned} s &\approx (1 - \alpha) \left(I + \alpha \mathbf{D}_N^{-\frac{1}{2}} \mathbf{N} \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T \mathbf{N} \mathbf{D}_N^{-\frac{1}{2}} \right) \mathbf{h} \\ \text{where } \mathbf{U} &= \mathbf{U}_1 \otimes \mathbf{U}_2 \\ \mathbf{\Lambda} &= [(\mathbf{\Lambda}_1 \otimes \mathbf{\Lambda}_2)^{-1} - \alpha \mathbf{U}^T \mathbf{N} \mathbf{D}_N^{-1} \mathbf{N} \mathbf{U}]^{-1} \end{aligned}$$

- **Complexity:** $O(n^2 r^4)$
- **Challenge:** it is still $O(n^2)$. Can we do better?

FINAL — Speed-up On-query Alignment

- **Obs:** only need one column, or one segment of S
- **Key Ideas:**
 - Low-rank approximation (same as for the full alignment)
 - Relax the degree matrix $\mathbf{D}_N = \mathbf{D}_1 \otimes \mathbf{D}_2$

- **Details:** $s_a = (1 - \alpha) \left[\mathbf{H}(:, a) + \alpha (\mathbf{D}_1(a, a) \mathbf{D}_2)^{-\frac{1}{2}} \mathbf{N}_a \right]$
 $\times \left[\underbrace{(\mathbf{U}_1(a, :) \otimes \mathbf{U}_2)}_{O(nr^2)} \hat{\Lambda} \underbrace{\mathbf{U}^T \mathbf{N} (\mathbf{D}_1 \otimes \mathbf{D}_2)^{-\frac{1}{2}} \mathbf{h}}_{O(n^2r^2)} \right]$
 (1) same trick as for full alignment

$$\mathbf{g} = \mathbf{U}^T \mathbf{N} (\mathbf{D}_1 \otimes \mathbf{D}_2)^{-\frac{1}{2}} \mathbf{h} = \sum_{i=1}^p \sum_{k=1}^K \sigma_i \underbrace{(\mathbf{U}_1^T \mathbf{N}_1^k \mathbf{D}_1^{-\frac{1}{2}} \mathbf{v}_i)}_{O(nr)} \otimes \underbrace{(\mathbf{U}_2^T \mathbf{N}_2^k \mathbf{D}_2^{-\frac{1}{2}} \mathbf{u}_i)}_{O(nr)}$$

(2) SVD on matrix $\mathbf{H} = \sum_{i=1}^p \sigma_i \mathbf{u}_i \mathbf{v}_i^T$

- **Benefits:** linear complexity $O((Kr^2 + pKr + p^2)n + mr + m_H p + r^6)$

Outline

- Motivations ✓
- Q1: FINAL Formulation ✓
- Q2: FINAL Algorithms ✓
- Q3: FINAL Speed-up Computation ✓
- Experimental Results
- Conclusions

Experimental Setup

■ Datasets:

- DBLP co-author networks (nodes: 9,143 vs. 9,143)
- Douban online & offline networks (nodes: 3,906 vs. 1,118)
- Flickr & Last.fm networks (nodes: 12,974 vs. 15,436)
- Flickr & Myspace networks (nodes: 6,714 vs. 10,733)

