Triangle switches: irreducibility and mixing

Catherine Greenhill

School of Mathematics and Statistics
UNSW Sydney

Joint work with Colin Cooper (Kings College London)
Martin Dyer (Leeds)

Random graphs, degree sequences and other properties

$G_{n,p}$ is the binomial random graph, with n vertices and each possible edge present with probability p, independently of all other edge/non-edge choices.

Random graphs, degree sequences and other properties

$G_{n,p}$ is the binomial random graph, with n vertices and each possible edge present with probability p, independently of all other edge/non-edge choices.

The degree of a vertex in $G_{n,p}$ has the binomial distribution $\mathrm{Bin}(n-1,p)$.

Random graphs, degree sequences and other properties

$G_{n,p}$ is the binomial random graph, with n vertices and each possible edge present with probability p, independently of all other edge/non-edge choices.

The degree of a vertex in $G_{n,p}$ has the binomial distribution $\text{Bin}(n-1, p)$.

If $p = p(n)$ is not too small then with high probability all degrees in $G_{n,p}$ are concentrated around the mean.

Random graphs, degree sequences and other properties

$G_{n,p}$ is the binomial random graph, with n vertices and each possible edge present with probability p, independently of all other edge/non-edge choices.

The degree of a vertex in $G_{n,p}$ has the binomial distribution $\mathrm{Bin}(n-1, p)$.

If $p = p(n)$ is not too small then with high probability all degrees in $G_{n,p}$ are concentrated around the mean.

Specifically, if $\varepsilon > 0$ and $p = \Omega\left(\frac{\log n}{n\varepsilon^2}\right)$ then with high probability all degrees lie in $[(1-\varepsilon)pn,\ (1+\varepsilon)pn]$.

Problem: Real-world networks often have heavy-tailed degree distributions.

Problem: Real-world networks often have heavy-tailed degree distributions.

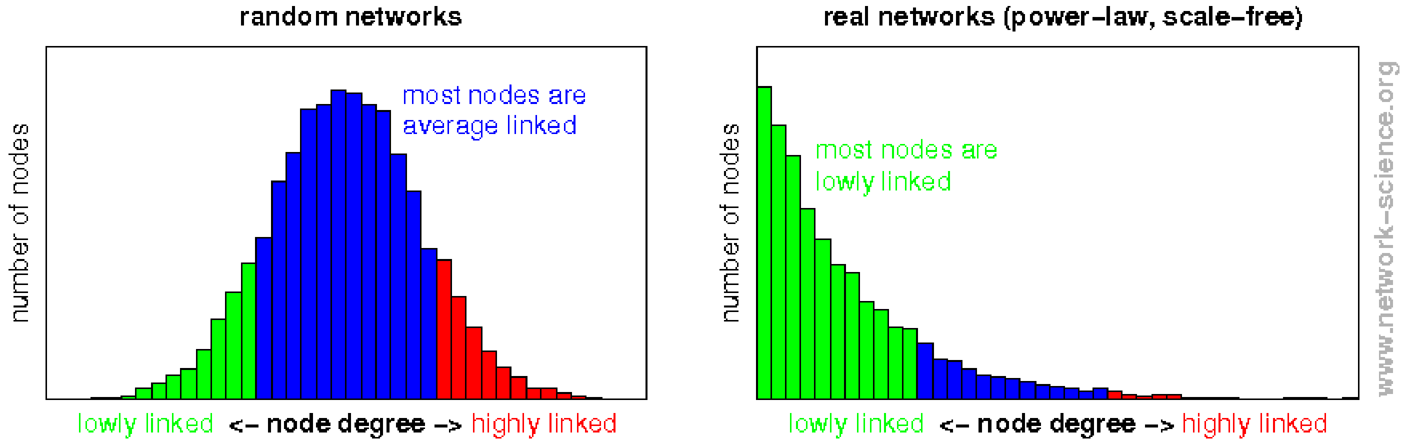

(Image from `network-science.org`)

Problem: Real-world networks often have heavy-tailed degree distributions.

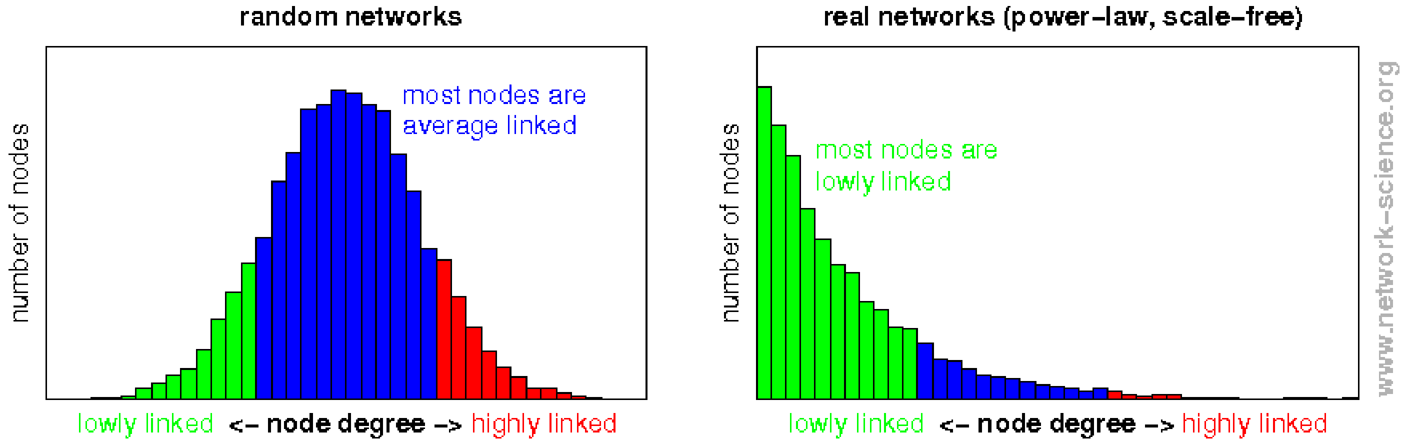

(Image from `network-science.org`)

So if you want to use random graphs to model some real-world network then $G_{n,p}$ might not be a good choice.

Instead, you might want more control over the degree sequence of the random graph.

We will look at the problem of generating random graphs with a given degree sequence $\boldsymbol{d}$.

Instead, you might want more control over the degree sequence of the random graph.

We will look at the problem of generating random graphs with a given degree sequence $\boldsymbol{d}$.

- This could then be applied to a random degree sequence, e.g. with each entry i.i.d. from some distribution.

Instead, you might want more control over the degree sequence of the random graph.

We will look at the problem of generating random graphs with a given degree sequence $\boldsymbol{d}$.

- This could then be applied to a random degree sequence, e.g. with each entry i.i.d. from some distribution.

- Also, Chung & Lu (2002) gave an efficient algorithm for generating random graphs with a given expected degree sequence.

See the excellent book by van der Hofstad (2016).

Graphs with given degrees

Graphs with given degrees

Let G be a graph with vertex set $[n] = \{1, 2, \dots, n\}$.
The degree sequence of G is $(d_1, \dots, d_n)$, where d_j is the degree of vertex j.

Graphs with given degrees

Let G be a graph with vertex set $[n] = \{1, 2, \ldots, n\}$.
The degree sequence of G is $(d_1, \ldots, d_n)$, where d_j is the degree of vertex j.

A sequence $\boldsymbol{d} = (d_1, \ldots, d_n)$ is graphical if there exists a graph with degree sequence $\boldsymbol{d}$.

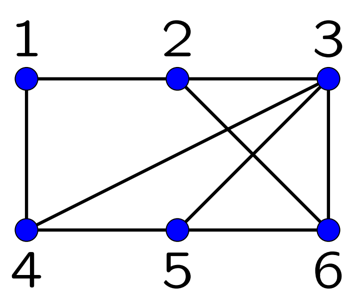

$\Rightarrow$ $(2, 3, 4, 3, 3, 3)$ is graphical

Let $G(\boldsymbol{d})$ be the set of all graphs on the vertex set $[n]$ with degree sequence $\boldsymbol{d}$. If $\boldsymbol{d} = (d, d, \ldots, d)$ is regular then write $G(n, d)$ instead.

Let $G(\boldsymbol{d})$ be the set of all graphs on the vertex set $[n]$ with degree sequence $\boldsymbol{d}$. If $\boldsymbol{d} = (d, d, \ldots, d)$ is regular then write $G(n, d)$ instead.

We really have a sequence of sets $G(\boldsymbol{d}(n))$, and we are interested in asymptotics as $n \to \infty$.

Lots of prior work on algorithms for sampling from $G(\boldsymbol{d})$.

See my BCC 2021 survey for some details (!!).

Let $G(\boldsymbol{d})$ be the set of all graphs on the vertex set $[n]$ with degree sequence $\boldsymbol{d}$. If $\boldsymbol{d} = (d, d, \ldots, d)$ is regular then write $G(n, d)$ instead.

We really have a sequence of sets $G(\boldsymbol{d}(n))$, and we are interested in asymptotics as $n \to \infty$.

Lots of prior work on algorithms for sampling from $G(\boldsymbol{d})$.

See my BCC 2021 survey for some details (!!).

We will discuss the Markov chain approach.

But if your maximum degree $d_{\max}$ is not too large then you should use the very fast exactly uniform sampling algorithms of Arman, Gao & Wormald (2021).

A switch is the smallest possible change you can make to the edge set of a graph without affecting the degree sequence.

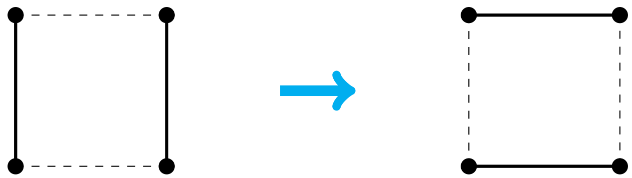

A switch is the smallest possible change you can make to the edge set of a graph without affecting the degree sequence.

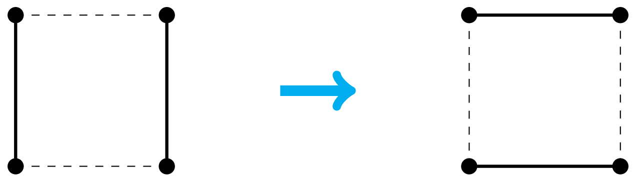

The switch chain is a particular Markov chain on $G(\boldsymbol{d})$ which performs a random switch at each time step.

The switch Markov chain

From current graph $G \in G(\boldsymbol{d})$:

- choose two non-adjacent edges uv, yz u.a.r.
- choose a perfect matching M of u, v, y, z u.a.r.
- if $(E(G) \setminus \{uv, yz\}) \cap M = \emptyset$ then
  delete edges uv, yz and insert edges M
- otherwise, stay at G.

Here u.a.r. means uniformly at random.

- Switches connect the set $G(\boldsymbol{d})$.

• Switches connect the set $G(\boldsymbol{d})$.

Some attribute this to Taylor (1981). He may have been the first to write down a proof for (simple) graphs with arbitrary $\boldsymbol{d}$.

- Switches connect the set $G(\boldsymbol{d})$.

Some attribute this to Taylor (1981). He may have been the first to write down a proof for (simple) graphs with arbitrary $\boldsymbol{d}$.

But Petersen (1891) proved that switches connect $G(n, d)$ when d is even (i.e., regular graphs of even degree).

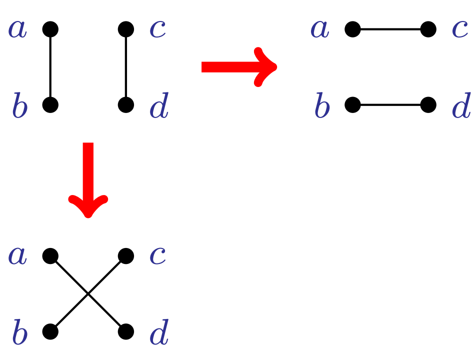

7. Aus einem *graph* geraden Grades $G_{2\alpha}$ können wir einen neuen $G'_{2\alpha}$ bilden, indem wir zwei nicht zusammenstossende Linien *ab* und *cd* entfernen und für diese zwei neue Linien *ac* und *bd* oder *ad* und *bc* hineinsetzen. Finden sich mehrere Linien *ab*, so wird nur die eine entfernt. Ob eine zugesetzte Linie sich schon im *graph* findet, ist ohne Bedeutung; sie bekommt dann eine um eins erhöhte Multiplicität. Ich werde die zwei *graphs gepaart* nennen.

