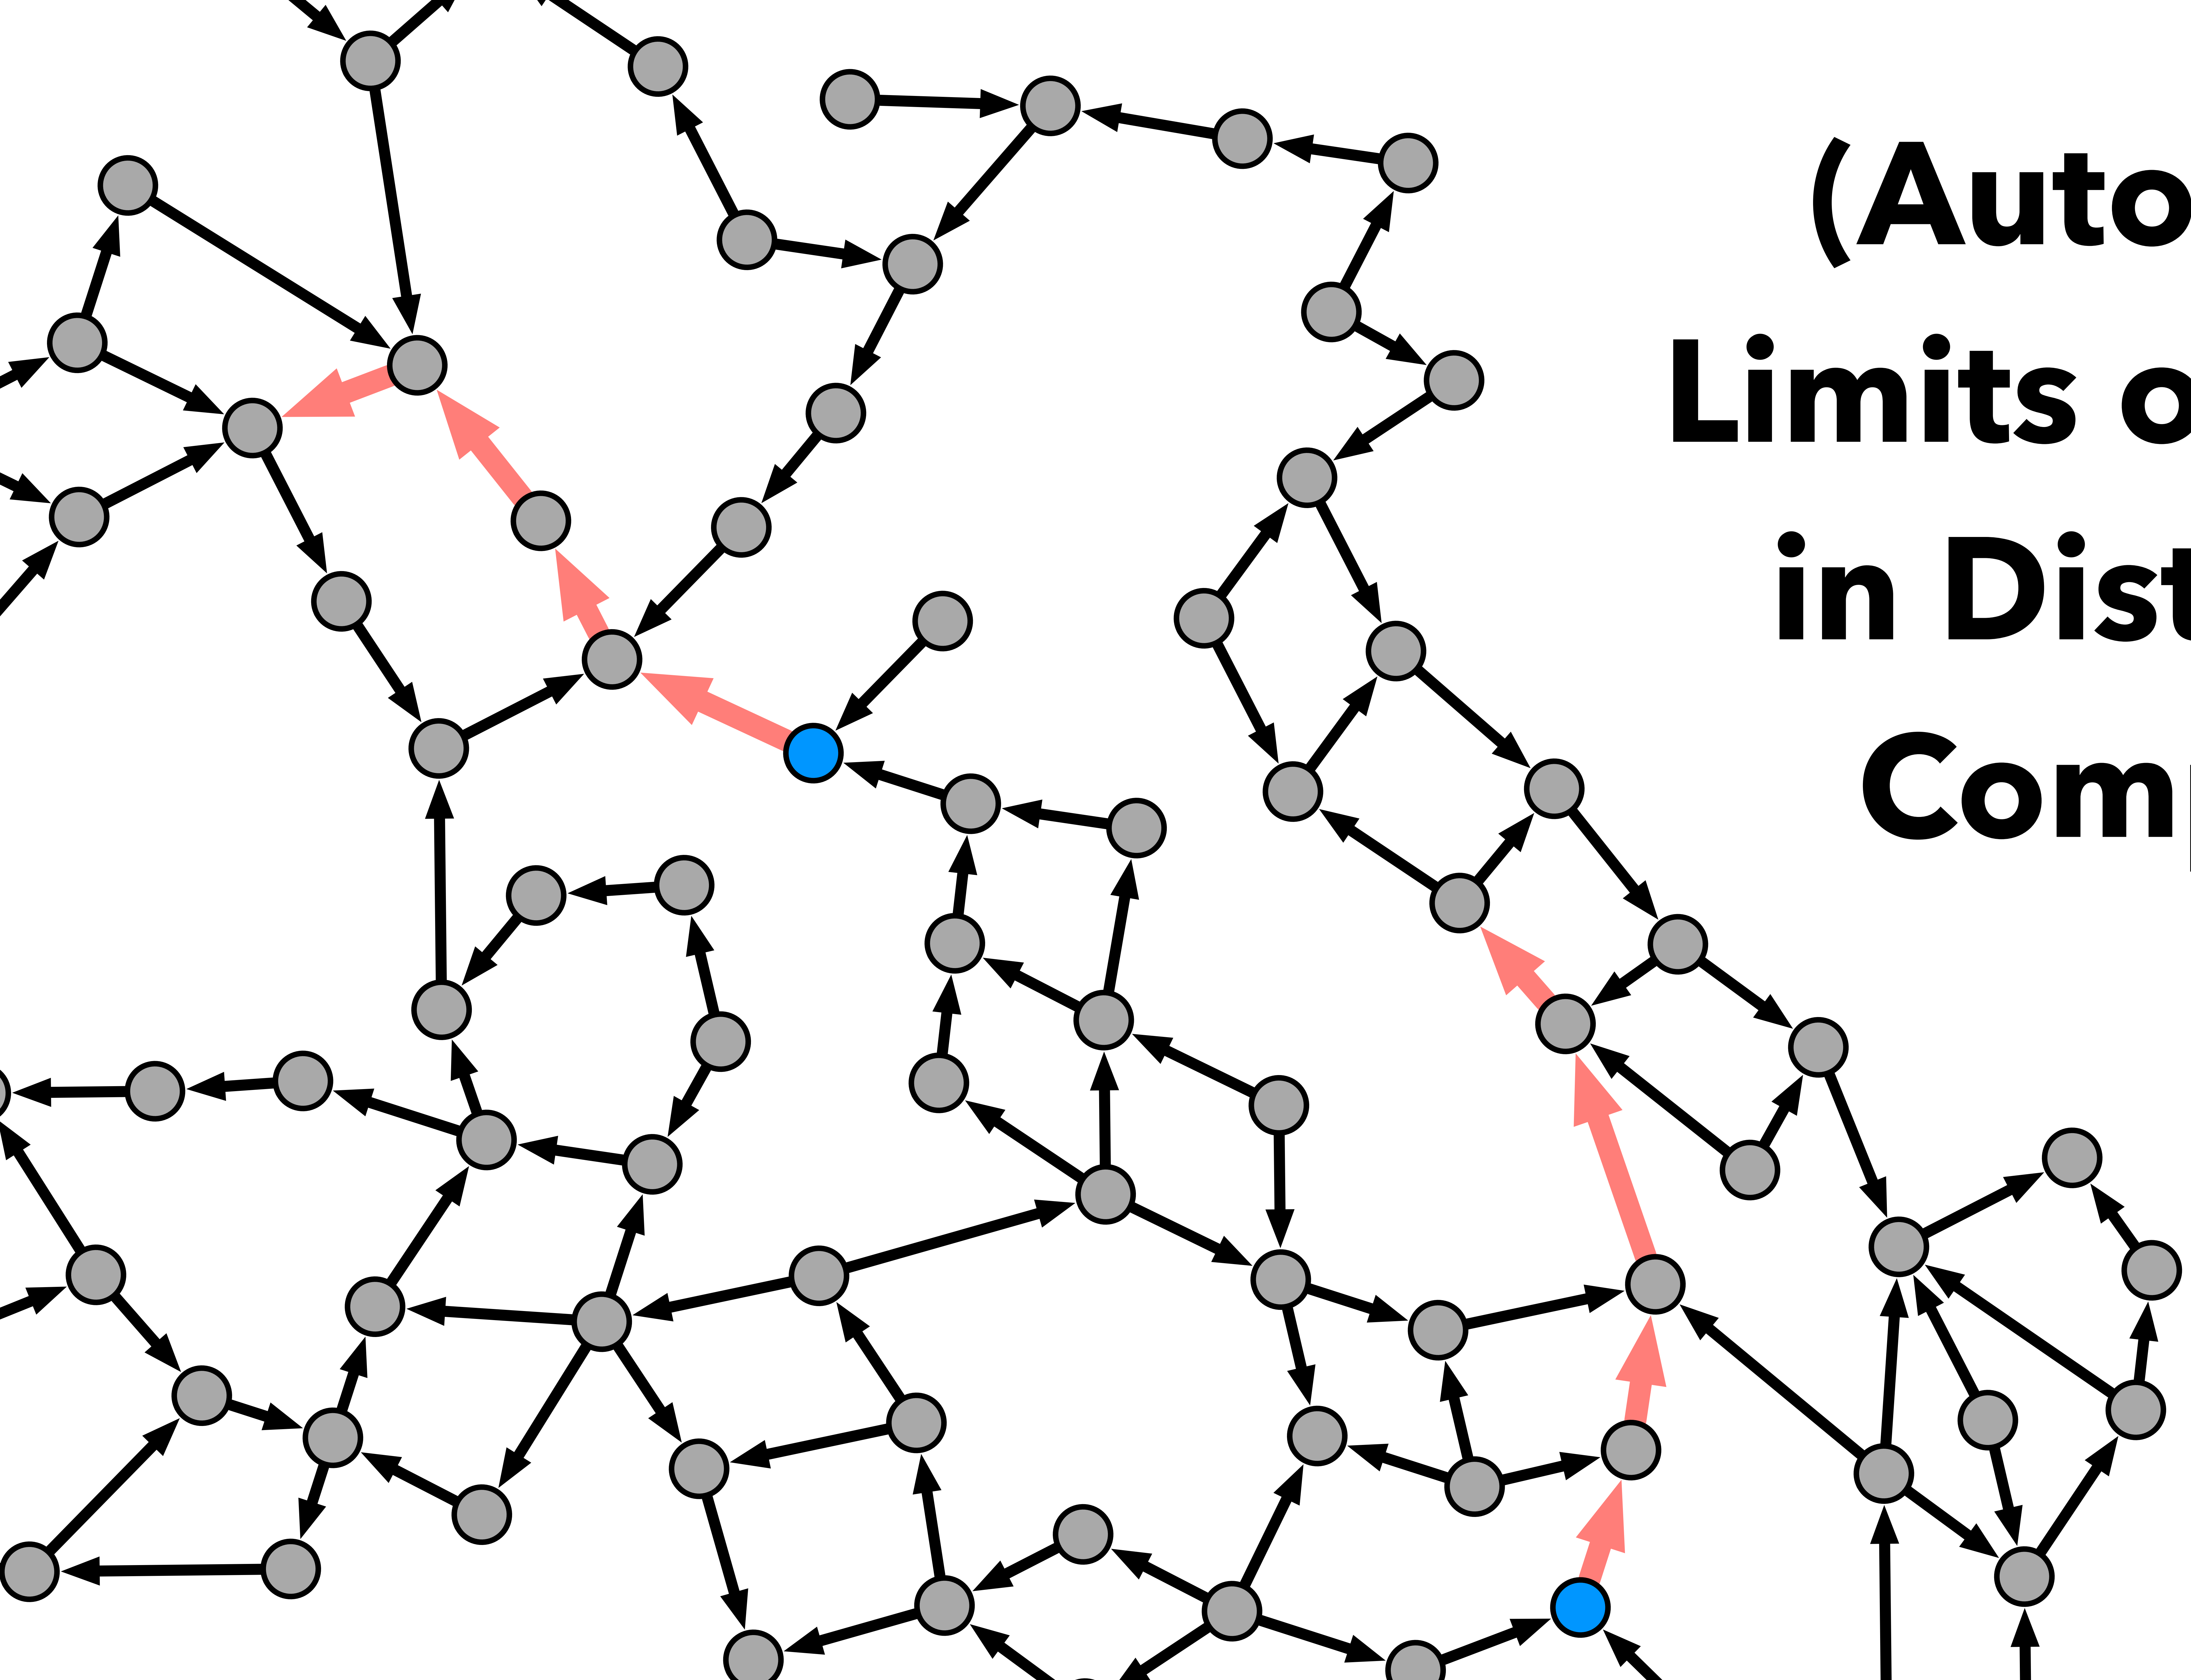

(Automated) Limits of