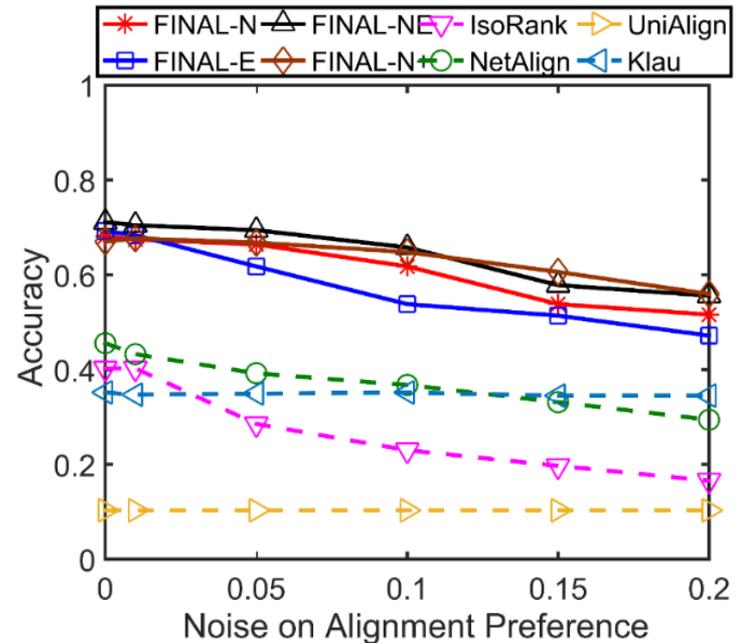
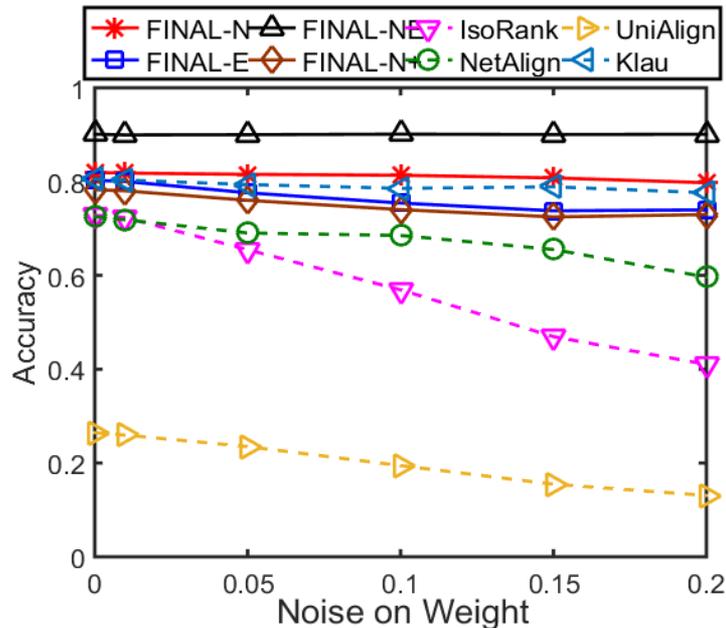
■ Evaluation Objectives:

- Effectiveness: one-to-one alignment accuracy
- Efficiency: running time

■ Comparison Methods:

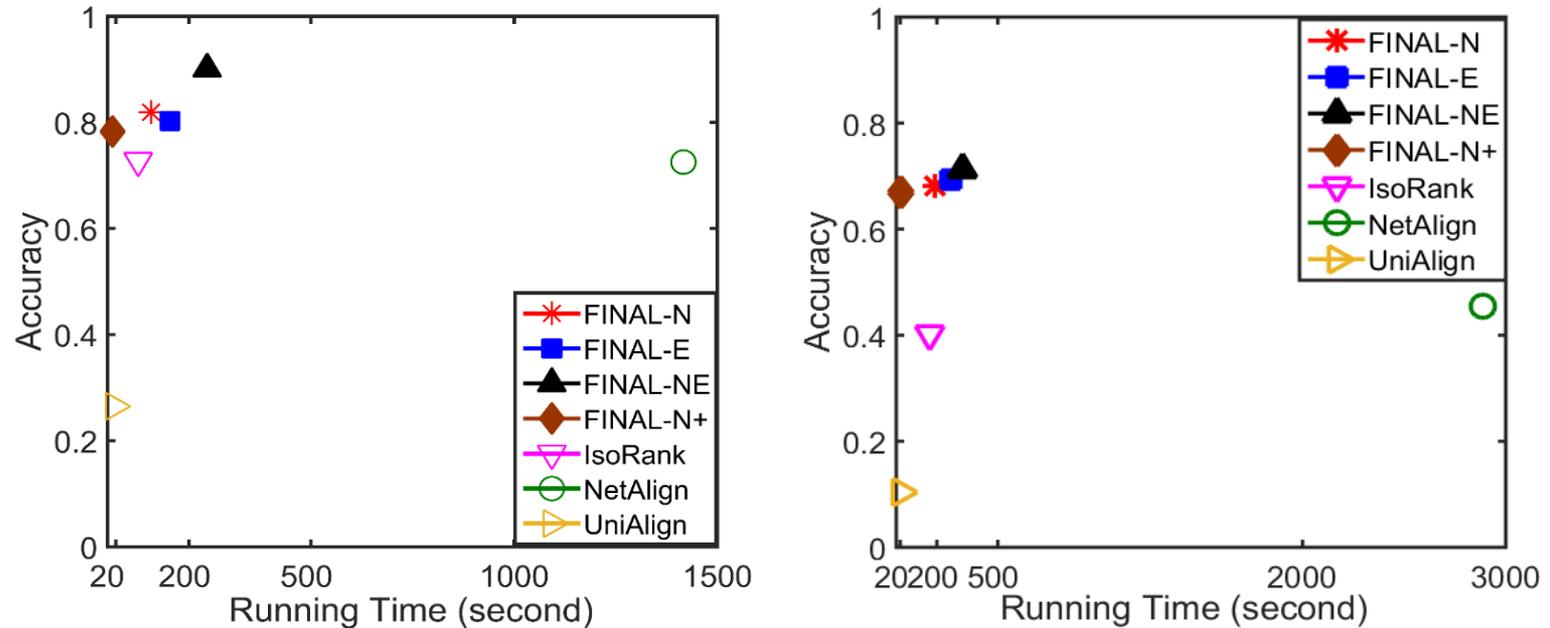
FINAL (Our Methods)	Baseline Methods
<ul style="list-style-type: none">❑ FINAL-NE (with node & edge attributes)❑ FINAL-N (with node attributes)❑ FINAL-E (with edge attributes)❑ FINAL-N+ (speed-up FINAL-N)	<ul style="list-style-type: none">❑ IsoRank [Singh'08]❑ NetAlign [Bayati'09]❑ UniAlign [Koutra'13]❑ Klau's Algorithm [Klau'09]