- Switches connect the set $G(\boldsymbol{d})$.
- The switch chain is also aperiodic as $P(G, G) \geq 1/3$ for all $G \in G(\boldsymbol{d})$.

- Switches connect the set $G(\boldsymbol{d})$.
- The switch chain is also aperiodic as $P(G, G) \geq 1/3$ for all $G \in G(\boldsymbol{d})$.
- Hence the switch chain is ergodic, so it has a unique stationary distribution which is a limiting distribution.

- Switches connect the set $G(\boldsymbol{d})$.

- The switch chain is also aperiodic as $P(G,G) \geq 1/3$ for all $G \in G(\boldsymbol{d})$.

- Hence the switch chain is ergodic, so it has a unique stationary distribution which is a limiting distribution.

- The stationary distribution π is uniform on $G(\boldsymbol{d})$. This follows from the detailed balance equations:

$$\pi(x)\, P(x,y) = \pi(y)\, P(y,x) \qquad \text{for all } x,y \in \Omega.$$

The switch chain is known to be rapidly mixing for a wide range of degree sequences, including those from

The switch chain is known to be rapidly mixing for a wide range of degree sequences, including those from

- ◇ P-stable families (Erdős et al., 2022), and
- ◇ certain heavy-tailed distributions (Gao & Greenhill, 2021).

("Rapidly" means in time polynomial in n and d_{max}.)

The switch chain is known to be rapidly mixing for a wide range of degree sequences, including those from

- ◇ P-stable families (Erdős et al., 2022), and
- ◇ certain heavy-tailed distributions (Gao & Greenhill, 2021).

(“Rapidly” means in time polynomial in n and d_{max}.)

Almost all proofs rest on a multicommodity flow argument, which is a generalisation of a canonical path argument. The resulting runtime bounds are very high degree polynomials and are believed to be very far from tight.

An exception...

Tikhomirov & Youssef, `arXiv.2206.12477` proved that the switch chain converges in time $C_d\, n \log n$ on d-regular bipartite graphs, where $2 \leq d \leq n/2$, for some constant $C_d > 0$ which depends only on d.

An exception...

Tikhomirov & Youssef, `arXiv.2206.12477` proved that the switch chain converges in time $C_d\, n \log n$ on d-regular bipartite graphs, where $2 \leq d \leq n/2$, for some constant $C_d > 0$ which depends only on d.

This bound is optimal up to the constant C_d, improves on the same authors' earlier work (*Probab. Theory Related Fields* 2022) and is a **huge** improvement on all earlier bounds.

An exception...

Tikhomirov & Youssef, `arXiv.2206.12477` proved that the switch chain converges in time $C_d\, n \log n$ on d-regular bipartite graphs, where $2 \leq d \leq n/2$, for some constant $C_d > 0$ which depends only on d.

This bound is optimal up to the constant C_d, improves on the same authors' earlier work (*Probab. Theory Related Fields* 2022) and is a **huge** improvement on all earlier bounds.

Proof involves establishing a new comparison result for the modified log-Sobolev inequality.

Quick reminder: Canonical paths

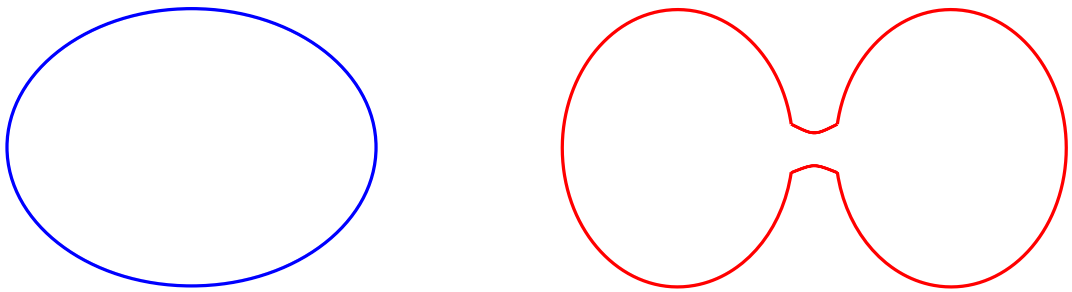

Constrictions in the state space make it difficult for the chain to escape: exponential time required to converge to stationary distribution.

Quick reminder: Canonical paths

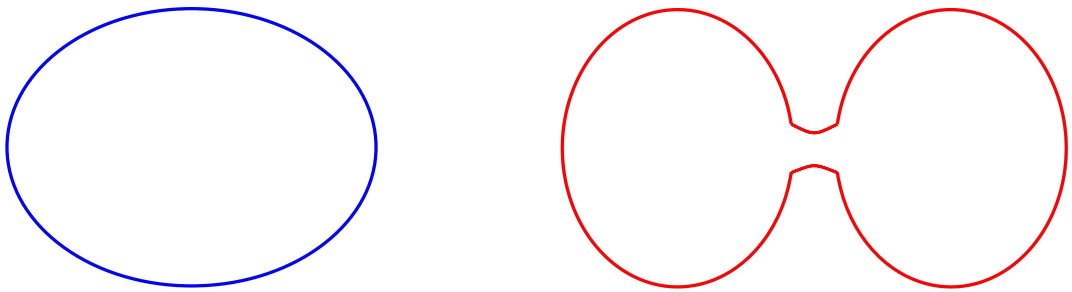

Constrictions in the state space make it difficult for the chain to escape: exponential time required to converge to stationary distribution.

Lack of constrictions allows chain to converge quickly. Results by Jerrum & Sinclair (1987) make this precise.

Quick reminder: Canonical paths

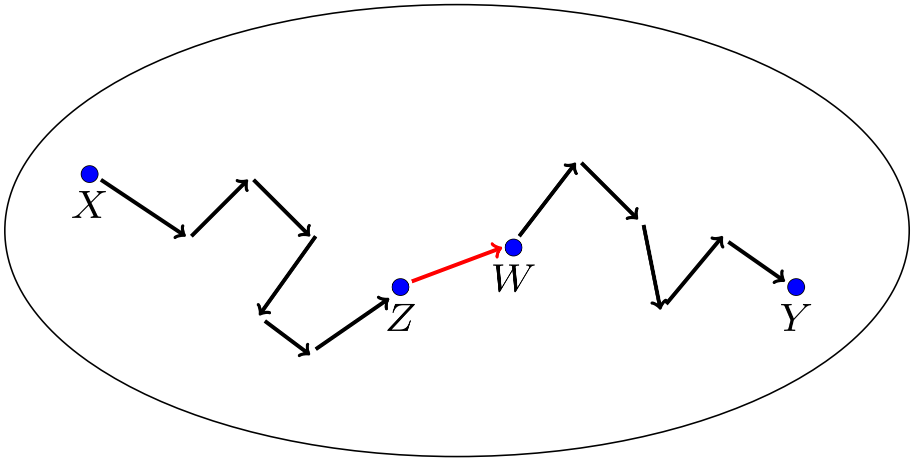

- For all pairs $(X, Y) \in \Omega^2$, define a path from X to Y, where each step is a transition of the Markov chain.

Quick reminder: Canonical paths

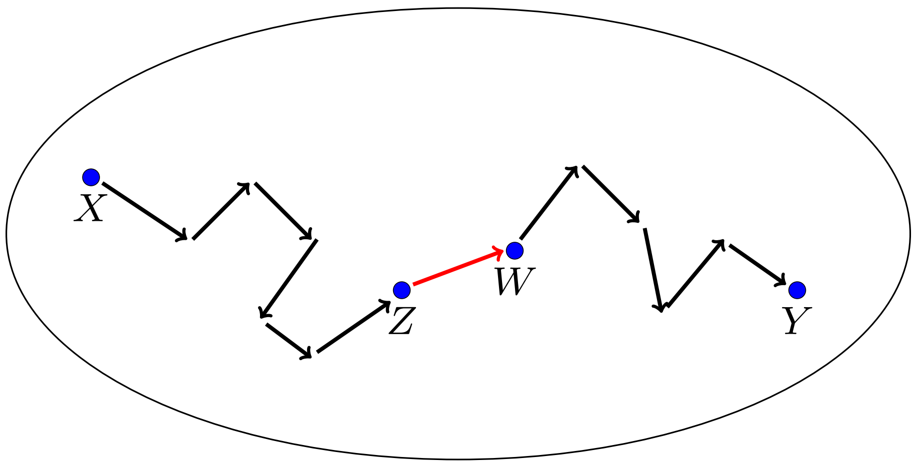

- For all pairs $(X, Y) \in \Omega^2$, define a path from X to Y, where each step is a transition of the Markov chain.

- Analyse the congestion of the set of all paths: are any transitions heavily loaded? Then apply Sinclair (1992).

Problem:
What if you also care about other features of the random graph, not just the degree sequence?

Problem:
What if you also care about other features of the random graph, not just the degree sequence?

For example, in social networks we see a lot of triangles.

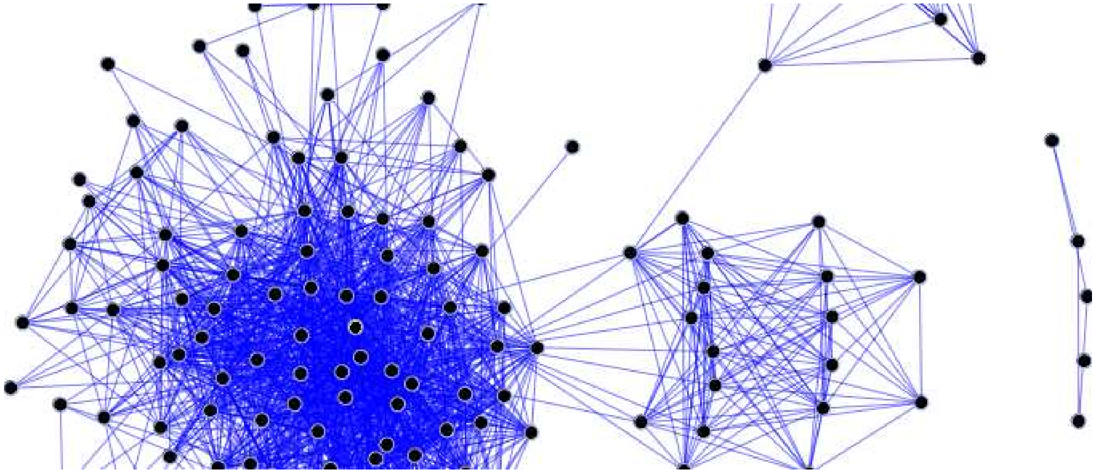

Image from https://commons.wikimedia.org/wiki/File:Sna_large.png

Problem:
What if you also care about other features of the random graph, not just the degree sequence?

For example, in social networks we see a lot of triangles.

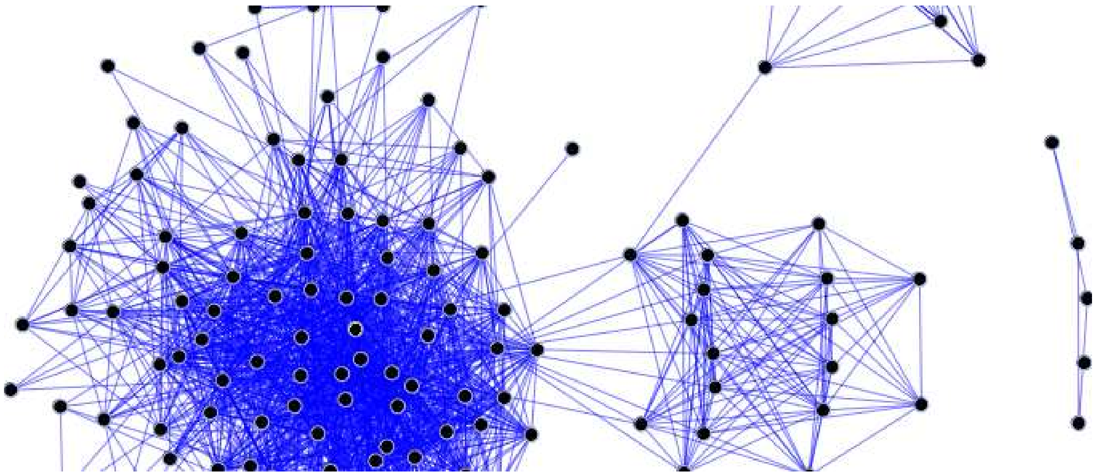

Image from https://commons.wikimedia.org/wiki/File:Sna_large.png

This phenomenon ("triadic closure") dates back to 1908: *Soziologie* by Georg Simmel.