Locality in Distributed Computing

Aalto CS Research Day 2019
Juho Hirvonen
Dennis Olivetti

Distributed Algorithms

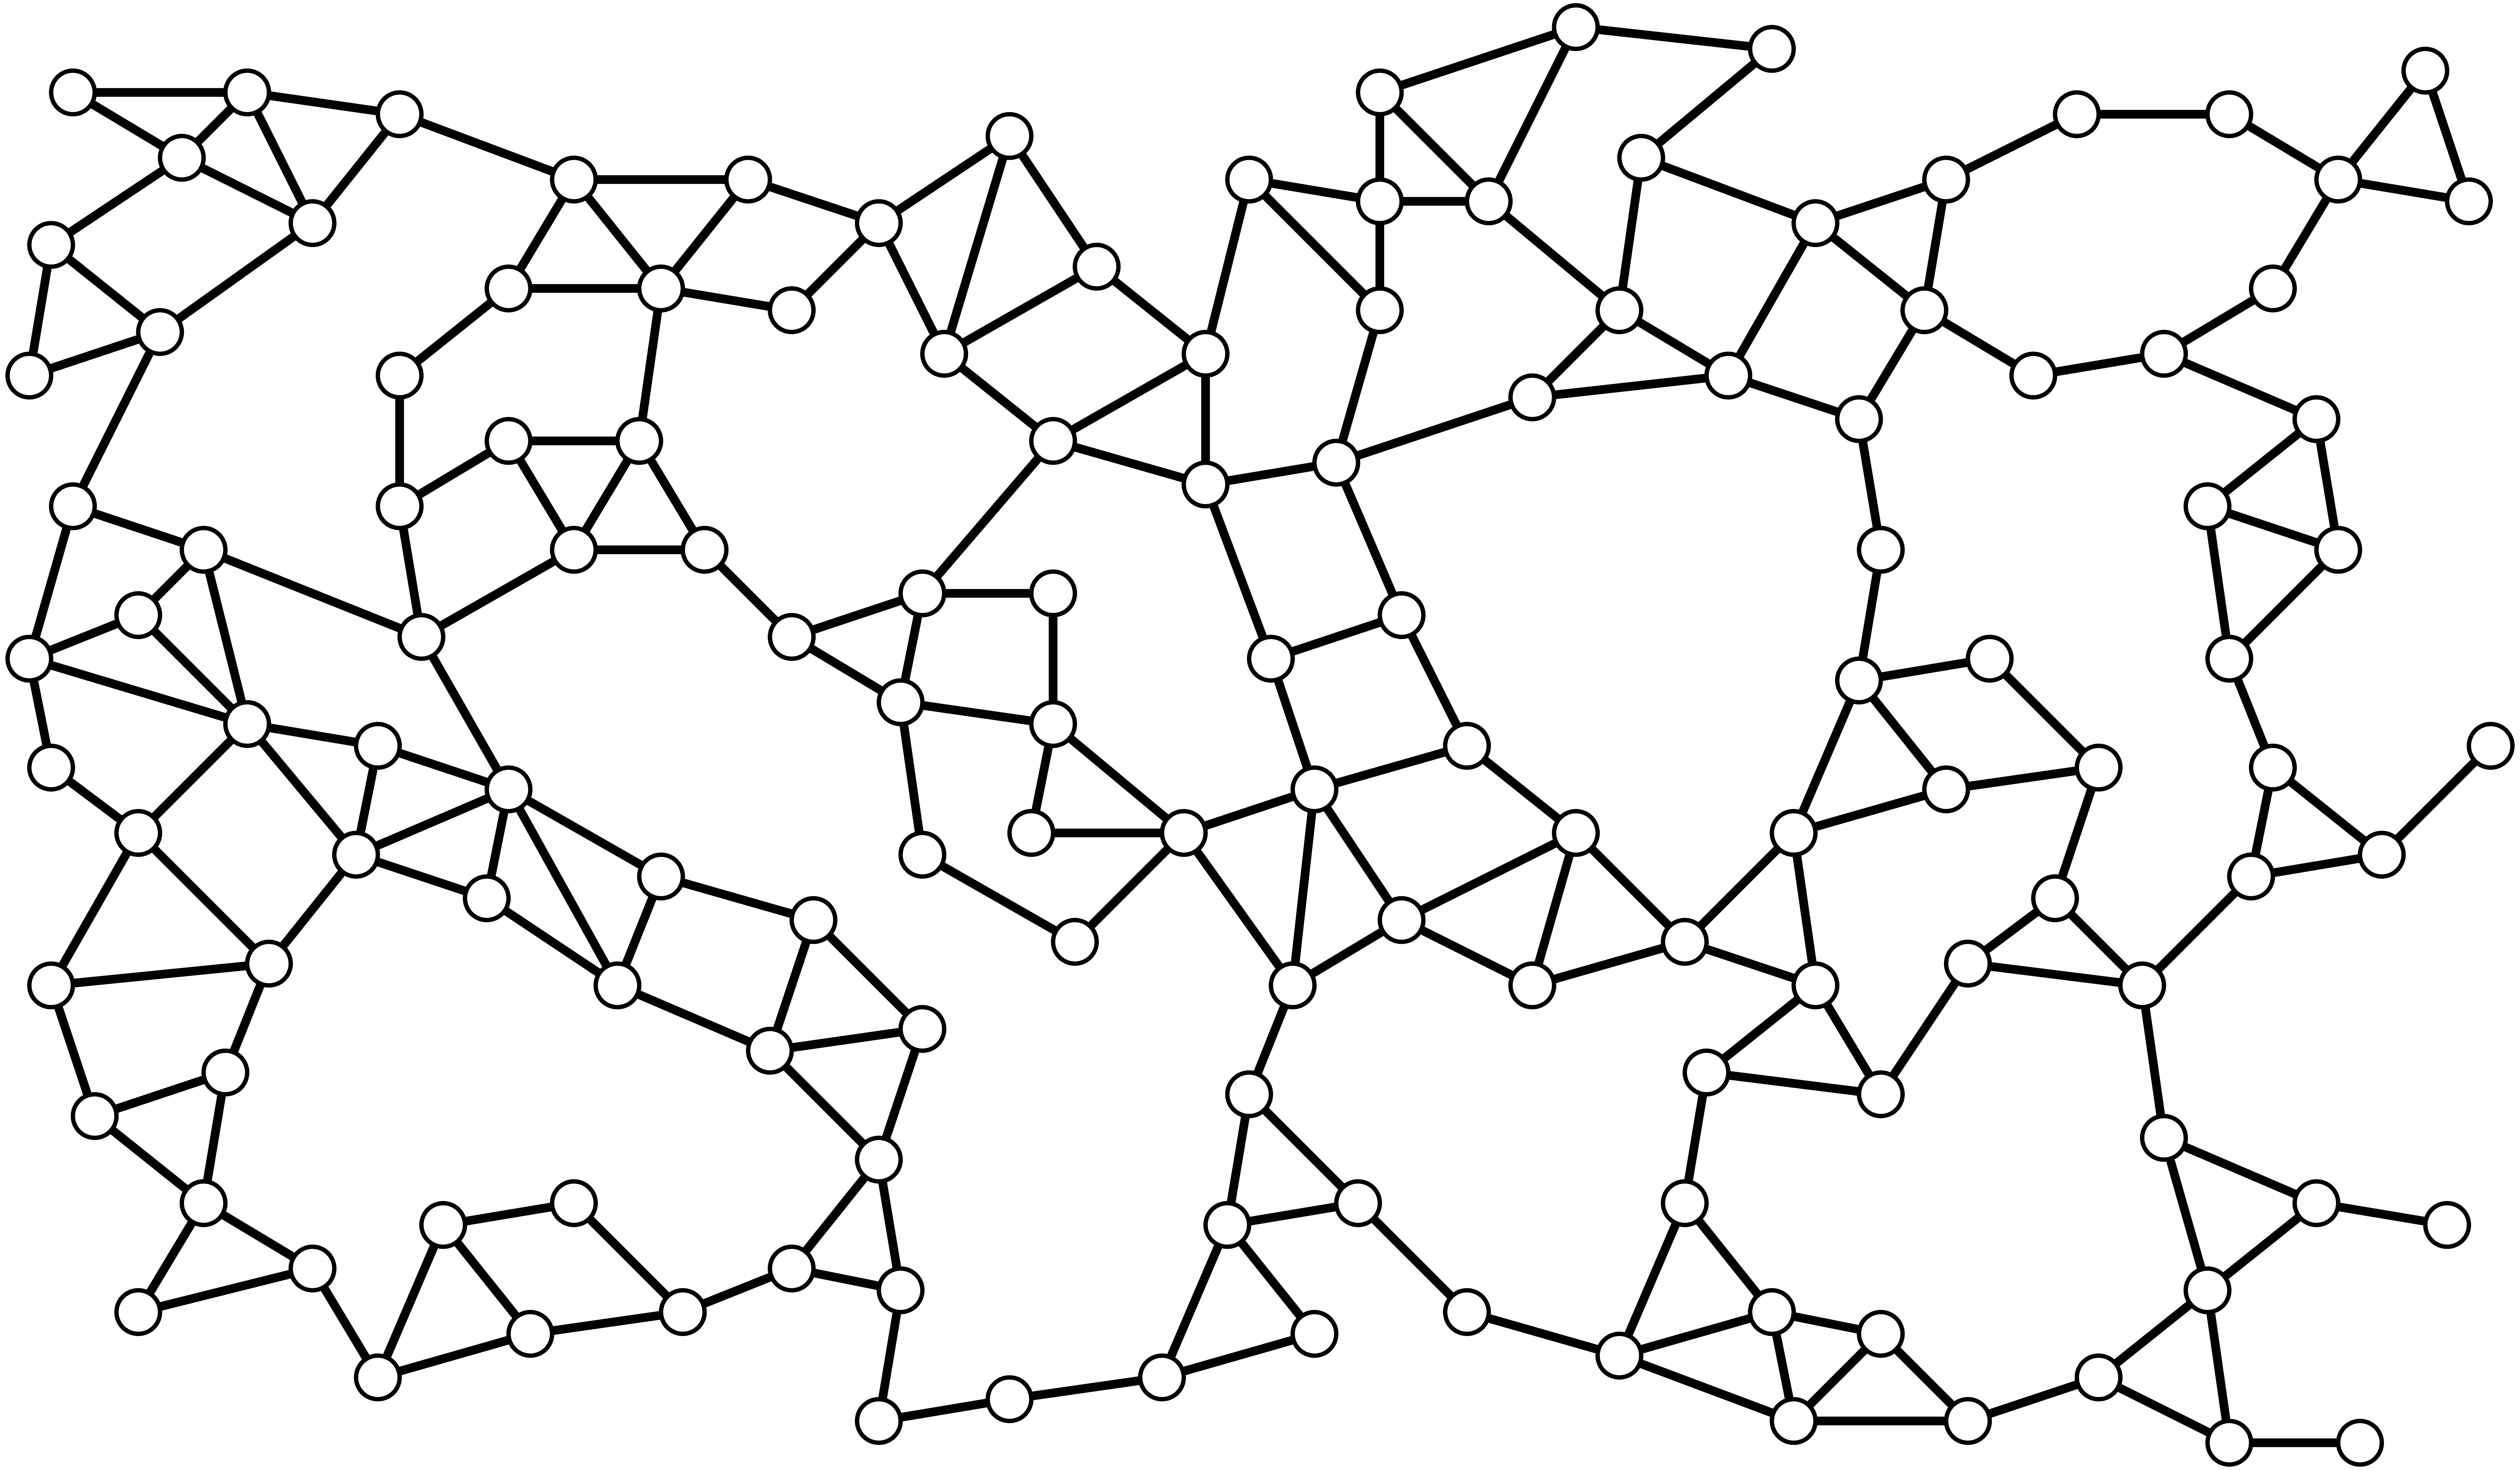

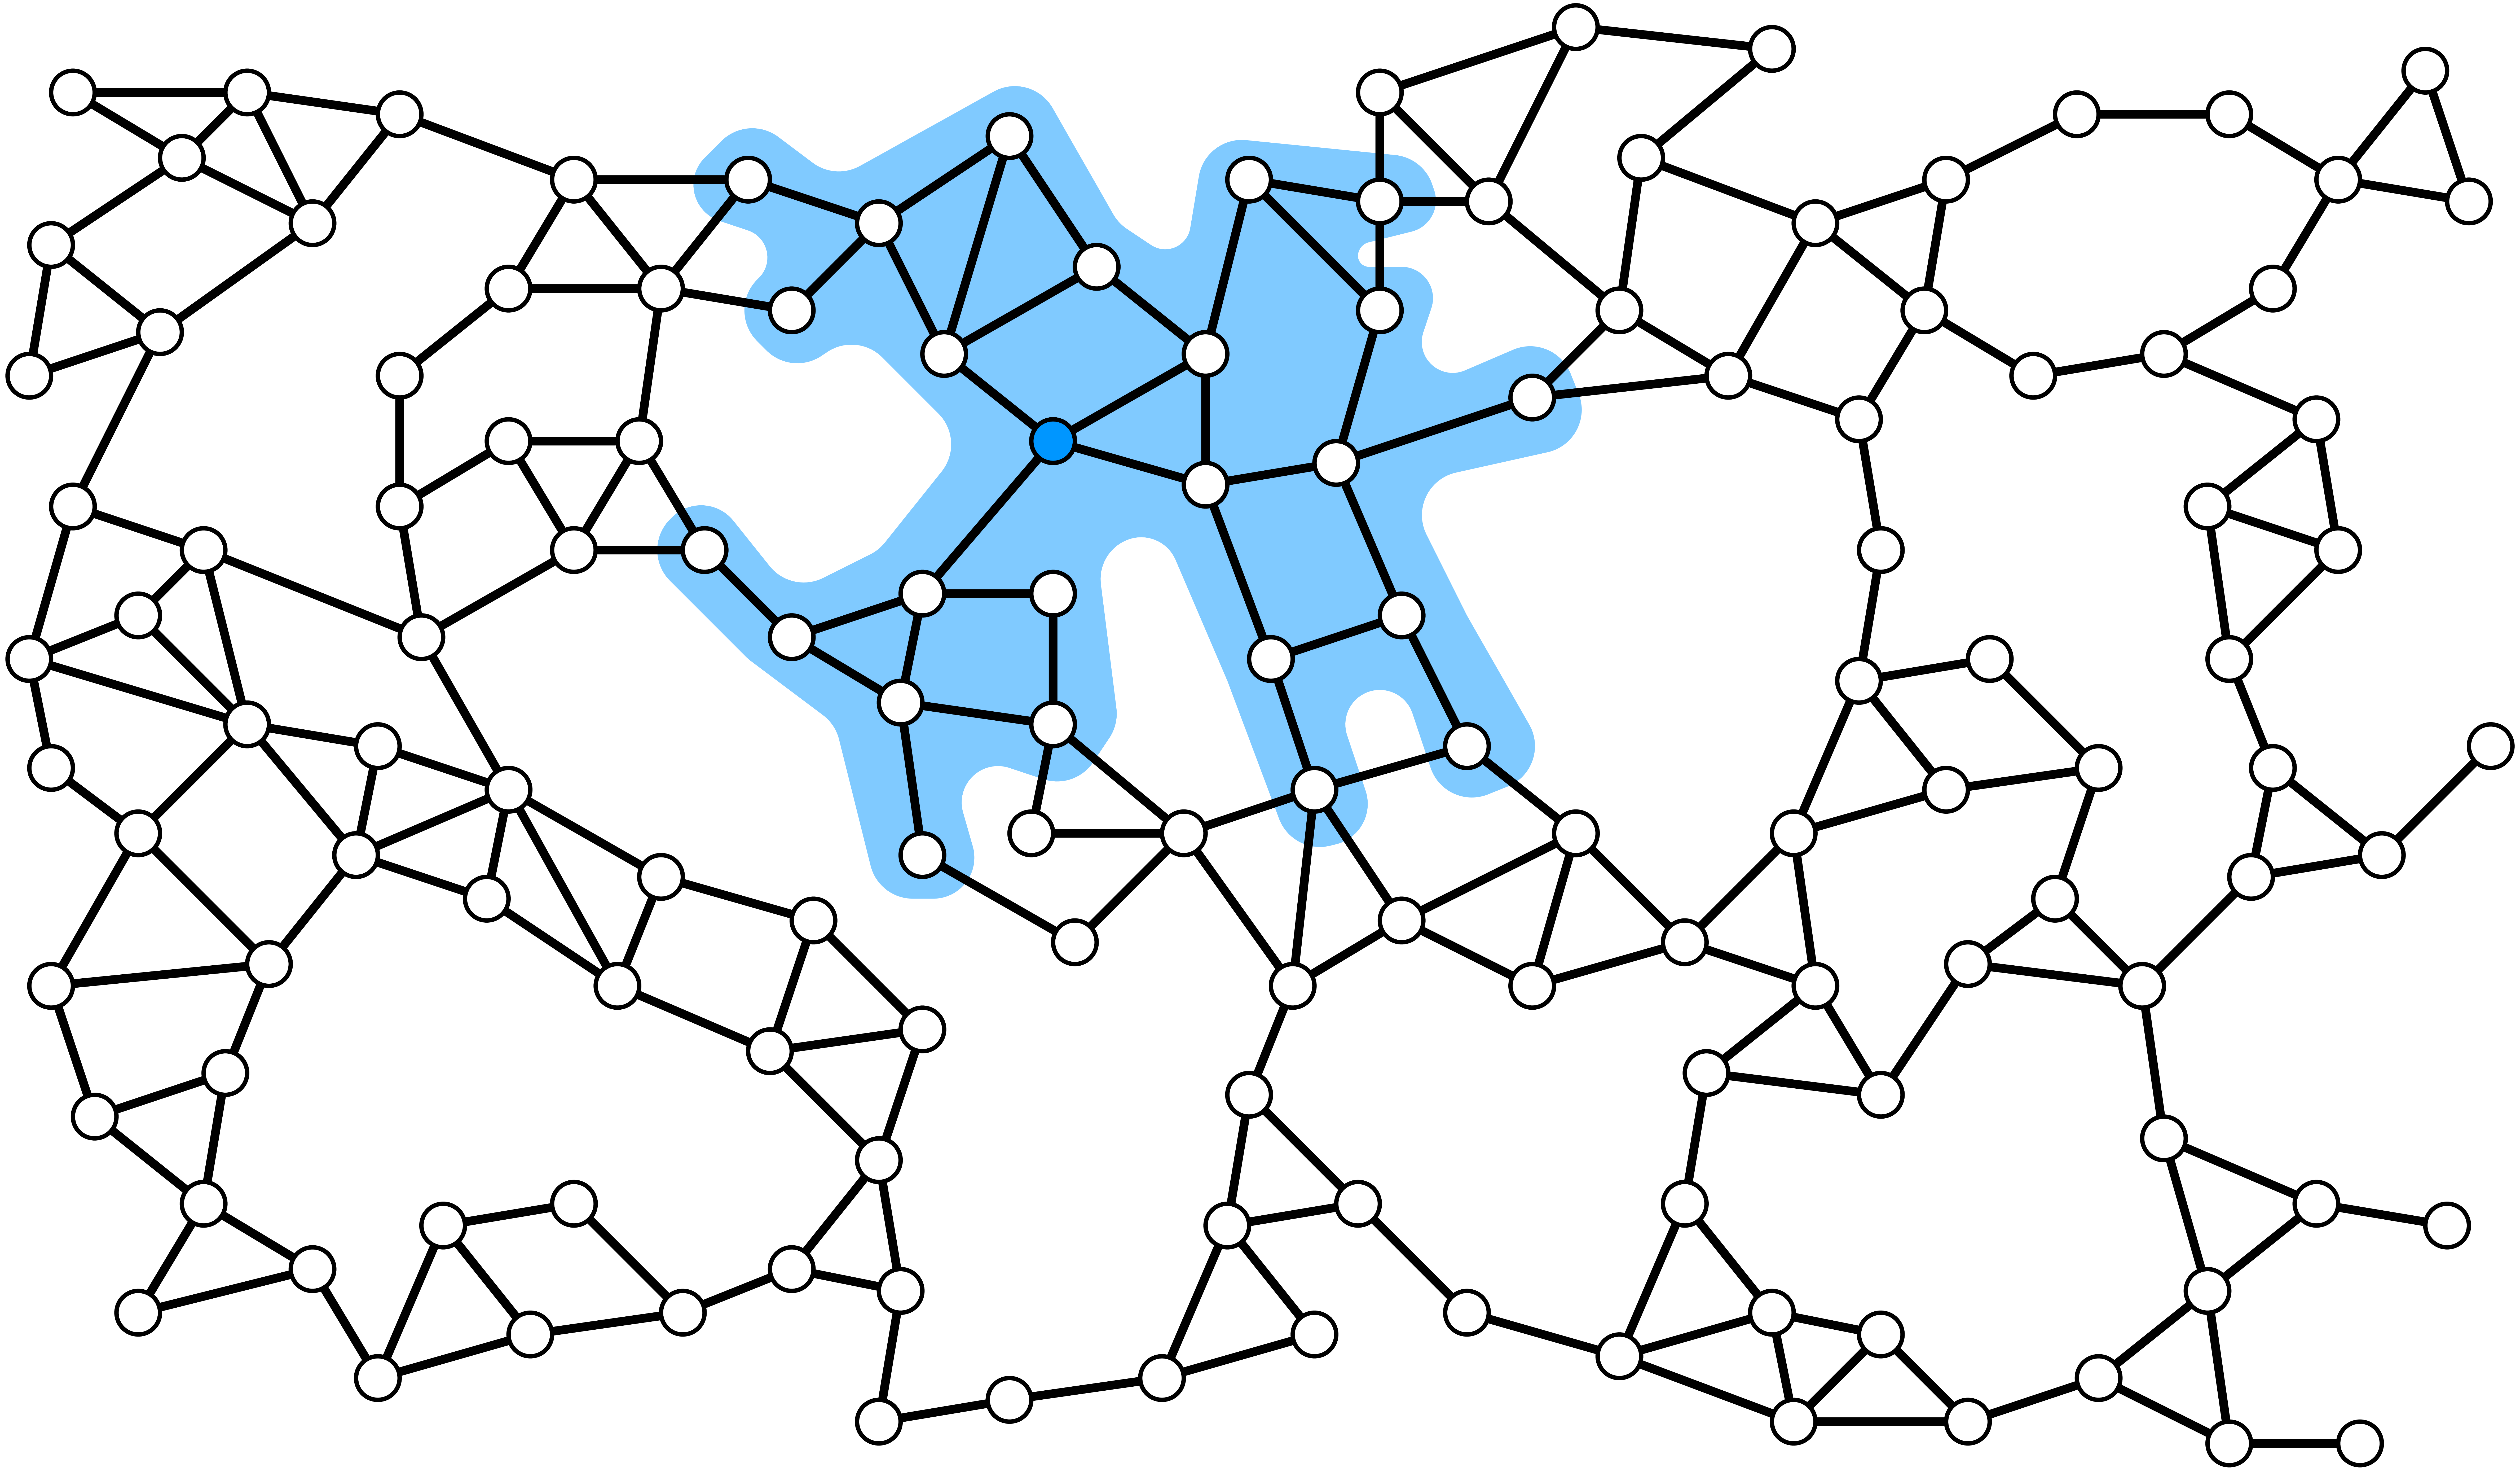