R1. Effectiveness Results



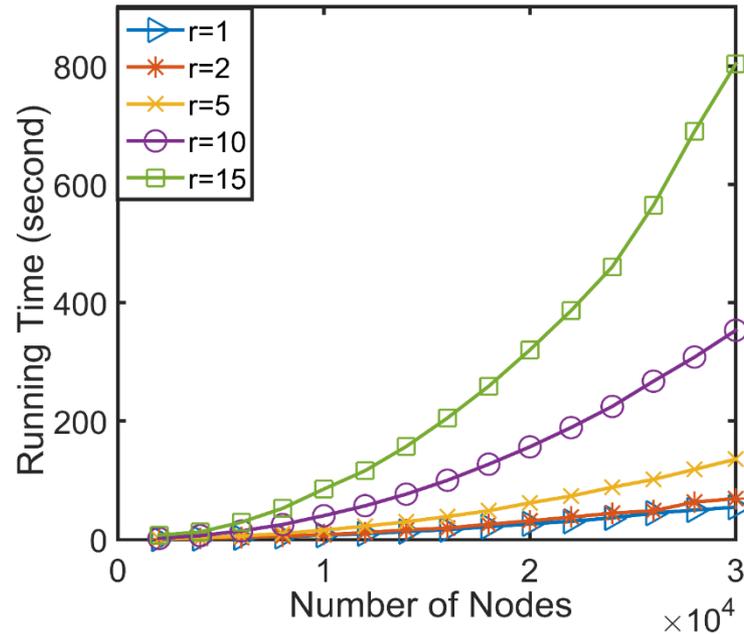
Obs: attributes help improve network alignment

R2. Quality-Speed Balance



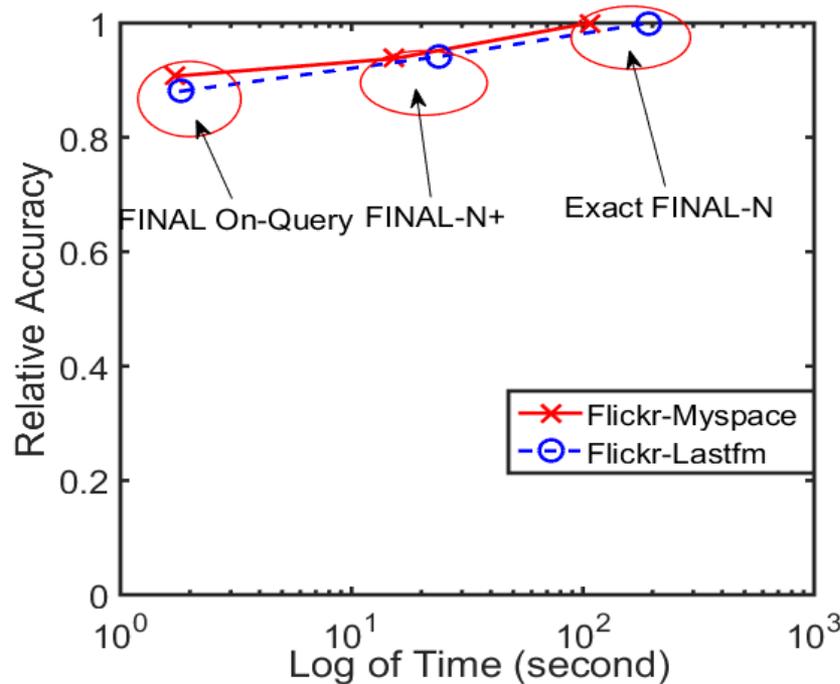
Obs: FINAL gain a better quality-speed balance.