But in a graph chosen randomly from $G(\boldsymbol{d})$, the expected number of triangles is asymptotically equal to

$$\mu(\boldsymbol{d}) := \frac{M_2^3}{6M^3}$$

where

$$M = M(\boldsymbol{d}) = \sum_{j \in [n]} d_j, \qquad M_2 = M_2(\boldsymbol{d}) = \sum_{j \in [n]} d_j(d_j - 1).$$

But in a graph chosen randomly from $G(\boldsymbol{d})$, the expected number of triangles is asymptotically equal to

$$\mu(\boldsymbol{d}) := \frac{M_2^3}{6M^3}$$

where

$$M = M(\boldsymbol{d}) = \sum_{j \in [n]} d_j, \qquad M_2 = M_2(\boldsymbol{d}) = \sum_{j \in [n]} d_j(d_j - 1).$$

Since $M_2 \leq d_{\max} M$, in particular this means that if $d_{\max}$ is constant then the expected number of triangles is at most $d_{\max}^3/6$, also constant.
(Distribution is asymptotically Poisson.)

But in a graph chosen randomly from $G(\boldsymbol{d})$, the expected number of triangles is asymptotically equal to

$$\mu(\boldsymbol{d}) := \frac{M_2^3}{6M^3}$$

where

$$M = M(\boldsymbol{d}) = \sum_{j \in [n]} d_j, \qquad M_2 = M_2(\boldsymbol{d}) = \sum_{j \in [n]} d_j(d_j - 1).$$

Since $M_2 \le d_{\max} M$, in particular this means that if $d_{\max}$ is constant then the expected number of triangles is at most $d_{\max}^3/6$, also constant.
(Distribution is asymptotically Poisson.)

So a random graph from $G(\boldsymbol{d})$ might not be a great model for a social network.

Idea: Can we adapt the switch chain to enable more control on the number of triangles?

Idea: Can we adapt the switch chain to enable more control on the number of triangles?

Definition: A switch is a $\triangle$-switch (triangle switch) if it changes the set of triangles.

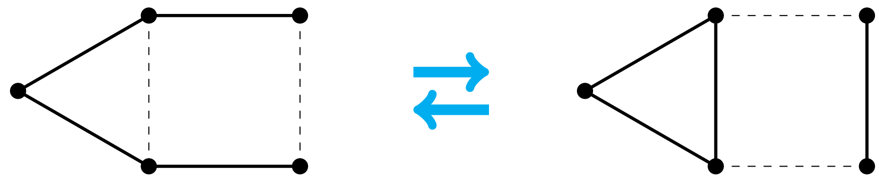

Idea: Can we adapt the switch chain to enable more control on the number of triangles?

Definition: A switch is a $\triangle$-switch (triangle switch) if it changes the set of triangles.

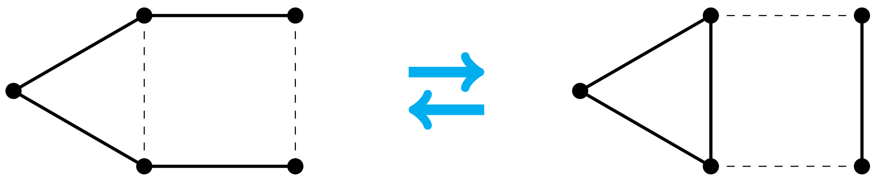

Q: Do $\triangle$-switches connect $G(\boldsymbol{d})$?

Note, one △-switch can affect a lot of triangles!

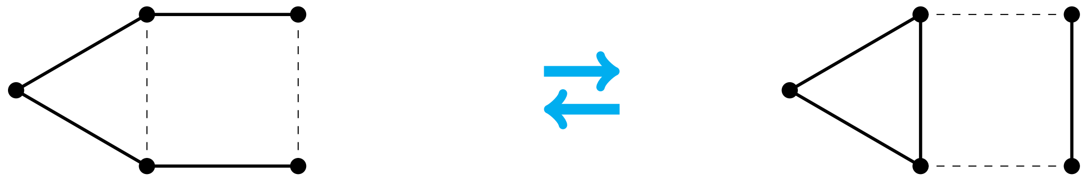

Note, one △-switch can affect a lot of triangles!

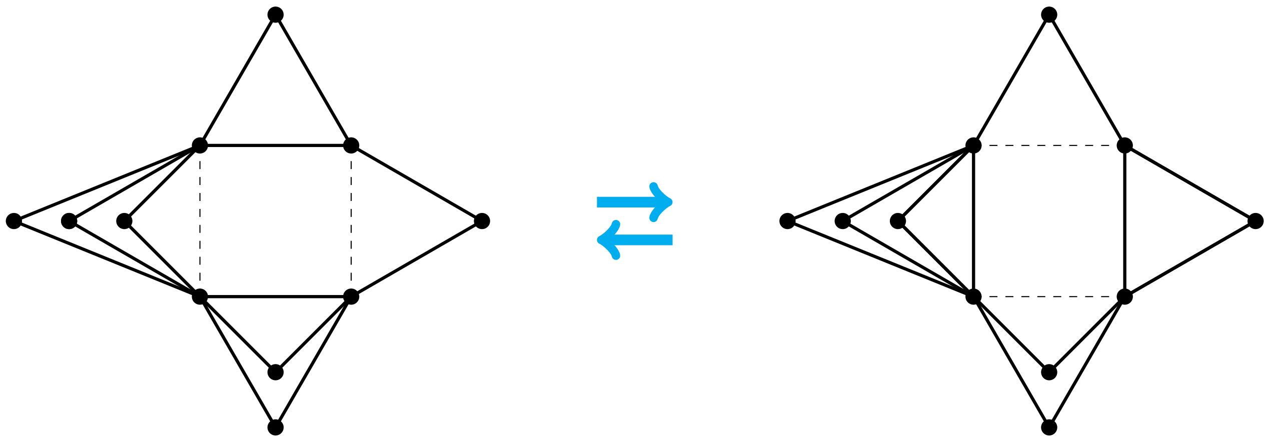

Sometimes this can complicate the analysis.

Cooper, Dyer, Greenhill (IWOCA 2021):
Triangle switches connect $G(n, d)$ whenever $d \geq 3$.

Cooper, Dyer, Greenhill (IWOCA 2021):
Triangle switches connect $G(n,d)$ whenever $d \geq 3$.

Proof: Given any $G \in G(n,d)$ we found a sequence of triangle switches which transformed G into a union of many disjoint copies of K_{d+1} and at most one "fragment" F with $d+1 < |F| < 2(d+1)$.

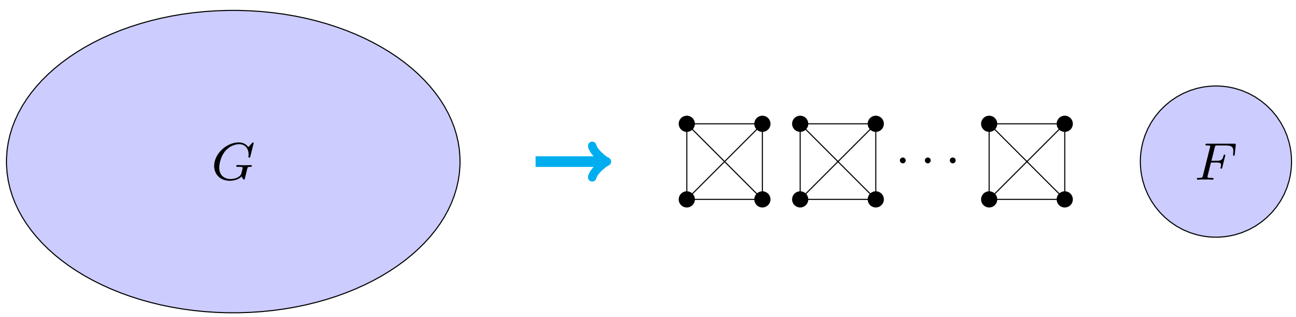

Cooper, Dyer, Greenhill (IWOCA 2021):
Triangle switches connect $G(n,d)$ whenever $d \geq 3$.

Proof: Given any $G \in G(n,d)$ we found a sequence of triangle switches which transformed G into a union of many disjoint copies of K_{d+1} and at most one "fragment" F with $d+1 < |F| < 2(d+1)$.

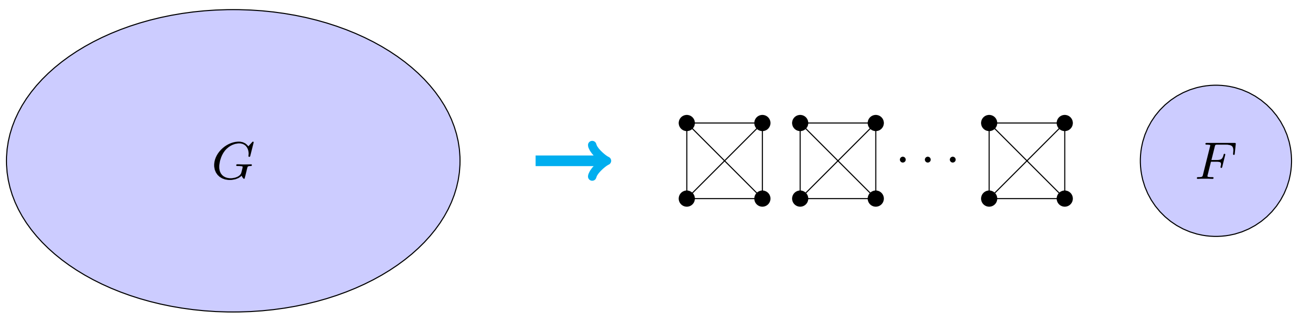

Then we showed that we could transform any two graphs of this form into each other, using triangle switches. □

But, if we wanted to try to prove rapid mixing using canonical paths...

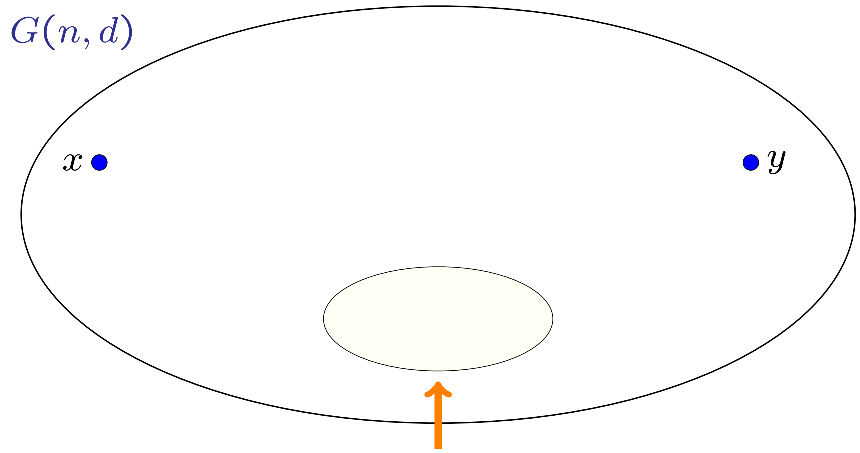

graphs consisting of lots of K_{d+1}'s and one fragment

But, if we wanted to try to prove rapid mixing using canonical paths...

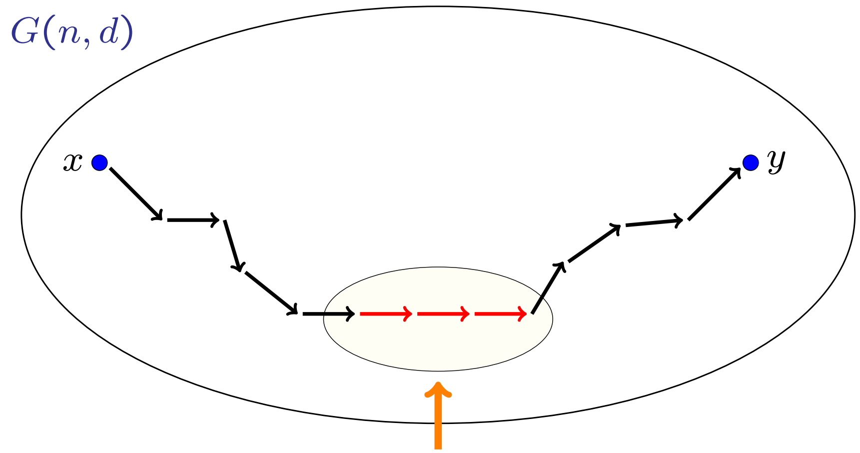

graphs consisting of lots of K_{d+1}'s and one fragment

But, if we wanted to try to prove rapid mixing using canonical paths...