maximal matching

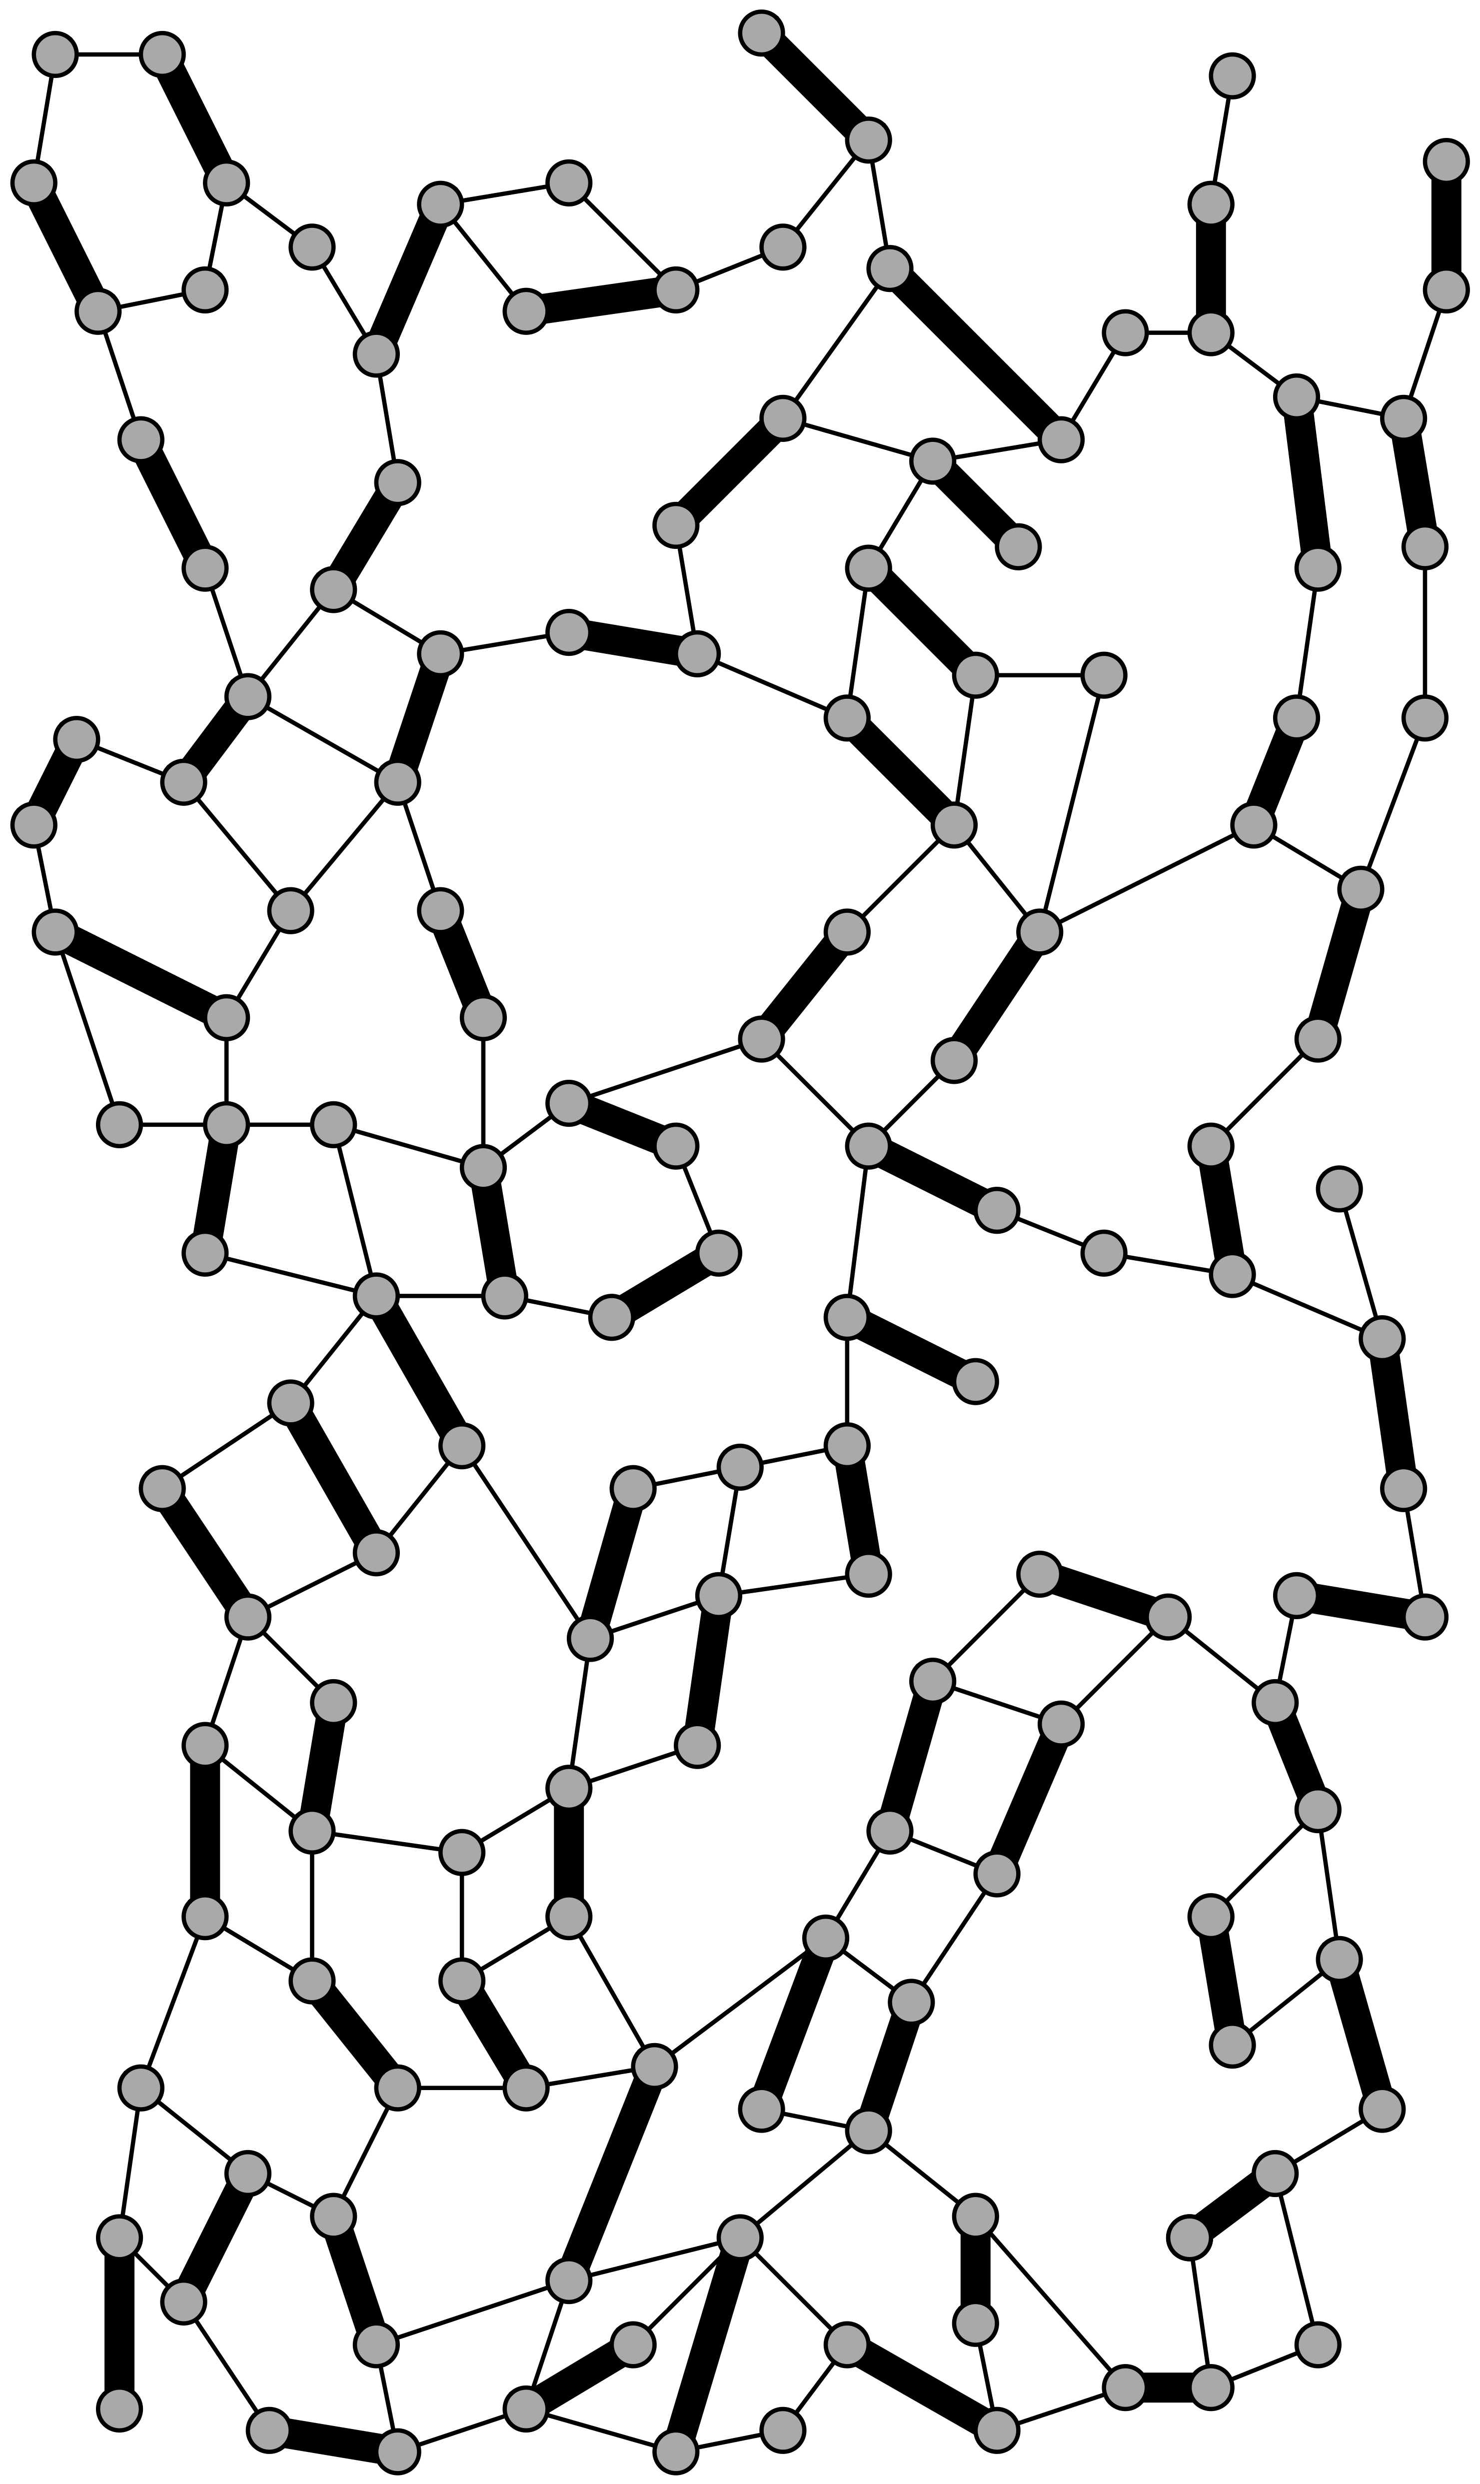