R3. Scalability of FINAL-N+



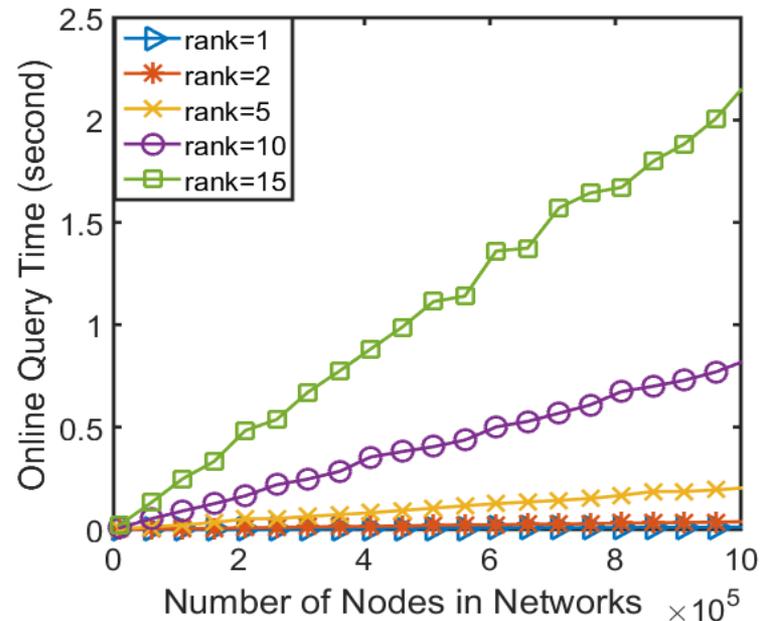
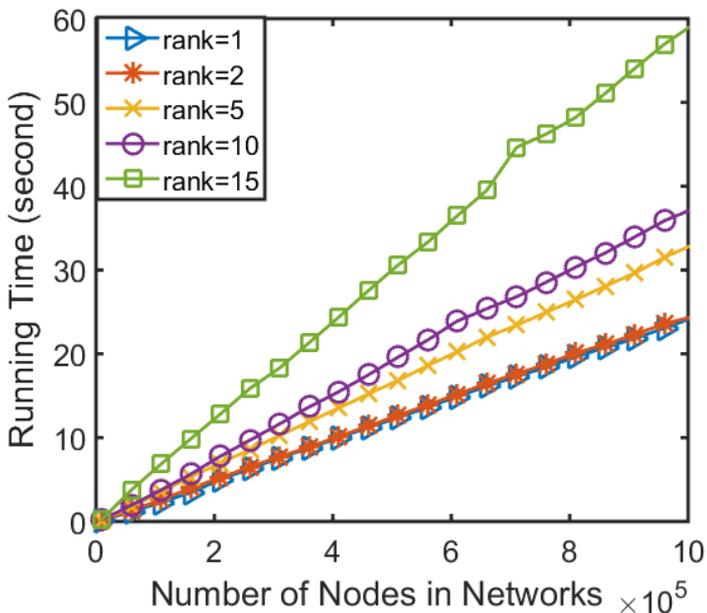
Obs: FINAL-N+ has a quadratic time complexity w.r.t the number of nodes.

R4. Quality-Speed of FINAL On-Query



Obs: FINAL On-Query gains around 90% accuracy relative to exact FINAL-N, but more than 100 times faster.

R5. Scalability of FINAL On-Query



Obs: FINAL On-Query has a **linear** time complexity

Outline

- Motivations ✓
- Q1: FINAL Formulation ✓
- Q2: FINAL Algorithms ✓
- Q3: FINAL Speed-up Computation ✓
- Experimental Results ✓
- **Conclusions**

Conclusions

- Attributed Network Alignment

- **Q1: Formulation**

- **A1: FINAL family**

- **Q2: Optimality**

$$\min_{\mathbf{S}} J(\mathbf{S}) = \sum_{a,b,x,y} \left[\frac{\mathbf{S}(x,a)}{\sqrt{f(x,a)}} - \frac{\mathbf{S}(y,b)}{\sqrt{f(y,b)}} \right]^2 A_1(a,b)A_2(x,y) \\ \times \mathbb{I}(N_1(a,a) = N_2(x,x))\mathbb{I}(N_1(b,b) = N_2(y,y)) \\ \times \mathbb{I}(E_1(a,b) = E_2(x,y))$$

- **A2: Convex optimization problem** → global optimal solution

- **Q3: Scalable computation**

- **A3: Fast algorithms (FINAL-N+ & FINAL On-Query)**

- Results

- FINAL outperform other baseline methods

- FINAL On-Query linear complexity

- More in Paper

- Proof of optimality & more experimental results