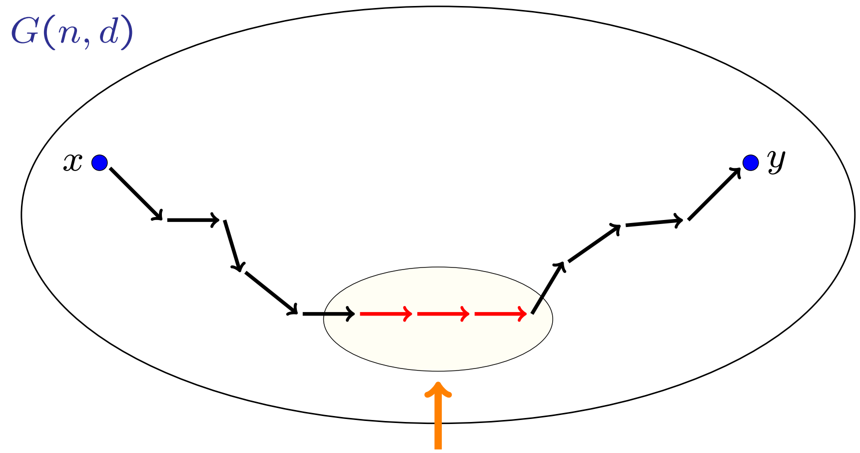

graphs consisting of lots of K_{d+1}'s and one fragment

... these paths would be a pretty bad choice.

New approach (Cooper, Dyer & me):

Let $\boldsymbol{d}$ be a graphical degree sequence with $d_{\text{min}} \geq 3$.

We proved that we can simulate any switch using a "simulation path" of at most five $\triangle$-switches.

New approach (Cooper, Dyer & me):

Let $\boldsymbol{d}$ be a graphical degree sequence with $d_{\mathsf{min}} \geq 3$.

We proved that we can simulate any switch using a "simulation path" of at most five $\triangle$-switches.

Proof: combinatorial argument, lots of cases.

New approach (Cooper, Dyer & me):

Let $\boldsymbol{d}$ be a graphical degree sequence with $d_{\text{min}} \geq 3$.

We proved that we can simulate any switch using a "simulation path" of at most five $\triangle$-switches.

Proof: combinatorial argument, lots of cases.

Since switches connect the space $G(\boldsymbol{d})$ (by Taylor 1981), it follows that $\triangle$-switches connect $G(\boldsymbol{d})$.

New approach (Cooper, Dyer & me):

Let $\boldsymbol{d}$ be a graphical degree sequence with $d_{\text{min}} \geq 3$.

We proved that we can simulate any switch using a "simulation path" of at most five $\triangle$-switches.

Proof: combinatorial argument, lots of cases.

Since switches connect the space $G(\boldsymbol{d})$ (by Taylor 1981), it follows that $\triangle$-switches connect $G(\boldsymbol{d})$.

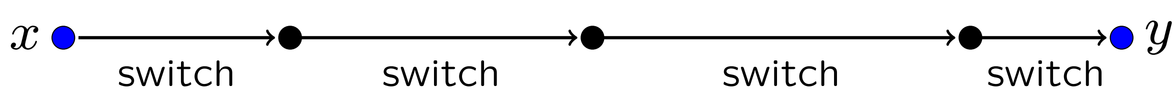

New approach (Cooper, Dyer & me):

Let $\boldsymbol{d}$ be a graphical degree sequence with $d_{\text{min}} \geq 3$.

We proved that we can simulate any switch using a "simulation path" of at most five $\triangle$-switches.

Proof: combinatorial argument, lots of cases.

Since switches connect the space $G(\boldsymbol{d})$ (by Taylor 1981), it follows that $\triangle$-switches connect $G(\boldsymbol{d})$.

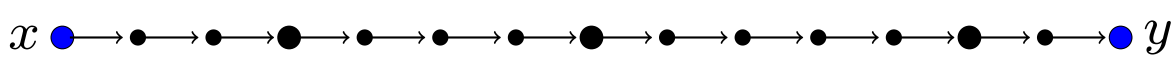

each switch is simulated by a sequence of $\triangle$-switches

How?

If the switch is also a $\triangle$-switch then we are done.

How?

If the switch is also a $\triangle$-switch then we are done.

Otherwise, we work through several cases depending on how many "diagonals" are present, on whether the endvertices of the diagonals have common neighbours, etc.

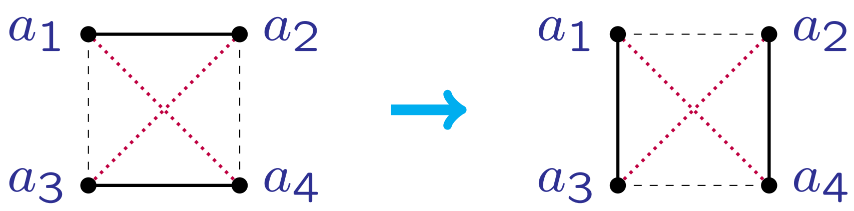

How?

If the switch is also a $\triangle$-switch then we are done.

Otherwise, we work through several cases depending on how many "diagonals" are present, on whether the endvertices of the diagonals have common neighbours, etc.

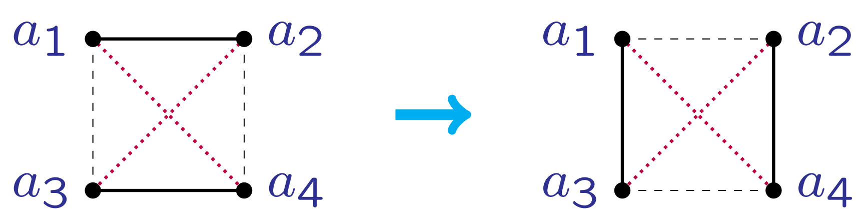

For example, suppose one diagonal a_1a_4 is present, that a_2 and a_3 have two common neighbours u, v which are adjacent, and a_1u, a_4v are non-edges. (This is Case V.)