sinkless orientation

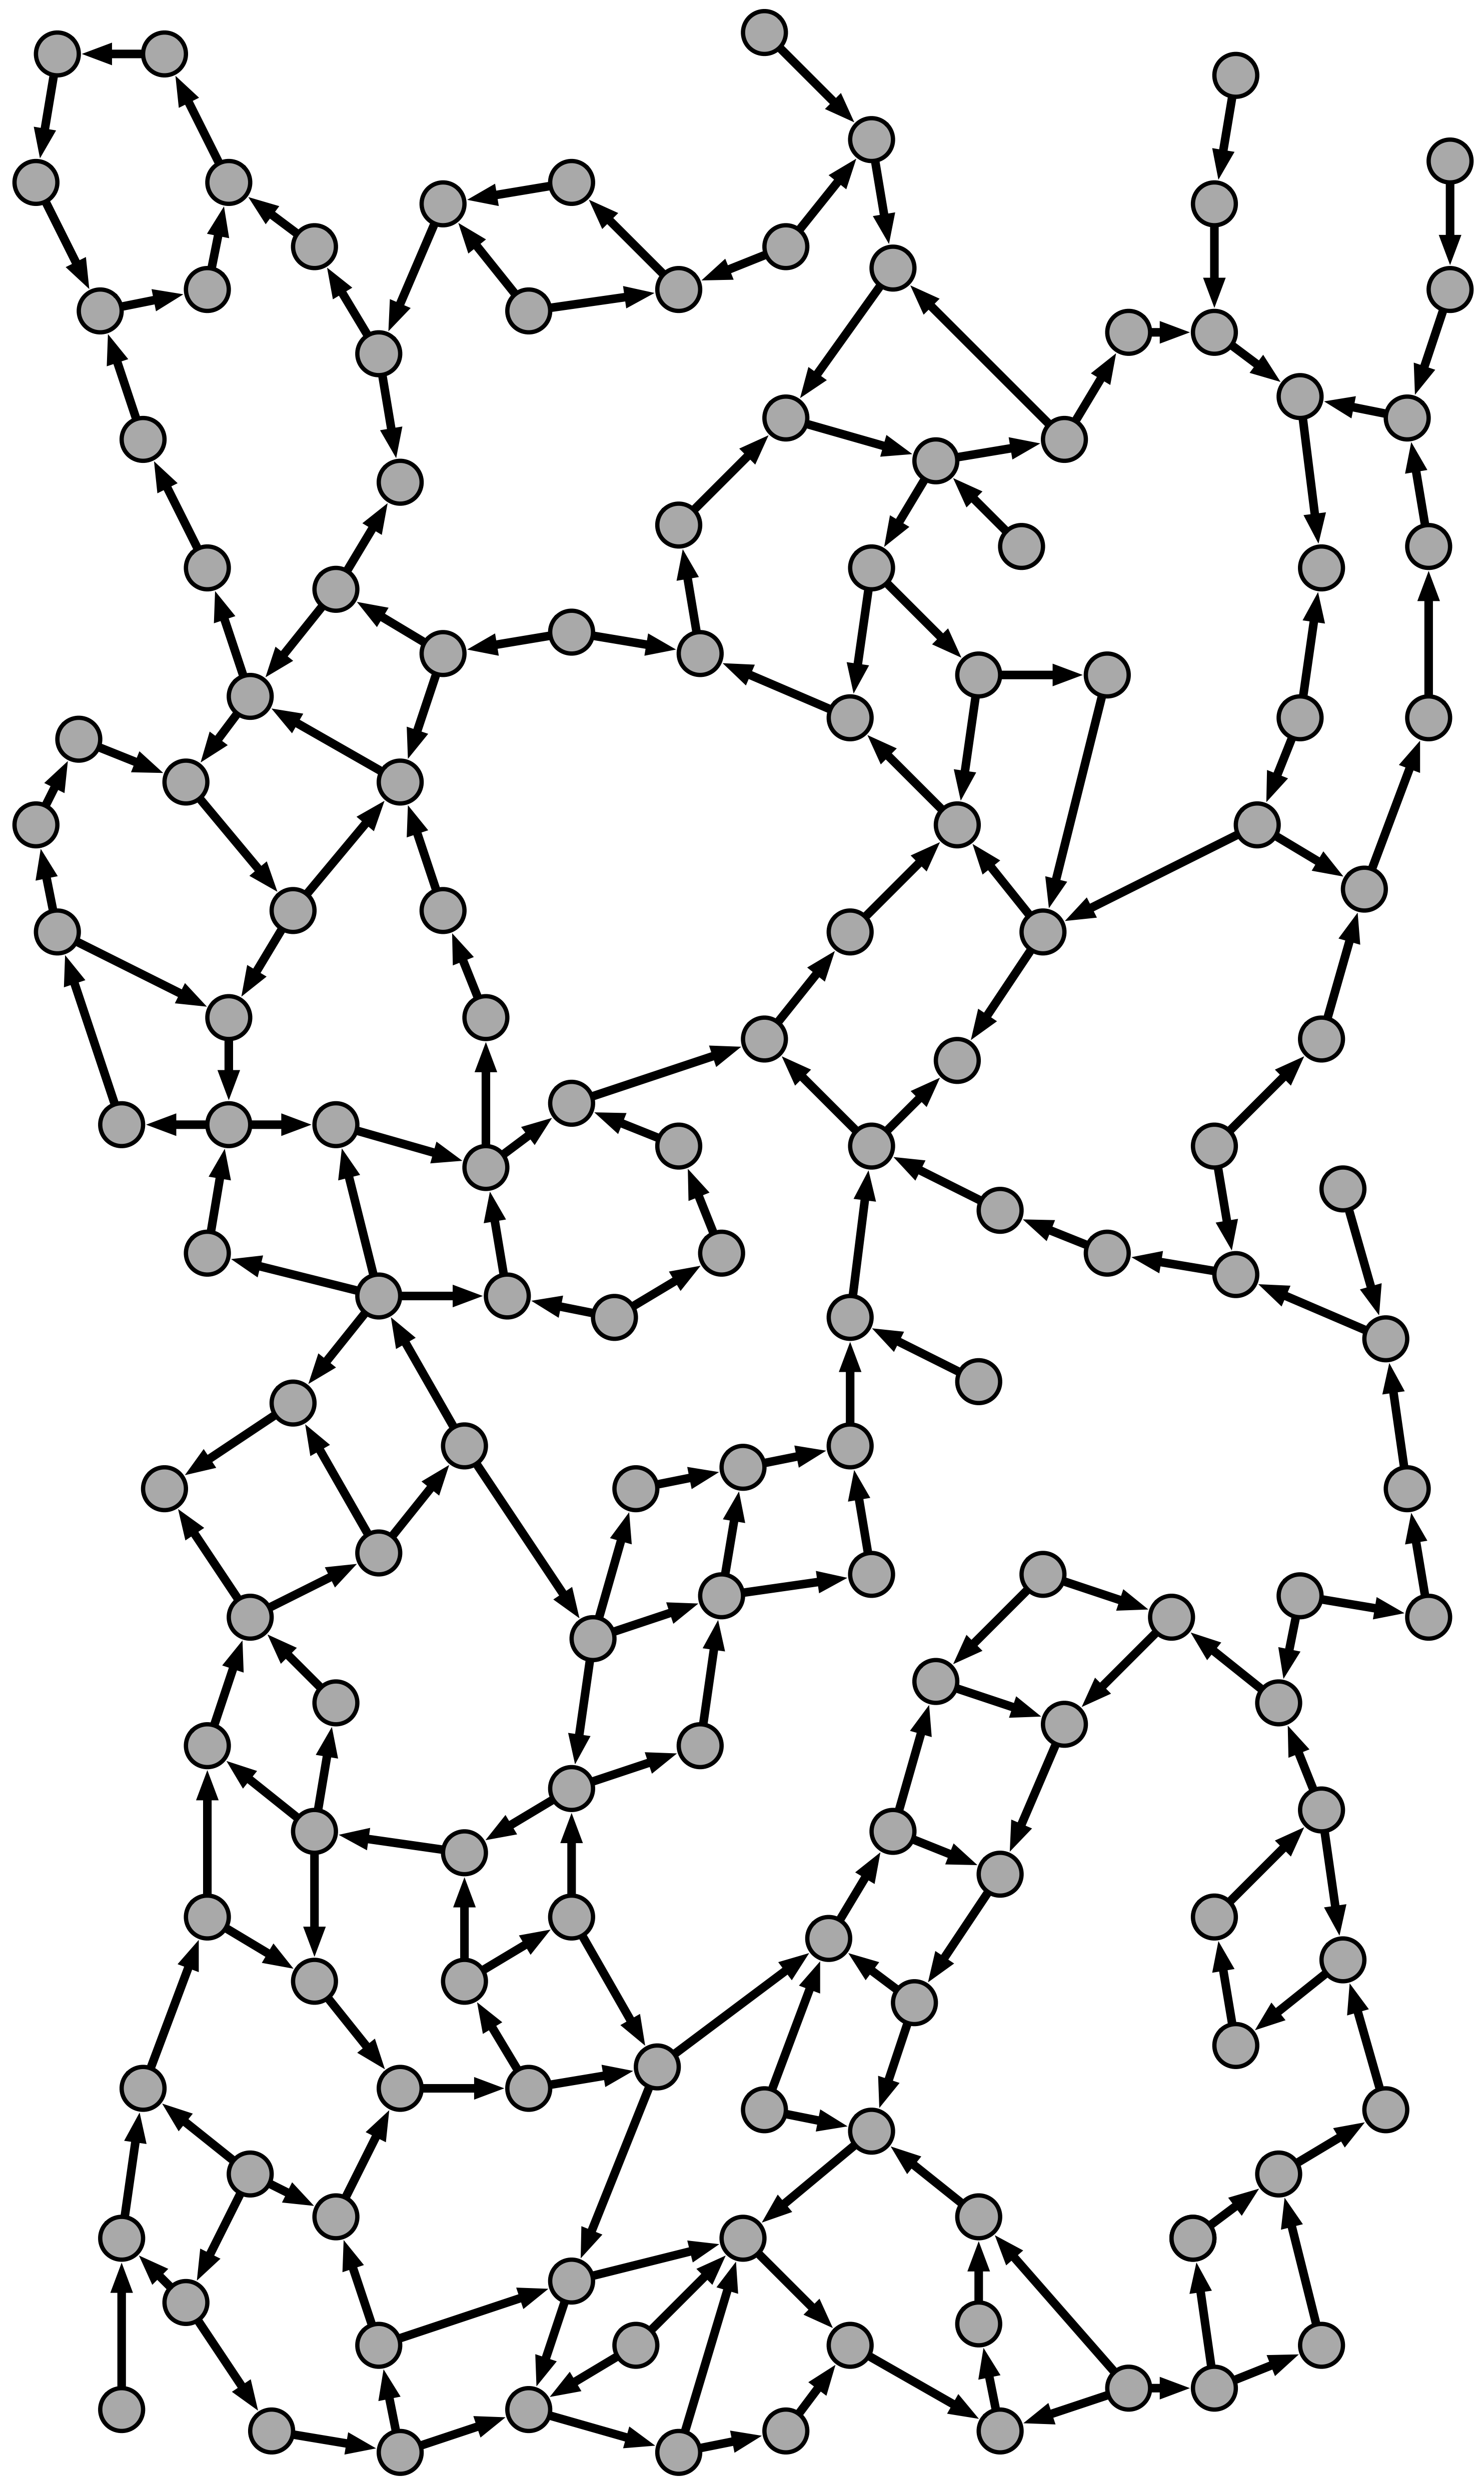