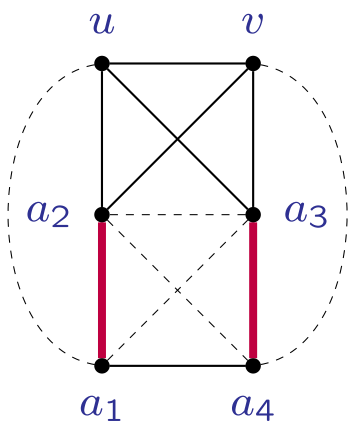
u
v
a_2
a_3
a_1
a_4

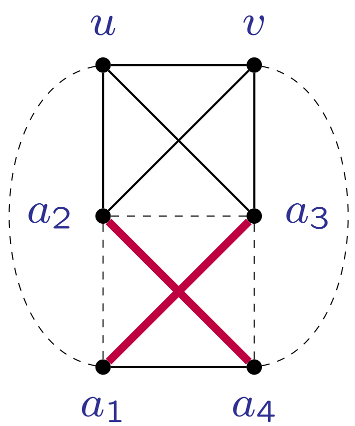
u
v
a_2
a_3
a_1
a_4

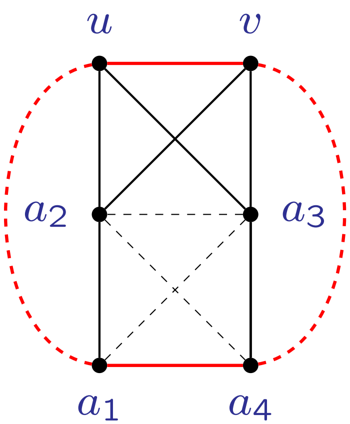

u
v
a_2
a_3
a_1
a_4

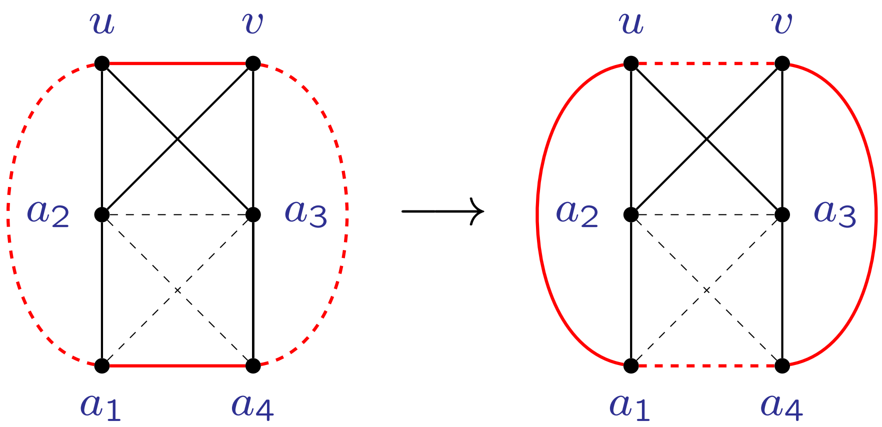

u
v
a2
a3
a1
a4
u
v
a2
a3
a1
a4

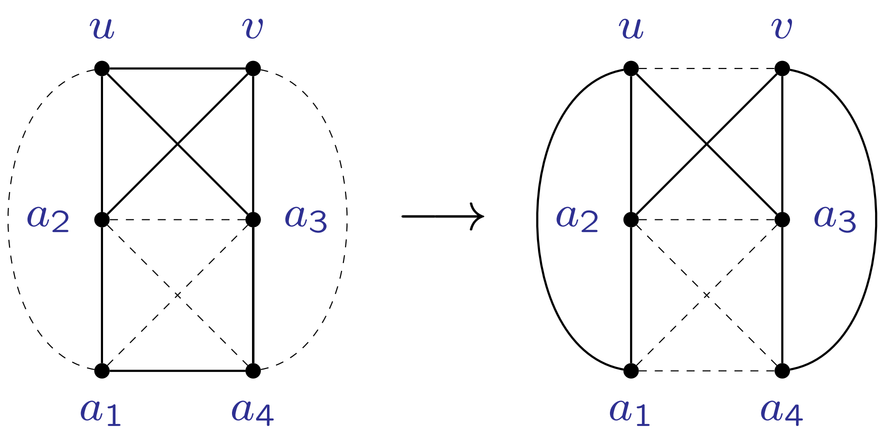
u
v
a_2
a_3
a_1
a_4
u
v
a_2
a_3
a_1
a_4

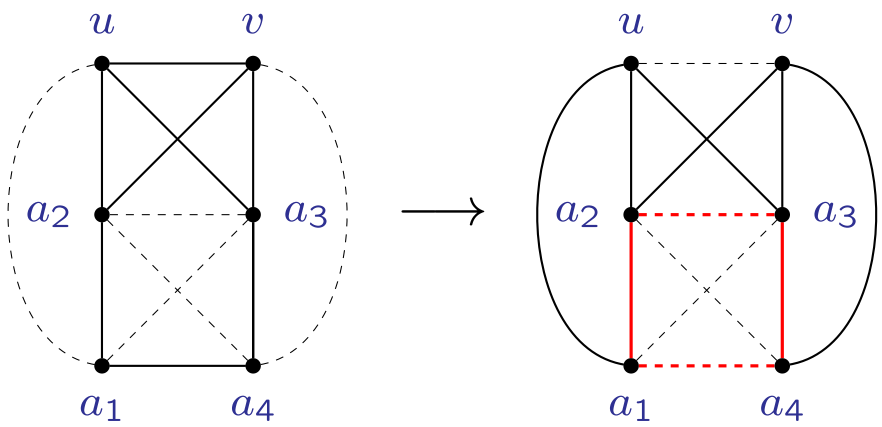

u
v
a2
a3
a1
a4
u
v
a2
a3
a1
a4

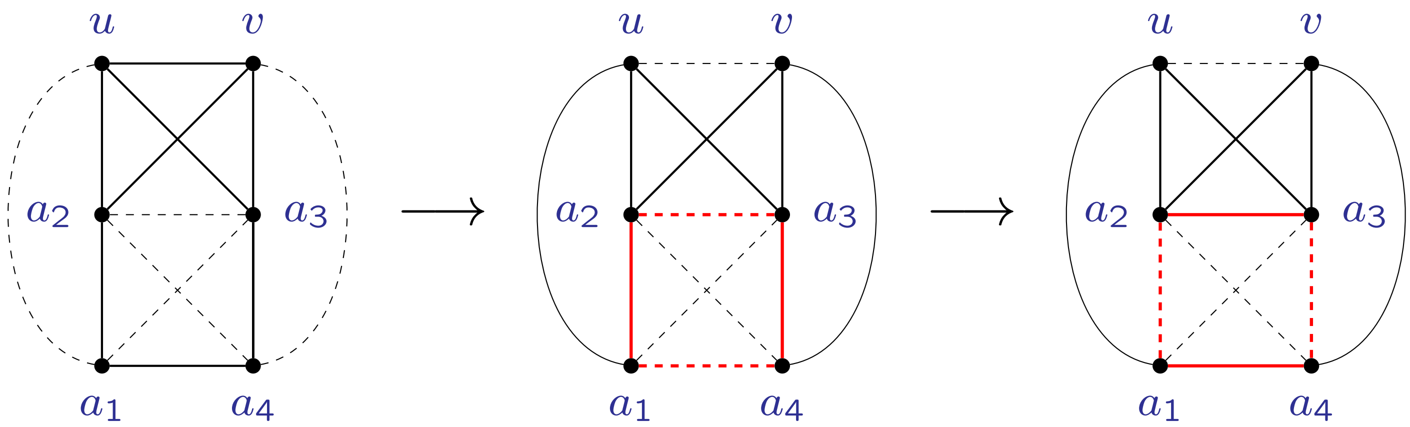

u
v
a_2
a_3
a_1
a_4
u
v
a_2
a_3
a_1
a_4
u
v
a_2
a_3
a_1
a_4

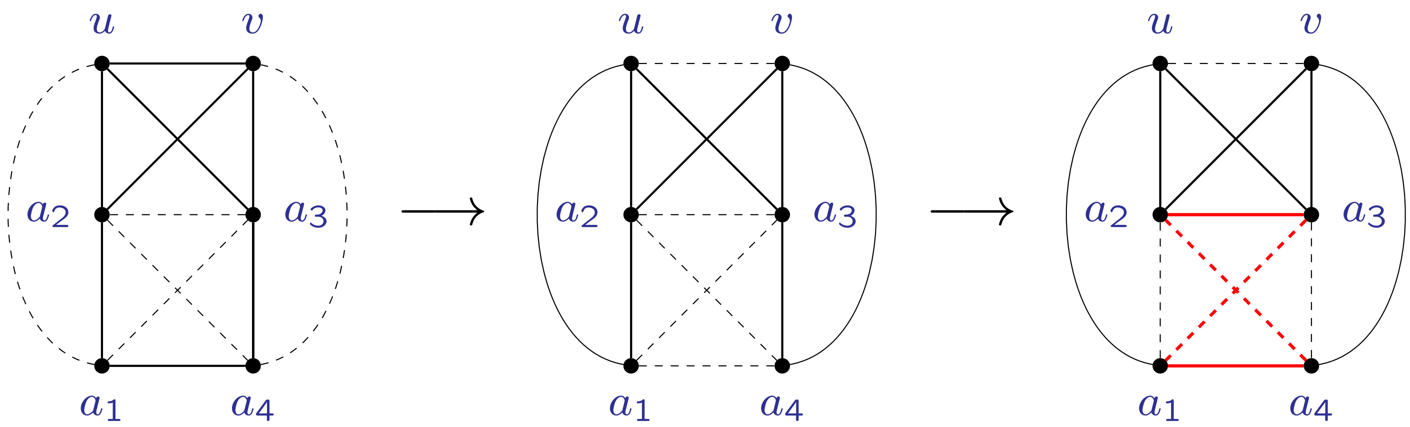

u
v
a2
a3
a1
a4
u
v
a2
a3
a1
a4
u
v
a2
a3
a1
a4

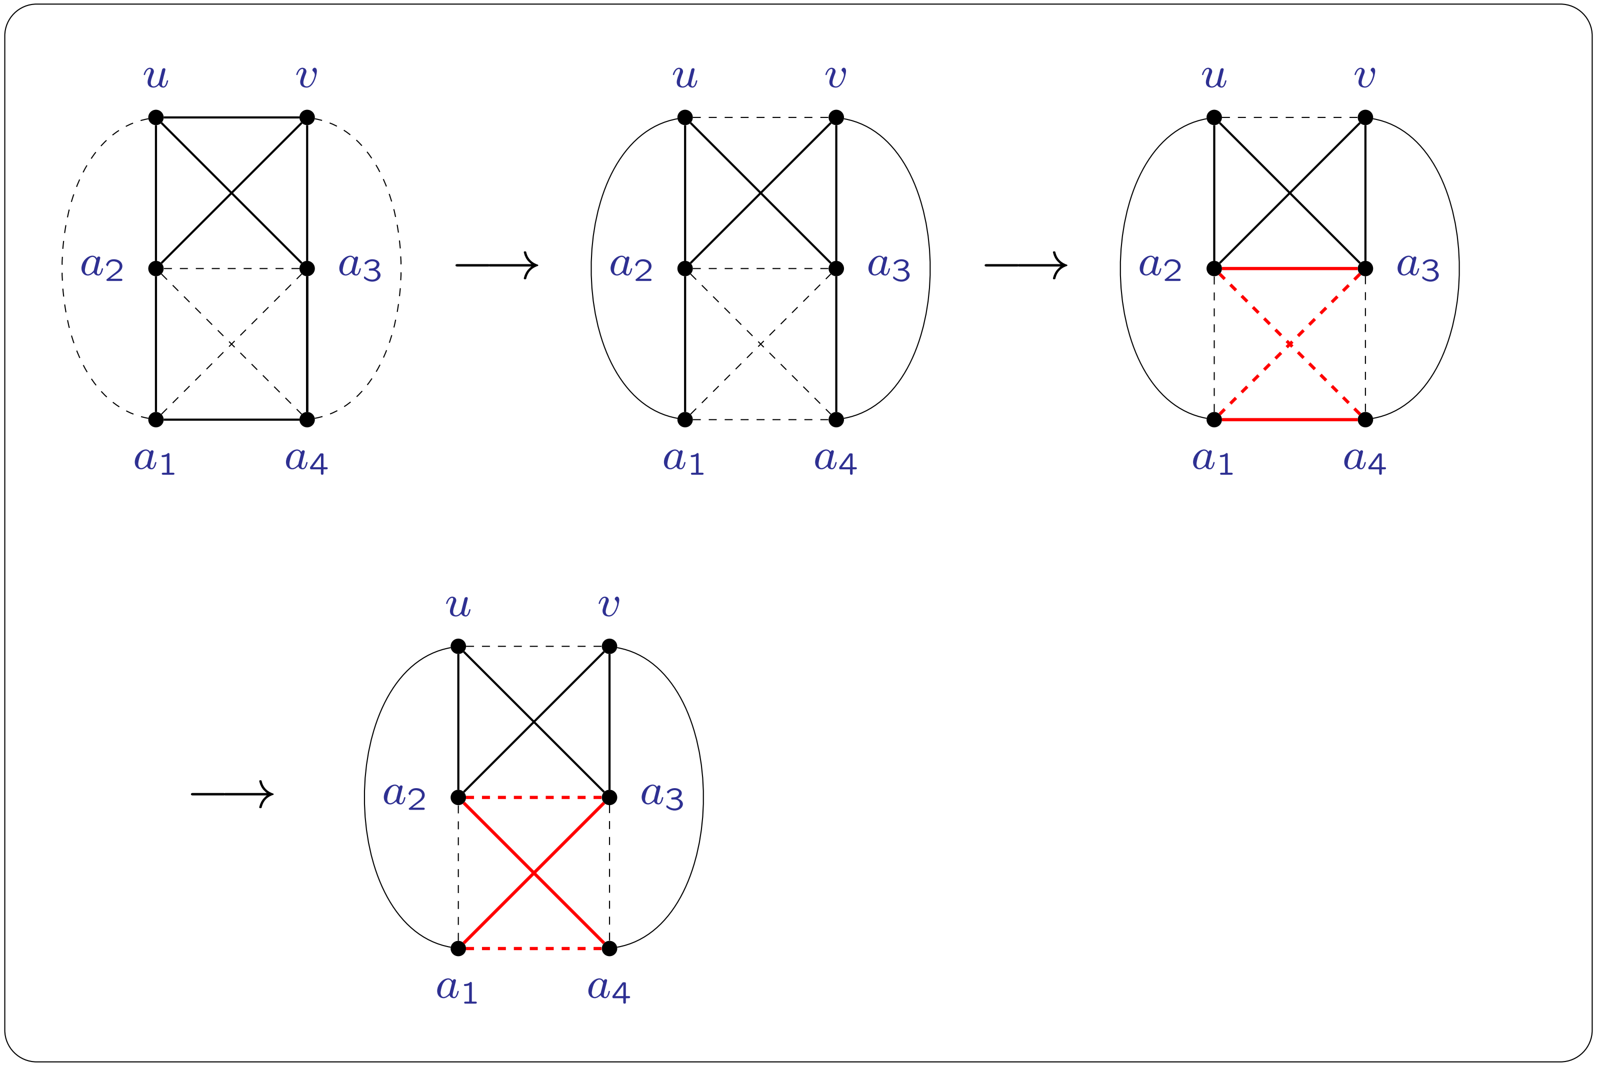

u v a_2 a_3 a_1 a_4
u v a_2 a_3 a_1 a_4
u v a_2 a_3 a_1 a_4
u v a_2 a_3 a_1 a_4

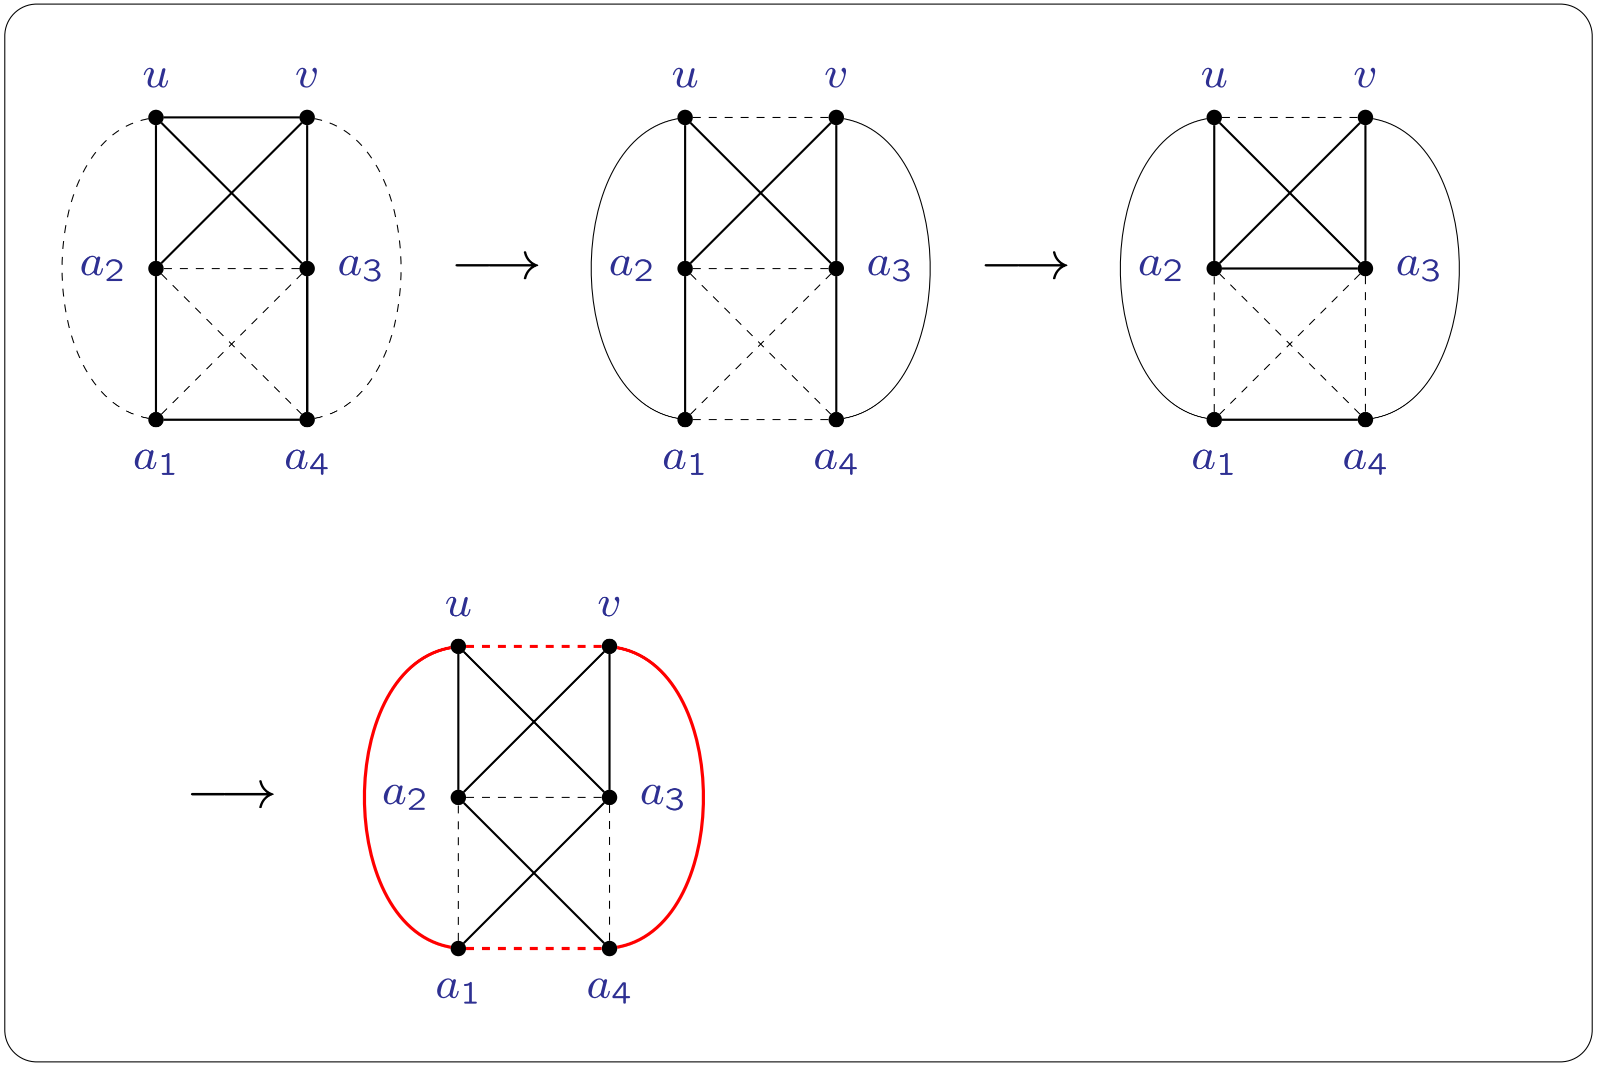

u
v
a_2
a_3
a_1
a_4

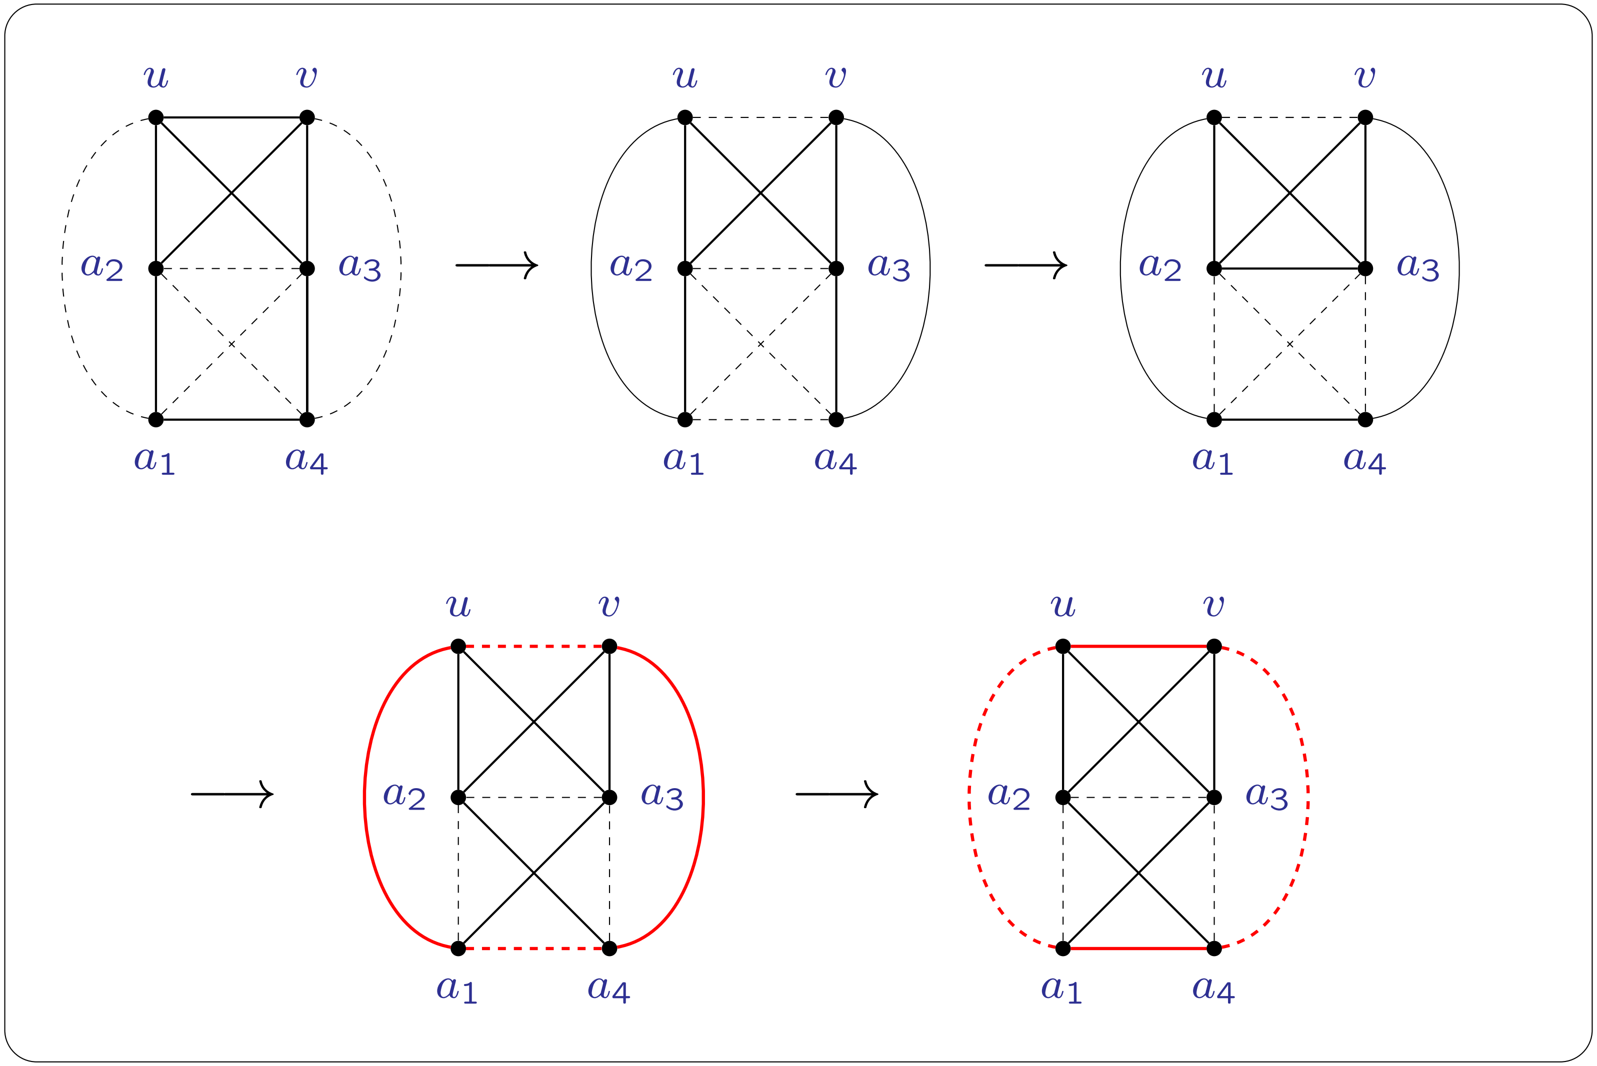

u
v
a_2
a_3
a_1
a_4

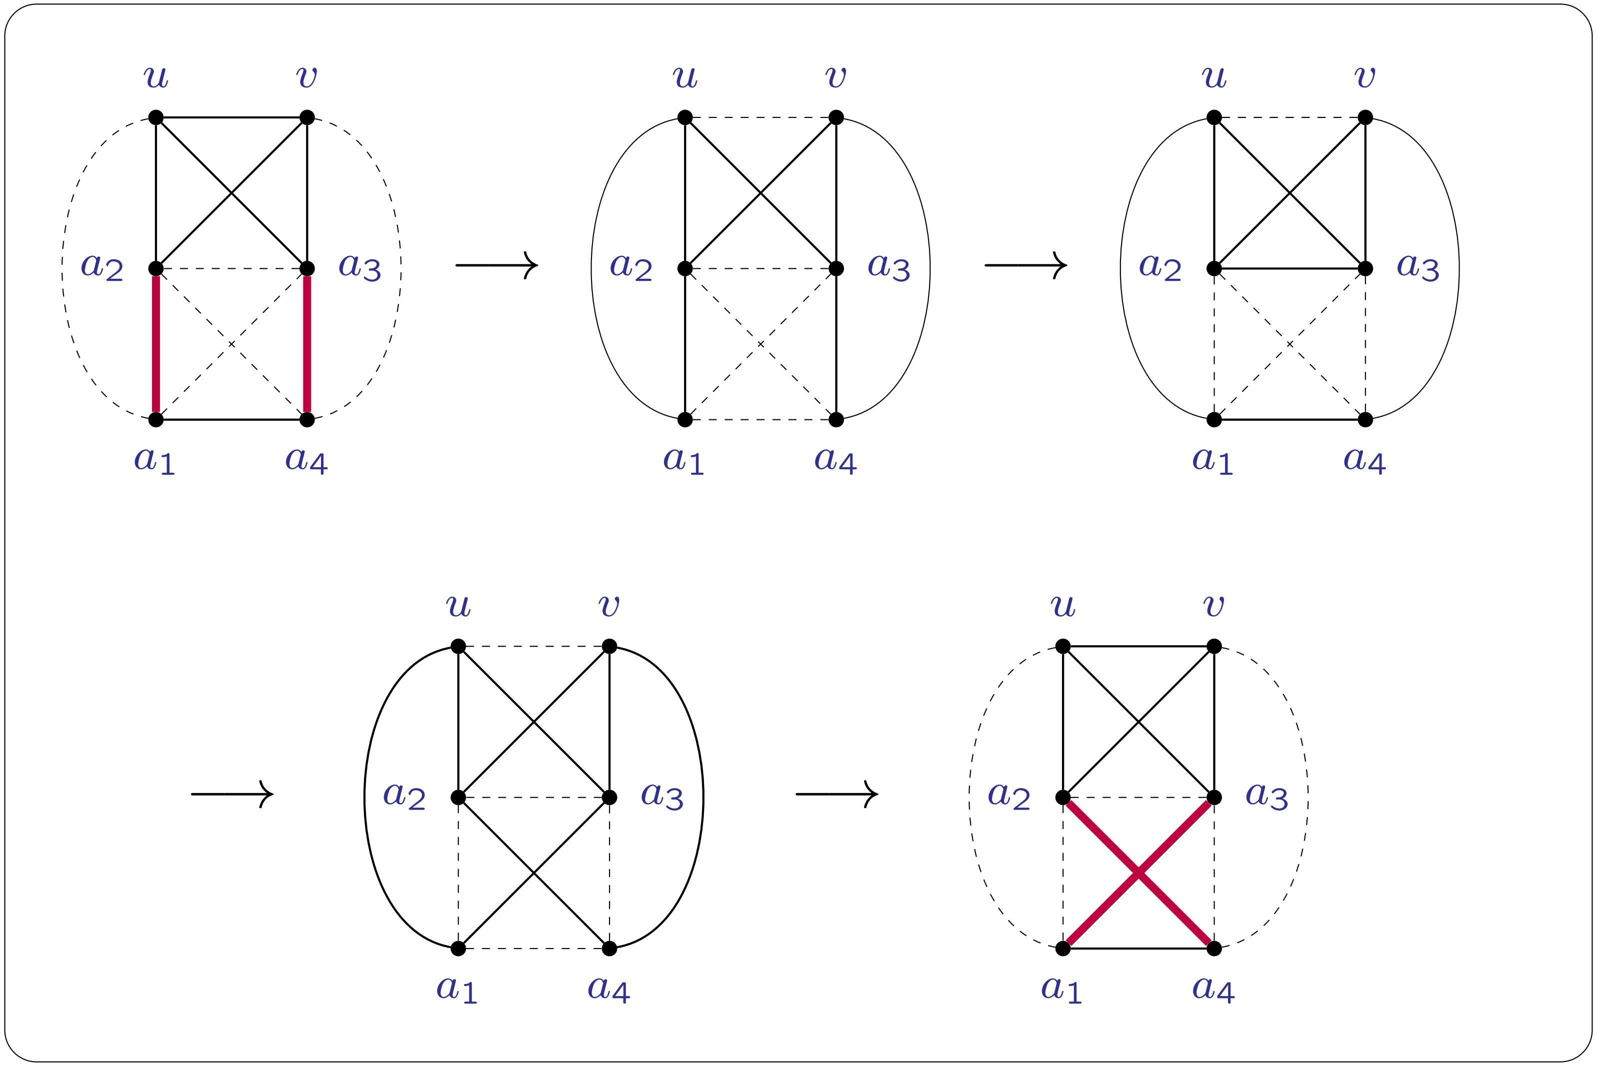

u
v
a2
a3
a1
a4

What about a variant of the switch chain that will encourage triangles? It will need a non-uniform stationary distribution.

What about a variant of the switch chain that will encourage triangles? It will need a non-uniform stationary distribution.

Let $t(G)$ be the number of triangles in the graph G, and fix $\lambda \geq 1$.

What about a variant of the switch chain that will encourage triangles? It will need a non-uniform stationary distribution.

Let $t(G)$ be the number of triangles in the graph G, and fix $\lambda \geq 1$. Let π be the probability distribution on $G(\boldsymbol{d})$ defined by

$$\pi(G) = \lambda^{t(G)}/Z_\lambda(\boldsymbol{d})$$

where

$$Z_\lambda(\boldsymbol{d}) = \sum_{G \in G(\boldsymbol{d})} \lambda^{t(G)}.$$

What about a variant of the switch chain that will encourage triangles? It will need a non-uniform stationary distribution.

Let $t(G)$ be the number of triangles in the graph G, and fix $\lambda \geq 1$. Let π be the probability distribution on $G(\boldsymbol{d})$ defined by

$$\pi(G) = \lambda^{t(G)}/Z_\lambda(\boldsymbol{d})$$

where

$$Z_\lambda(\boldsymbol{d}) = \sum_{G \in G(\boldsymbol{d})} \lambda^{t(G)}.$$

If $\lambda > 1$ then triangles are encouraged.

Metropolis $\triangle$-switch chain with parameter $\lambda \geq 1$

From current graph $G \in G(\boldsymbol{d})$:

- choose two non-adjacent edges uv, yz u.a.r.
- choose a perfect matching M of u, v, y, z u.a.r.
- let H be obtained from G by deleting edges uv, yz and inserting M
- if $G \mapsto H$ is a valid $\triangle$-switch then

  the next state is H with probability $\min\left\{1,\ \lambda^{t(H)-t(G)}\right\}$

  otherwise, stay at G.

For any $G \in \mathcal{G}(\boldsymbol{d})$, the number of choices of two non-adjacent edges in G is exactly

$$a(\boldsymbol{d}) = \binom{M/2}{2} - \frac{1}{2}M_2.$$

For any $G \in G(\boldsymbol{d})$, the number of choices of two non-adjacent edges in G is exactly

$$a(\boldsymbol{d}) = \binom{M/2}{2} - \tfrac{1}{2} M_2.$$

If G and H differ by one $\triangle$-switch then the corresponding transition probability is

$$P(G, H) = \frac{1}{3\, a(\boldsymbol{d})} \min \left\{ 1,\, \lambda^{t(H)-t(G)} \right\}.$$

For any $G \in G(\boldsymbol{d})$, the number of choices of two non-adjacent edges in G is exactly

$$a(\boldsymbol{d}) = \binom{M/2}{2} - \tfrac{1}{2}M_2.$$

If G and H differ by one $\triangle$-switch then the corresponding transition probability is

$$P(G,H) = \frac{1}{3\,a(\boldsymbol{d})}\,\min\left\{1,\, \lambda^{t(H)-t(G)}\right\}.$$

Hence the Metropolis $\triangle$-switch chain satisfies the detailed balance equations

$$\pi(G)\,P(G,H) = \pi(H)\,P(H,G) \qquad \text{for all } G,H \in G(\boldsymbol{d}).$$

For any $G \in G(\boldsymbol{d})$, the number of choices of two non-adjacent edges in G is exactly

$$a(\boldsymbol{d}) = \binom{M/2}{2} - \tfrac{1}{2} M_2.$$

If G and H differ by one $\triangle$-switch then the corresponding transition probability is

$$P(G,H) = \frac{1}{3\,a(\boldsymbol{d})}\,\min\left\{1,\ \lambda^{t(H)-t(G)}\right\}.$$

Hence the Metropolis $\triangle$-switch chain satisfies the detailed balance equations

$$\pi(G)\,P(G,H) = \pi(H)\,P(H,G) \qquad \text{for all } G,H \in G(\boldsymbol{d}).$$

This implies that π is the unique stationary distribution of the chain.

Recall that $\mu = \mu(\boldsymbol{d})$ is the expected number of triangles in a uniformly random element of $G(\boldsymbol{d})$.

We proved that if $d_{\mathrm{max}} \log \lambda = o(\log n)$ then for G drawn from the distribution π on $G(\boldsymbol{d})$,

$$\Pr(t(G) = s) \sim \frac{e^{-\lambda\,\mu}(\lambda\mu)^s}{s!}$$

if s is not too large (more precisely, if $s = o(n^{\varepsilon})$).

Recall that $\mu = \mu(\boldsymbol{d})$ is the expected number of triangles in a uniformly random element of $G(\boldsymbol{d})$.

We proved that if $d_{\text{max}} \log \lambda = o(\log n)$ then for G drawn from the distribution π on $G(\boldsymbol{d})$,

$$\Pr(t(G) = s) \sim \frac{e^{-\lambda\mu}(\lambda\mu)^s}{s!}$$

if s is not too large (more precisely, if $s = o(n^{\varepsilon})$).

So in the stationary distribution of the Metropolis $\triangle$-switch chain, the number of triangles is (roughly) asymptotically Poisson with mean $\lambda\mu$.

Recall that $\mu = \mu(\boldsymbol{d})$ is the expected number of triangles in a uniformly random element of $G(\boldsymbol{d})$.

We proved that if $d_{\max} \log \lambda = o(\log n)$ then for G drawn from the distribution π on $G(\boldsymbol{d})$,

$$\Pr(t(G) = s) \sim \frac{e^{-\lambda\mu}(\lambda\mu)^s}{s!}$$

if s is not too large (more precisely, if $s = o(n^\varepsilon)$).

So in the stationary distribution of the Metropolis $\triangle$-switch chain, the number of triangles is (roughly) asymptotically Poisson with mean $\lambda\mu$.

$\Rightarrow$ It looks like we should take λ as large as possible?

[First we looked at the distribution of triangles under the uniform distribution on $G(\boldsymbol{d})$, using a switching argument to prove asymptotic Poisson-ness.

Prior work e.g. McKay, Wormald & Wysocka (2004), Gao (2021) for the regular case.]

Here is our main mixing/convergence result.

Theorem Let $\mathcal{D}$ be a family of graphical sequences which all have minimum degree at least 3. Let $\lambda = \lambda(n) \geq 1$. Suppose that there exists $\alpha \in (0, 1)$ such that

$$\lambda\, \mu(\boldsymbol{d}) \leq \log^{\alpha} n$$

for every $\boldsymbol{d} \in \mathcal{D}$ of length n.

Here is our main mixing/convergence result.

Theorem Let $\mathcal{D}$ be a family of graphical sequences which all have minimum degree at least 3. Let $\lambda = \lambda(n) \geq 1$. Suppose that there exists $\alpha \in (0, 1)$ such that

$$\lambda\, \mu(\boldsymbol{d}) \leq \log^{\alpha} n$$

for every $\boldsymbol{d} \in \mathcal{D}$ of length n.

If the switch chain is rapidly mixing on $G(\boldsymbol{d})$ for all $\boldsymbol{d} \in \mathcal{D}$ then the same is true for the modified Metropolis $\triangle$-switch chain with parameter λ.

Here is our main mixing/convergence result.

Theorem Let $\mathcal{D}$ be a family of graphical sequences which all have minimum degree at least 3. Let $\lambda = \lambda(n) \geq 1$. Suppose that there exists $\alpha \in (0, 1)$ such that

$$\lambda\, \mu(\boldsymbol{d}) \leq \log^{\alpha} n$$

for every $\boldsymbol{d} \in \mathcal{D}$ of length n.

If the switch chain is rapidly mixing on $G(\boldsymbol{d})$ for all $\boldsymbol{d} \in \mathcal{D}$ then the same is true for the modified Metropolis $\triangle$-switch chain with parameter λ.

(I'll explain "modified" soon.)

Proof. If the switch chain is rapidly mixing on $G(\boldsymbol{d})$ then there is a set Γ' of canonical paths for $G(\boldsymbol{d})$ with low congestion. (See Sinclair 1992, Guruswami 2000.)

Proof. If the switch chain is rapidly mixing on $G(\boldsymbol{d})$ then there is a set Γ' of canonical paths for $G(\boldsymbol{d})$ with low congestion. (See Sinclair 1992, Guruswami 2000.)

Replacing each switch in these paths by the corresponding simulation path of at most 5 $\triangle$-switches gives a set of canonical paths Γ for the Metropolis $\triangle$-switch chain.

Proof. If the switch chain is rapidly mixing on $G(\boldsymbol{d})$ then there is a set Γ' of canonical paths for $G(\boldsymbol{d})$ with low congestion. (See Sinclair 1992, Guruswami 2000.)

Replacing each switch in these paths by the corresponding simulation path of at most 5 $\triangle$-switches gives a set of canonical paths Γ for the Metropolis $\triangle$-switch chain.

Applying a result from Cooper, Dyer, Greenhill, Handley (2019) [with an error fixed!] to the Metropolis $\triangle$-switch chain guarantees that the congestion of Γ can be no bigger than the congestion of Γ' times an adjustment factor of

$$100\, d_{\max}^2 \left(2M + d_{\max}^2\right) n^{2\alpha}\, \frac{Z_\lambda(\boldsymbol{d})}{|G(\boldsymbol{d})|}.$$

Machinery from Cooper, Dyer, Greenhill, Handley (2019):
the two-stage direct canonical path construction method.

Machinery from Cooper, Dyer, Greenhill, Handley (2019): the two-stage direct canonical path construction method.

Say we have canonical paths Γ' for a Markov chain $\mathcal{M}'$, and we have a simulation path for each transition of $\mathcal{M}'$ using a sequence of transitions of $\mathcal{M}$. This gives a set Γ of canonical paths for $\mathcal{M}$.

Machinery from Cooper, Dyer, Greenhill, Handley (2019): the two-stage direct canonical path construction method.

Say we have canonical paths Γ' for a Markov chain $\mathcal{M}'$, and we have a simulation path for each transition of $\mathcal{M}'$ using a sequence of transitions of $\mathcal{M}$. This gives a set Γ of canonical paths for $\mathcal{M}$.

Let P, π denote the transition matrix and stationary distribution of $\mathcal{M}$, and similarly for $\mathcal{M}'$.

(We assume here that $\mathcal{M}$ and $\mathcal{M}'$ have the same state space.)

Extra factors in the congestion bound for Γ:

- maximum length of the simulation paths

- maximum number of simulation paths through a transition of $\mathcal{M}$

- simulation gap $\max \frac{\pi'(z)\, P'(z,w)}{\pi(u)\, P(u,v)}$

- stationary ratio $\left(\max \frac{\pi(x)}{\pi'(x)}\right)^2$

(In CDGH we forgot about the stationary ratio.)

Example: In how many ways can a $\triangle$-switch (X, Y) be used in a simulation path in Case V? If the switch is $G \mapsto H$ then the simulation path is $G \mapsto G_1 \mapsto G_2 \mapsto G_3 \mapsto H$.

Example: In how many ways can a $\triangle$-switch (X, Y) be used in a simulation path in Case V? If the switch is $G \mapsto H$ then the simulation path is $G \mapsto G_1 \mapsto G_2 \mapsto G_3 \mapsto H$.

First step: $(X, Y) = (G, G_1)$.

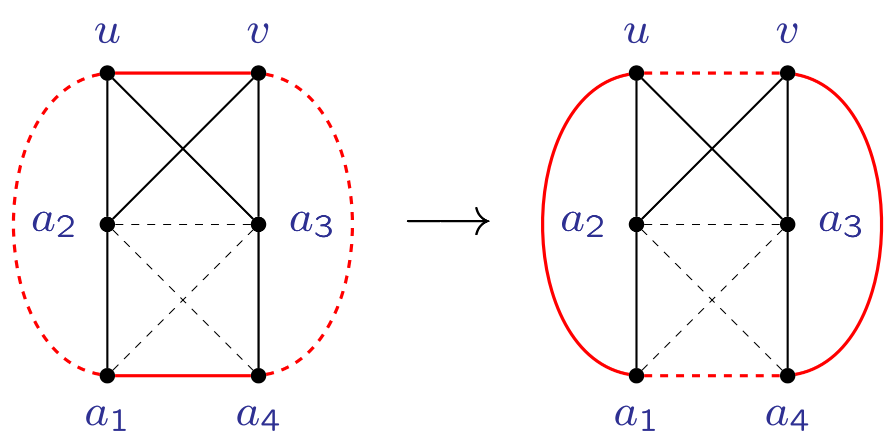

Example: In how many ways can a $\triangle$-switch (X, Y) be used in a simulation path in Case V? If the switch is $G \mapsto H$ then the simulation path is $G \mapsto G_1 \mapsto G_2 \mapsto G_3 \mapsto H$.

First step: $(X, Y) = (G, G_1)$.

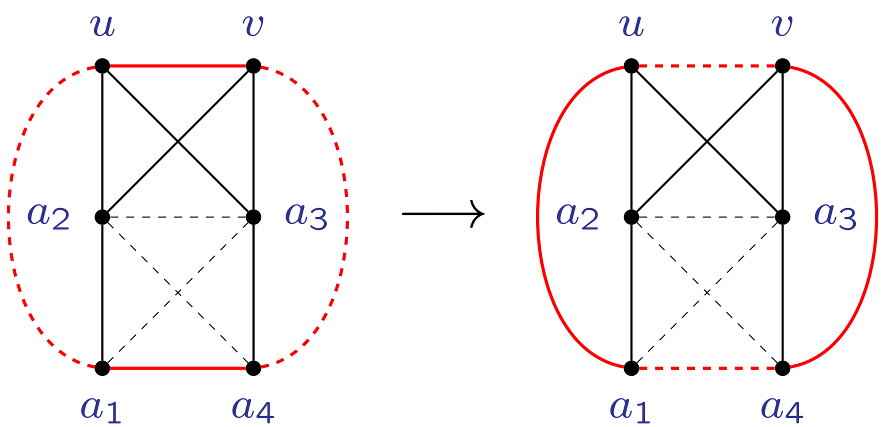

- 2 choices for which deleted edge is $a_1 a_4$,
- at most $d_{\max}^2$ choices for a_2, a_3.

Second step: $(X, Y) = (G_1, G_2)$.

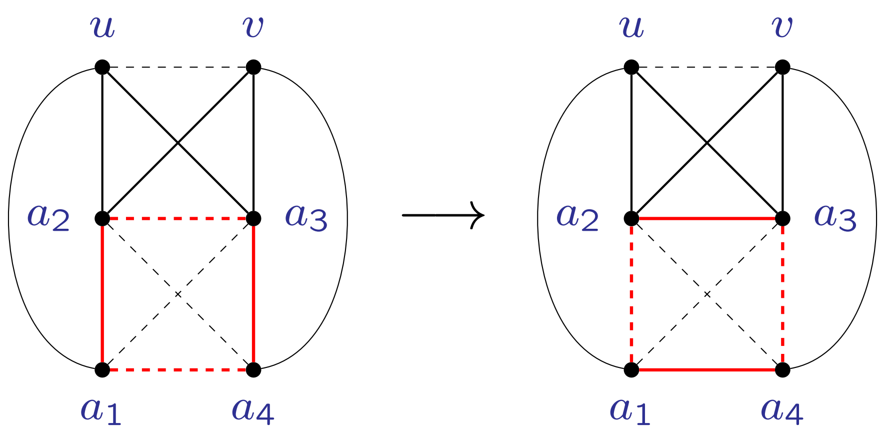

Second step: $(X, Y) = (G_1, G_2)$.

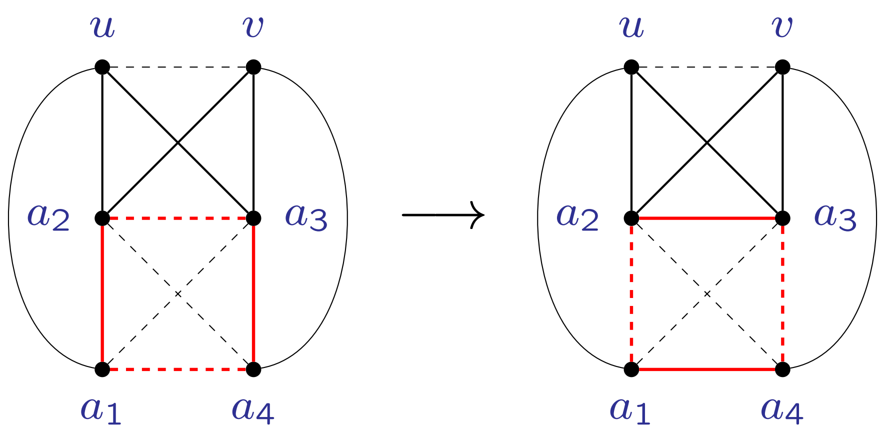

- All switch vertices uniquely identified (up to symmetry),
- at most $d_{\max}^2$ choices for u, v, needed to reconstruct G from G_1.

Second step: $(X, Y) = (G_1, G_2)$.

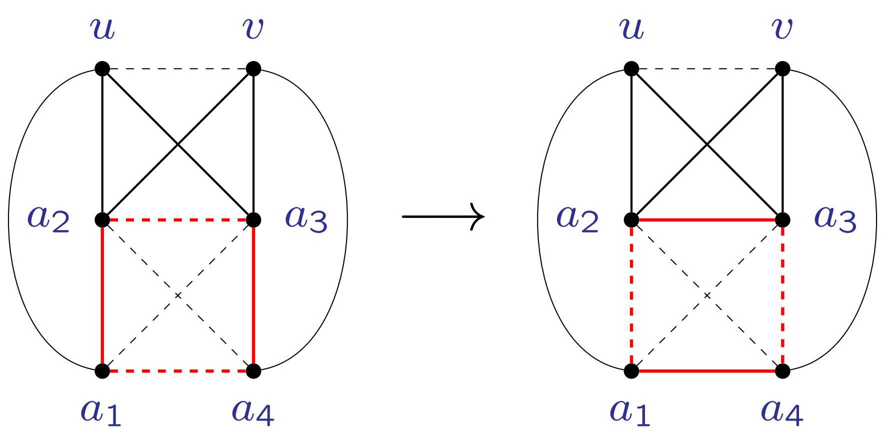

- All switch vertices uniquely identified (up to symmetry),
- at most $d_{\max}^2$ choices for u, v, needed to reconstruct G from G_1.

Overall a $\triangle$-switch can be part of a simulation path in Case (V) in at most $6d_{\max}^2$ ways.

The simulation gap is bounded above by $Z_\lambda(\boldsymbol{d})/|G(\boldsymbol{d})|$, which is precisely the expected value of $\lambda^{t(G)}$ when $G \in G(\boldsymbol{d})$ is chosen uniformly at random.

The simulation gap is bounded above by $Z_\lambda(\boldsymbol{d})/|\mathcal{G}(\boldsymbol{d})|$, which is precisely the expected value of $\lambda^{t(G)}$ when $G \in \mathcal{G}(\boldsymbol{d})$ is chosen uniformly at random.

Problem:
If λ is large then the upper tail of the distribution can contribute significantly, and this factor will not be polynomially-bounded.

The simulation gap is bounded above by $Z_\lambda(\boldsymbol{d})/|G(\boldsymbol{d})|$, which is precisely the expected value of $\lambda^{t(G)}$ when $G \in G(\boldsymbol{d})$ is chosen uniformly at random.

Problem:
If λ is large then the upper tail of the distribution can contribute significantly, and this factor will not be polynomially-bounded.

Our solution: Modify the Metropolis $\triangle$-switch chain so that the stationary distribution is proportional to $\lambda^{\min\{t(G),\nu\}}$, where $\nu := \log n/(\log\log n)$.

The simulation gap is bounded above by $Z_\lambda(\boldsymbol{d})/|G(\boldsymbol{d})|$, which is precisely the expected value of $\lambda^{t(G)}$ when $G \in G(\boldsymbol{d})$ is chosen uniformly at random.

Problem:
If λ is large then the upper tail of the distribution can contribute significantly, and this factor will not be polynomially-bounded.

Our solution: Modify the Metropolis △-switch chain so that the stationary distribution is proportional to $\lambda^{\min\{t(G),\nu\}}$, where $\nu := \log n/(\log\log n)$.

This new distribution is polynomial-time indistinguishable from π.

In the modified Metropolis $\triangle$-switch chain, the simulation gap is the expected value of $\lambda^{\min\{t(G),\nu\}}$, which is at most $\lambda^{\nu} \leq n^{\alpha}$.

In the modified Metropolis $\triangle$-switch chain, the simulation gap is the expected value of $\lambda^{\min\{t(G),\nu\}}$, which is at most $\lambda^{\nu} \leq n^{\alpha}$.

Hence the congestion of the set of canonical paths Γ for the modified Metropolis $\triangle$-switch chain is at most a factor

$$100\, d_{\max}^2 \left(2M + d_{\max}^2\right) n^{3\alpha}$$

worse than the congestion of the set of canonical paths Γ' for the switch chain.

In the modified Metropolis $\triangle$-switch chain, the simulation gap is the expected value of $\lambda^{\min\{t(G),\nu\}}$, which is at most $\lambda^{\nu} \leq n^{\alpha}$.

Hence the congestion of the set of canonical paths Γ for the modified Metropolis $\triangle$-switch chain is at most a factor

$$100\, d_{\max}^{2} \left(2M + d_{\max}^{2}\right) n^{3\alpha}$$

worse than the congestion of the set of canonical paths Γ' for the switch chain.

This factor is polynomial (great!!) and shows up as a factor in the bound on the mixing time.

In the modified Metropolis $\triangle$-switch chain, the simulation gap is the expected value of $\lambda^{\min\{t(G),\nu\}}$, which is at most $\lambda^\nu \leq n^\alpha$.

Hence the congestion of the set of canonical paths Γ for the modified Metropolis $\triangle$-switch chain is at most a factor

$$100\, d_{\max}^2 \left(2M + d_{\max}^2\right) n^{3\alpha}$$

worse than the congestion of the set of canonical paths Γ' for the switch chain.

This factor is polynomial (great!!) and shows up as a factor in the bound on the mixing time.

But, the condition $\lambda\mu(\boldsymbol{d}) \leq \log^\alpha n$ is more restrictive than we would like.

Final questions:

- Is there a better choice than Metropolis for the transitions of a triangle switch chain?

Final questions:

- Is there a better choice than Metropolis for the transitions of a triangle switch chain?

- Is there a better choice for the analysis than comparision to the switch chain?

Final questions:

- Is there a better choice than Metropolis for the transitions of a triangle switch chain?

- Is there a better choice for the analysis than comparision to the switch chain?

[We also claimed that this is the "first rigorous analysis of a Markov chain algorithm for generating graphs from a known non-uniform distribution." If you know a counter-example to this please let me know!]

Final questions:

- Is there a better choice than Metropolis for the transitions of a triangle switch chain?

- Is there a better choice for the analysis than comparision to the switch chain?

[We also claimed that this is the "first rigorous analysis of a Markov chain algorithm for generating graphs from a known non-uniform distribution." If you know a counter-example to this please let me know!]

$*$ Thank you! $*$