Proposal algorithm: maximal matching

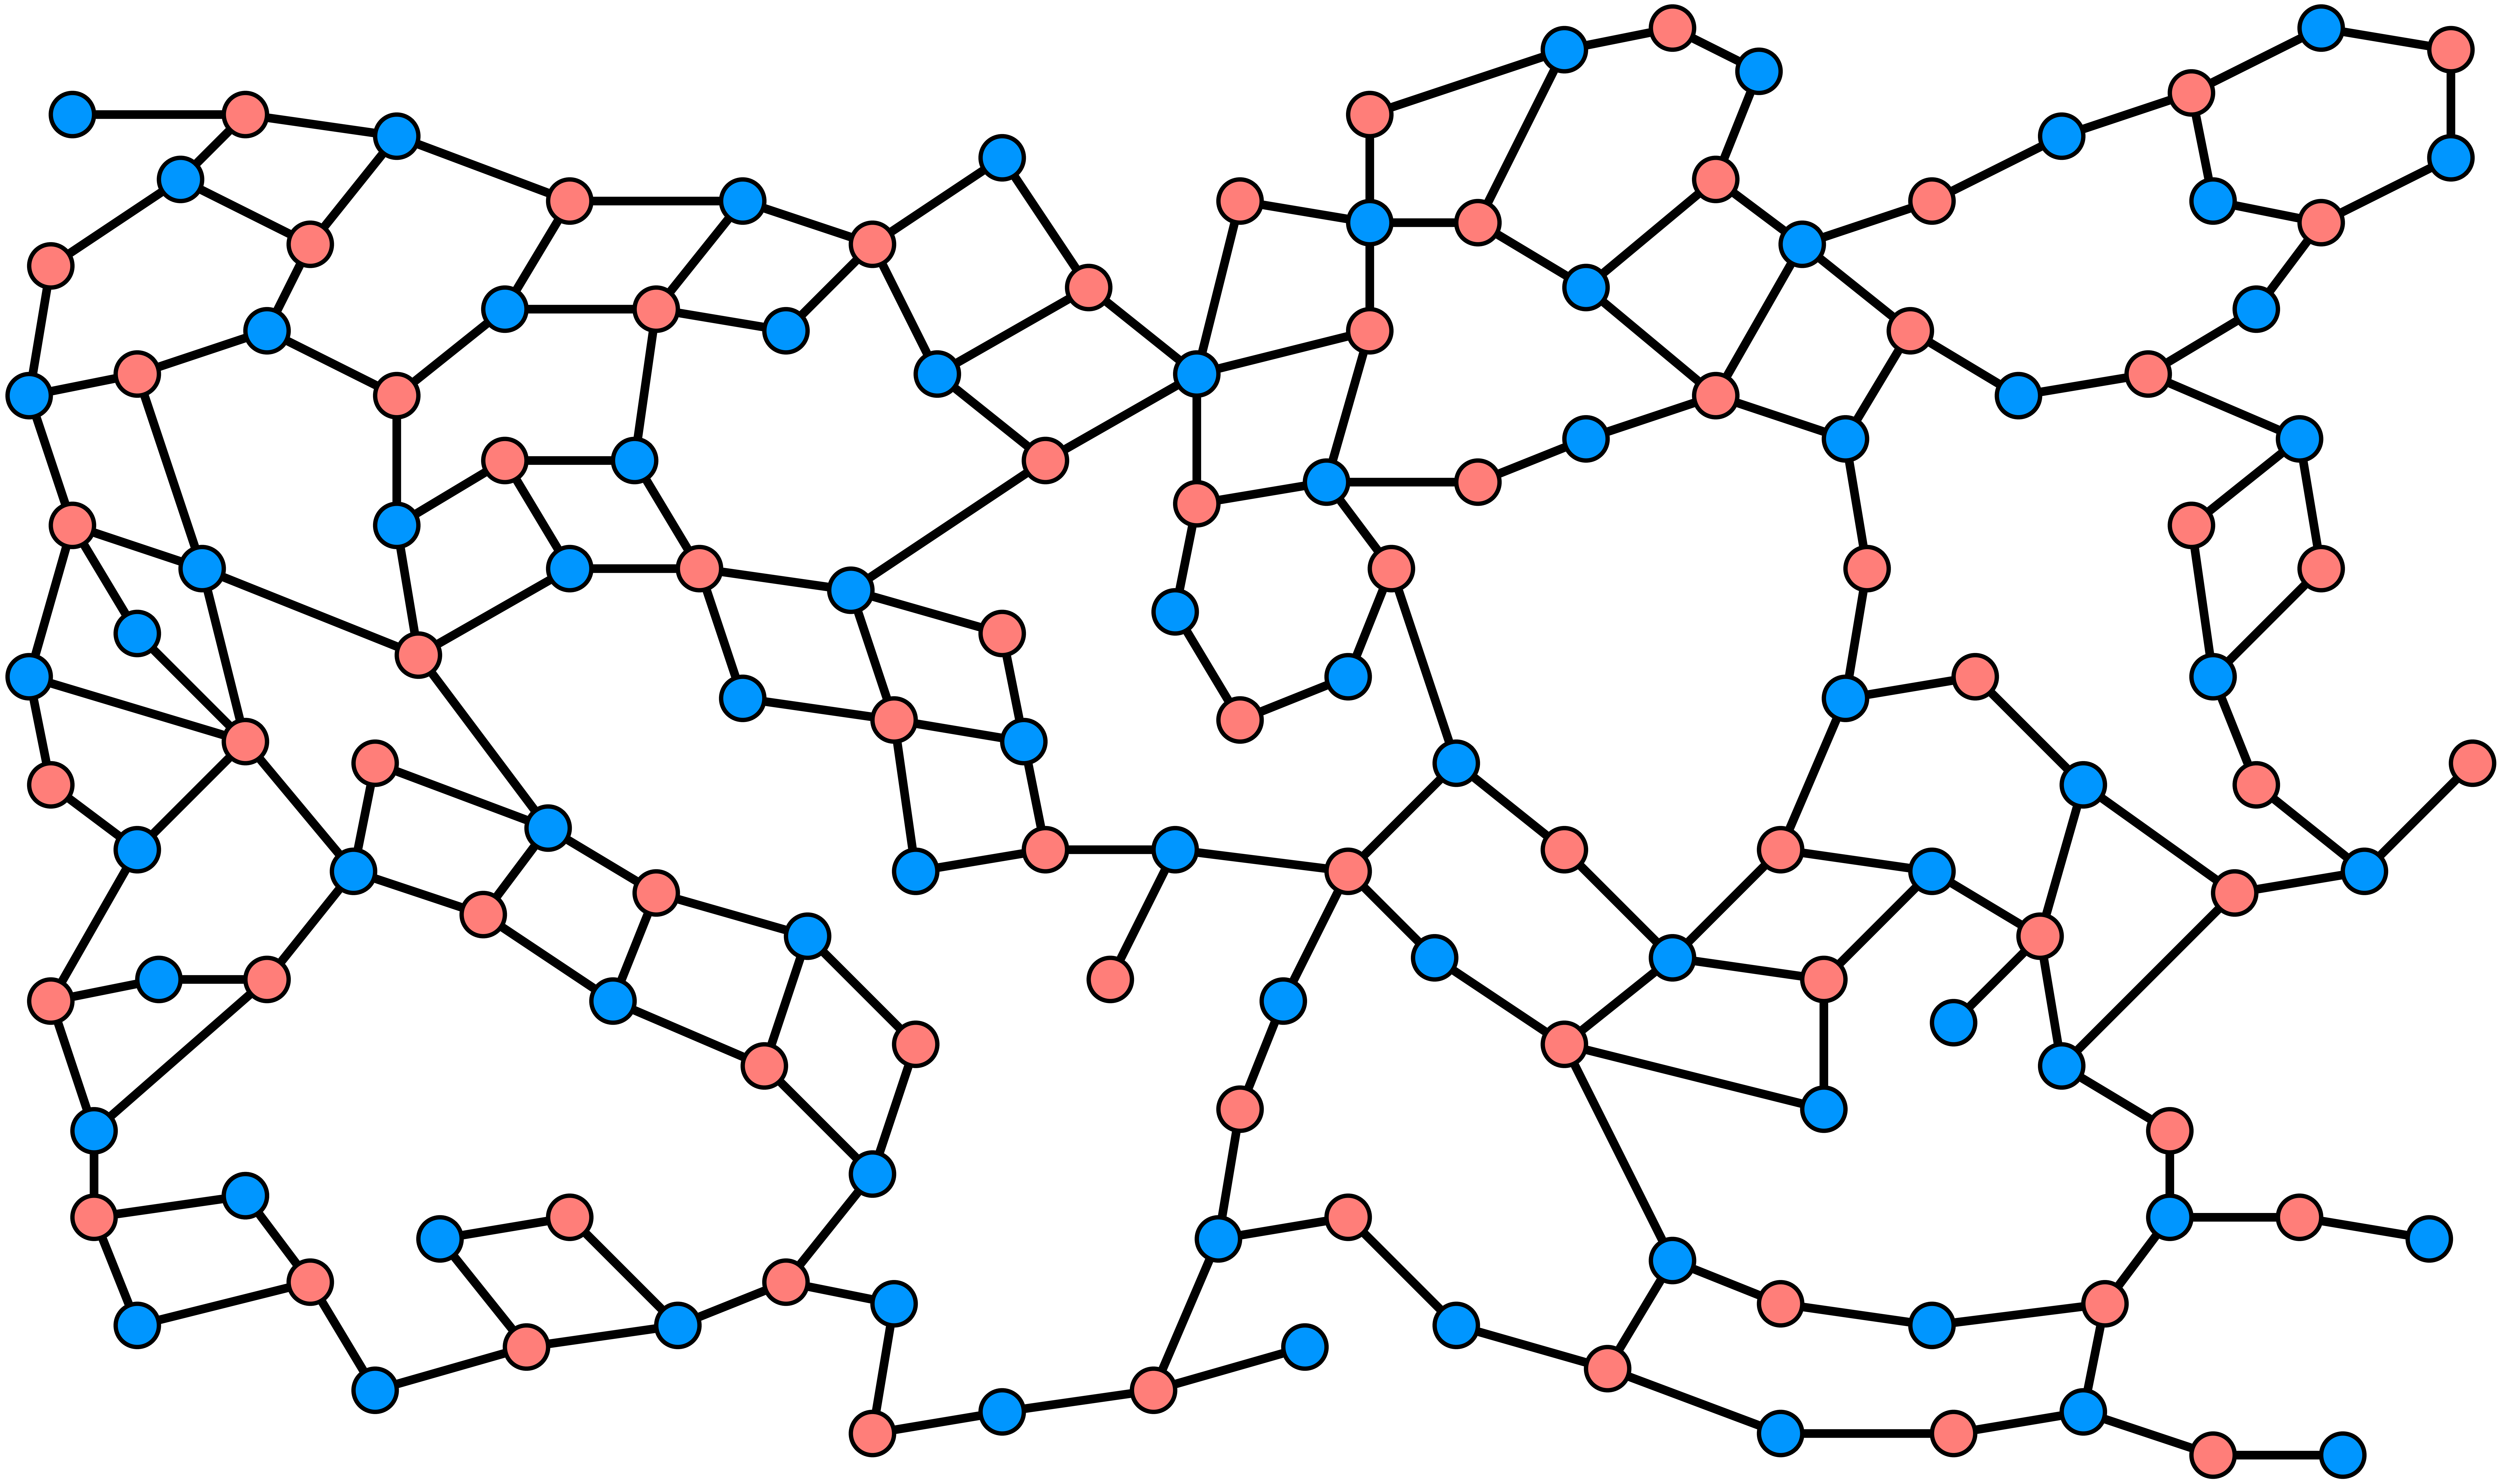

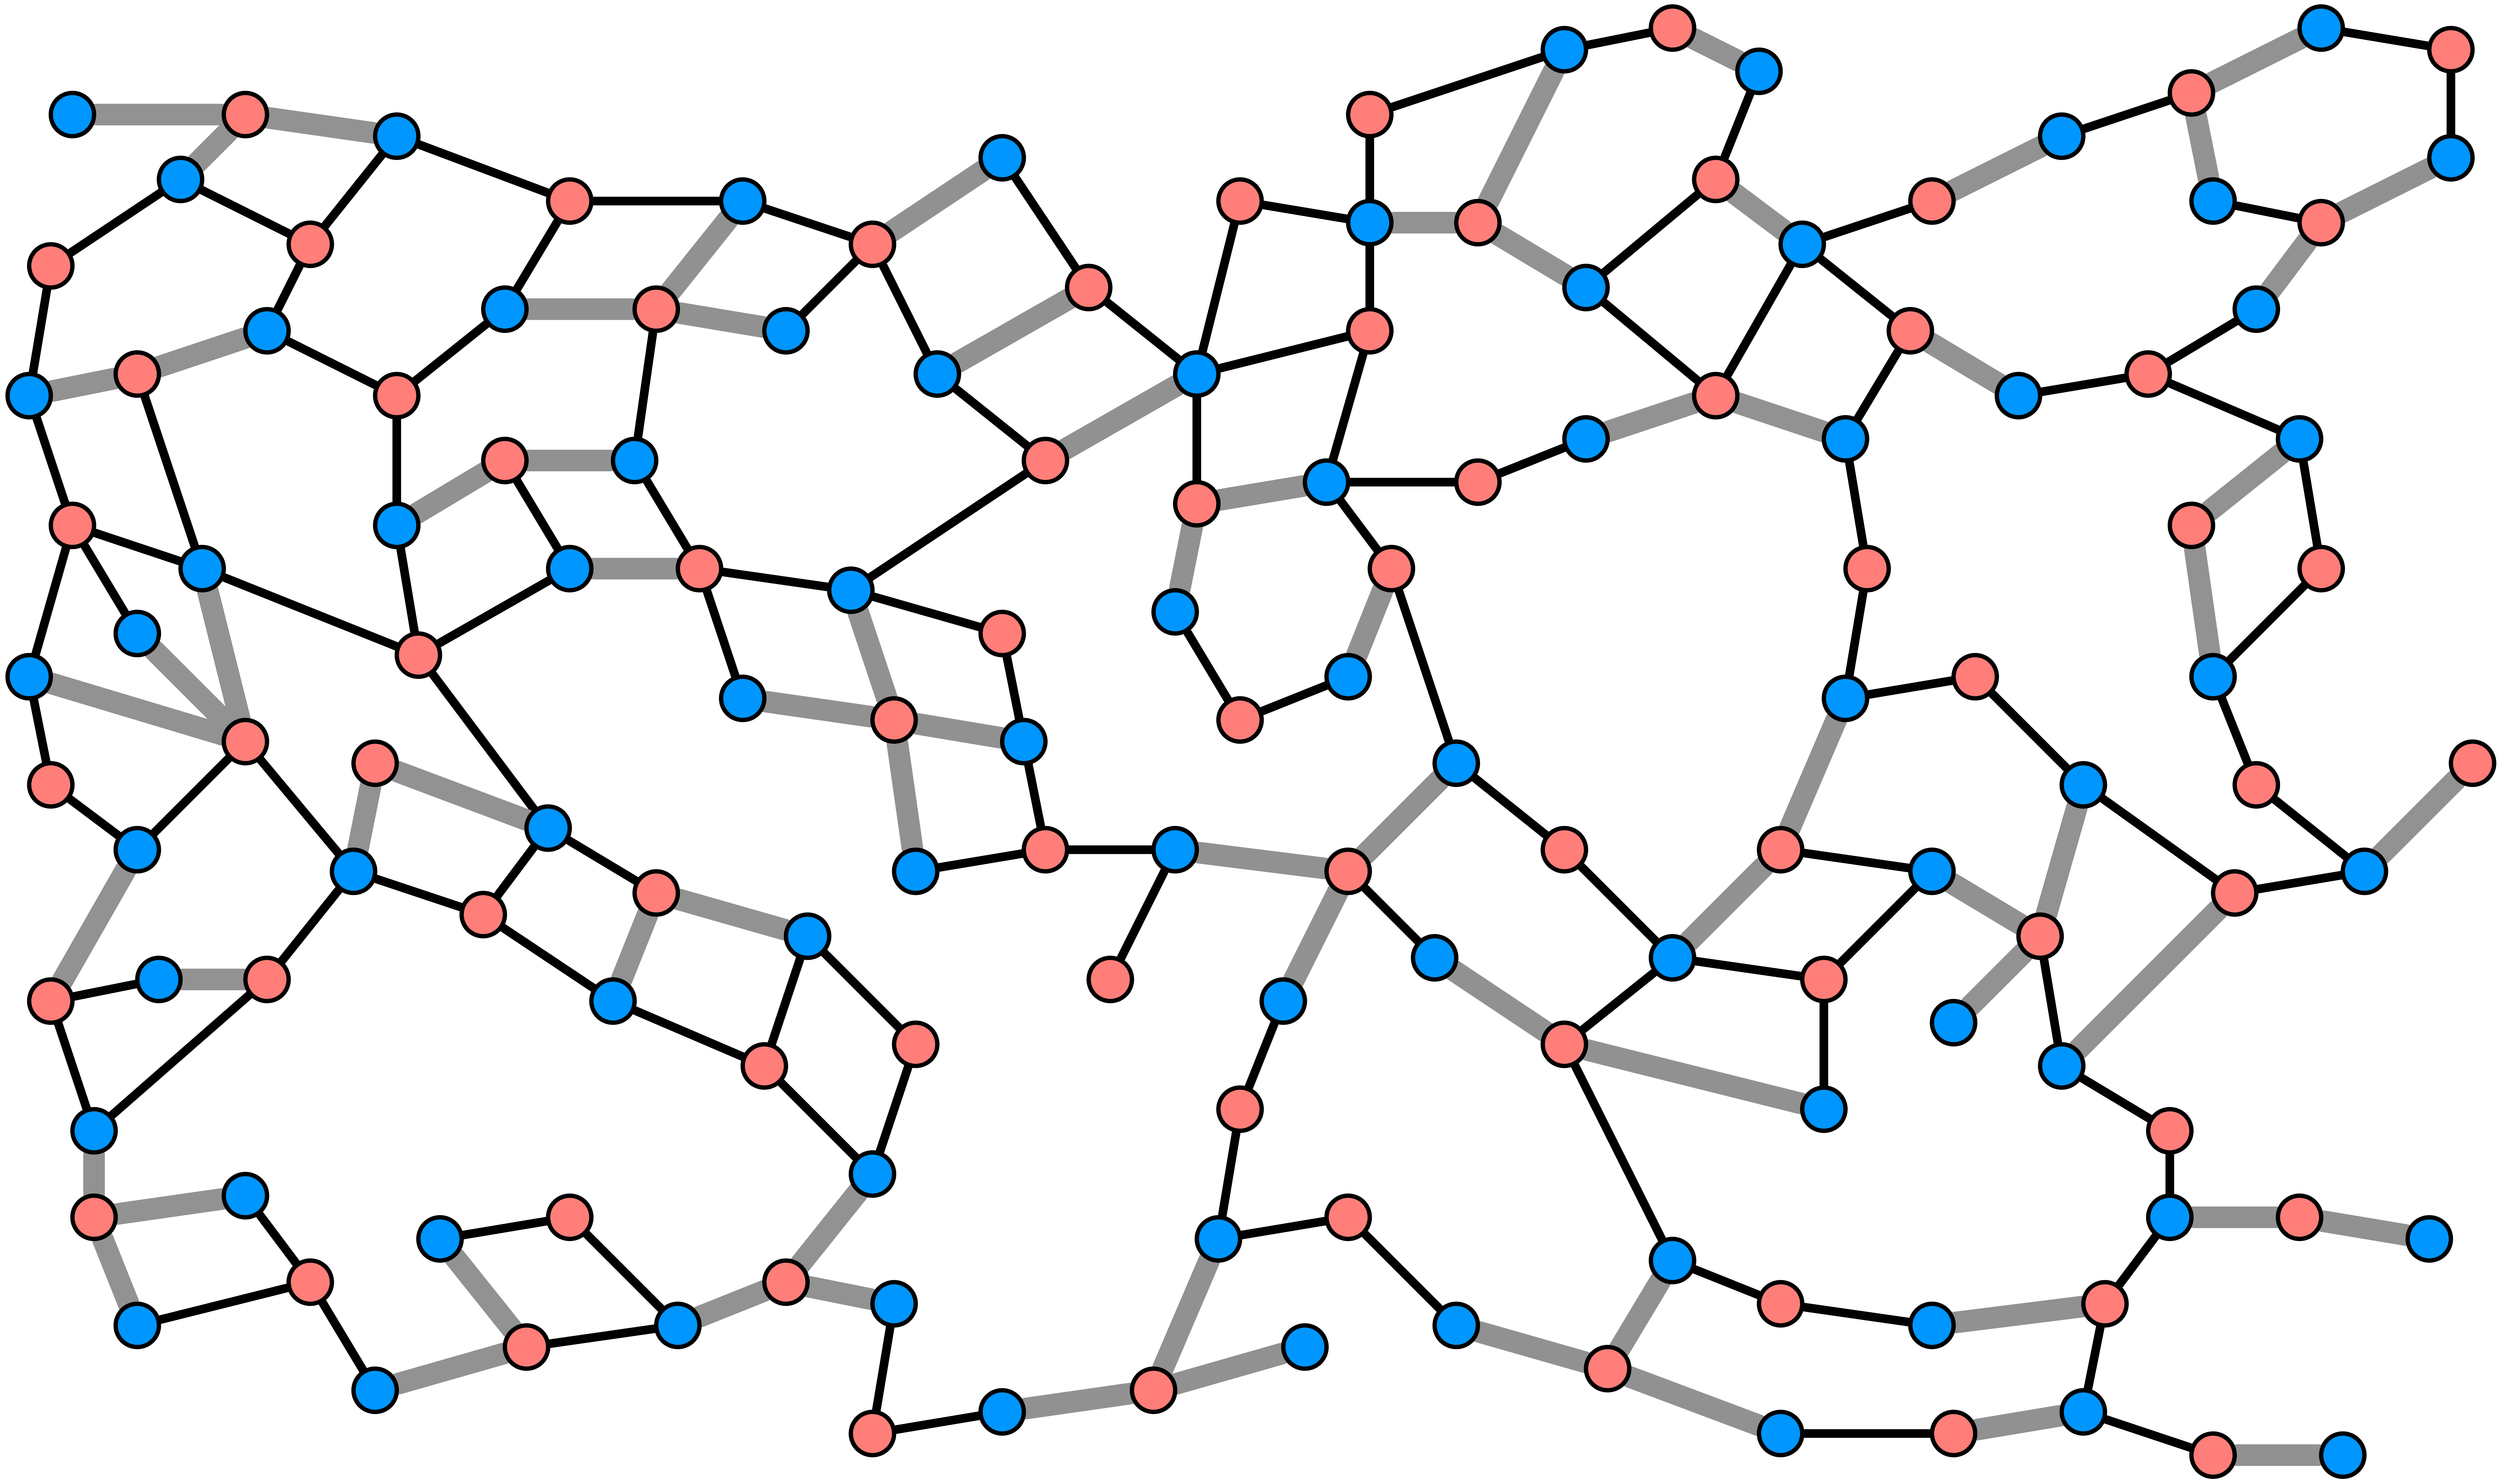

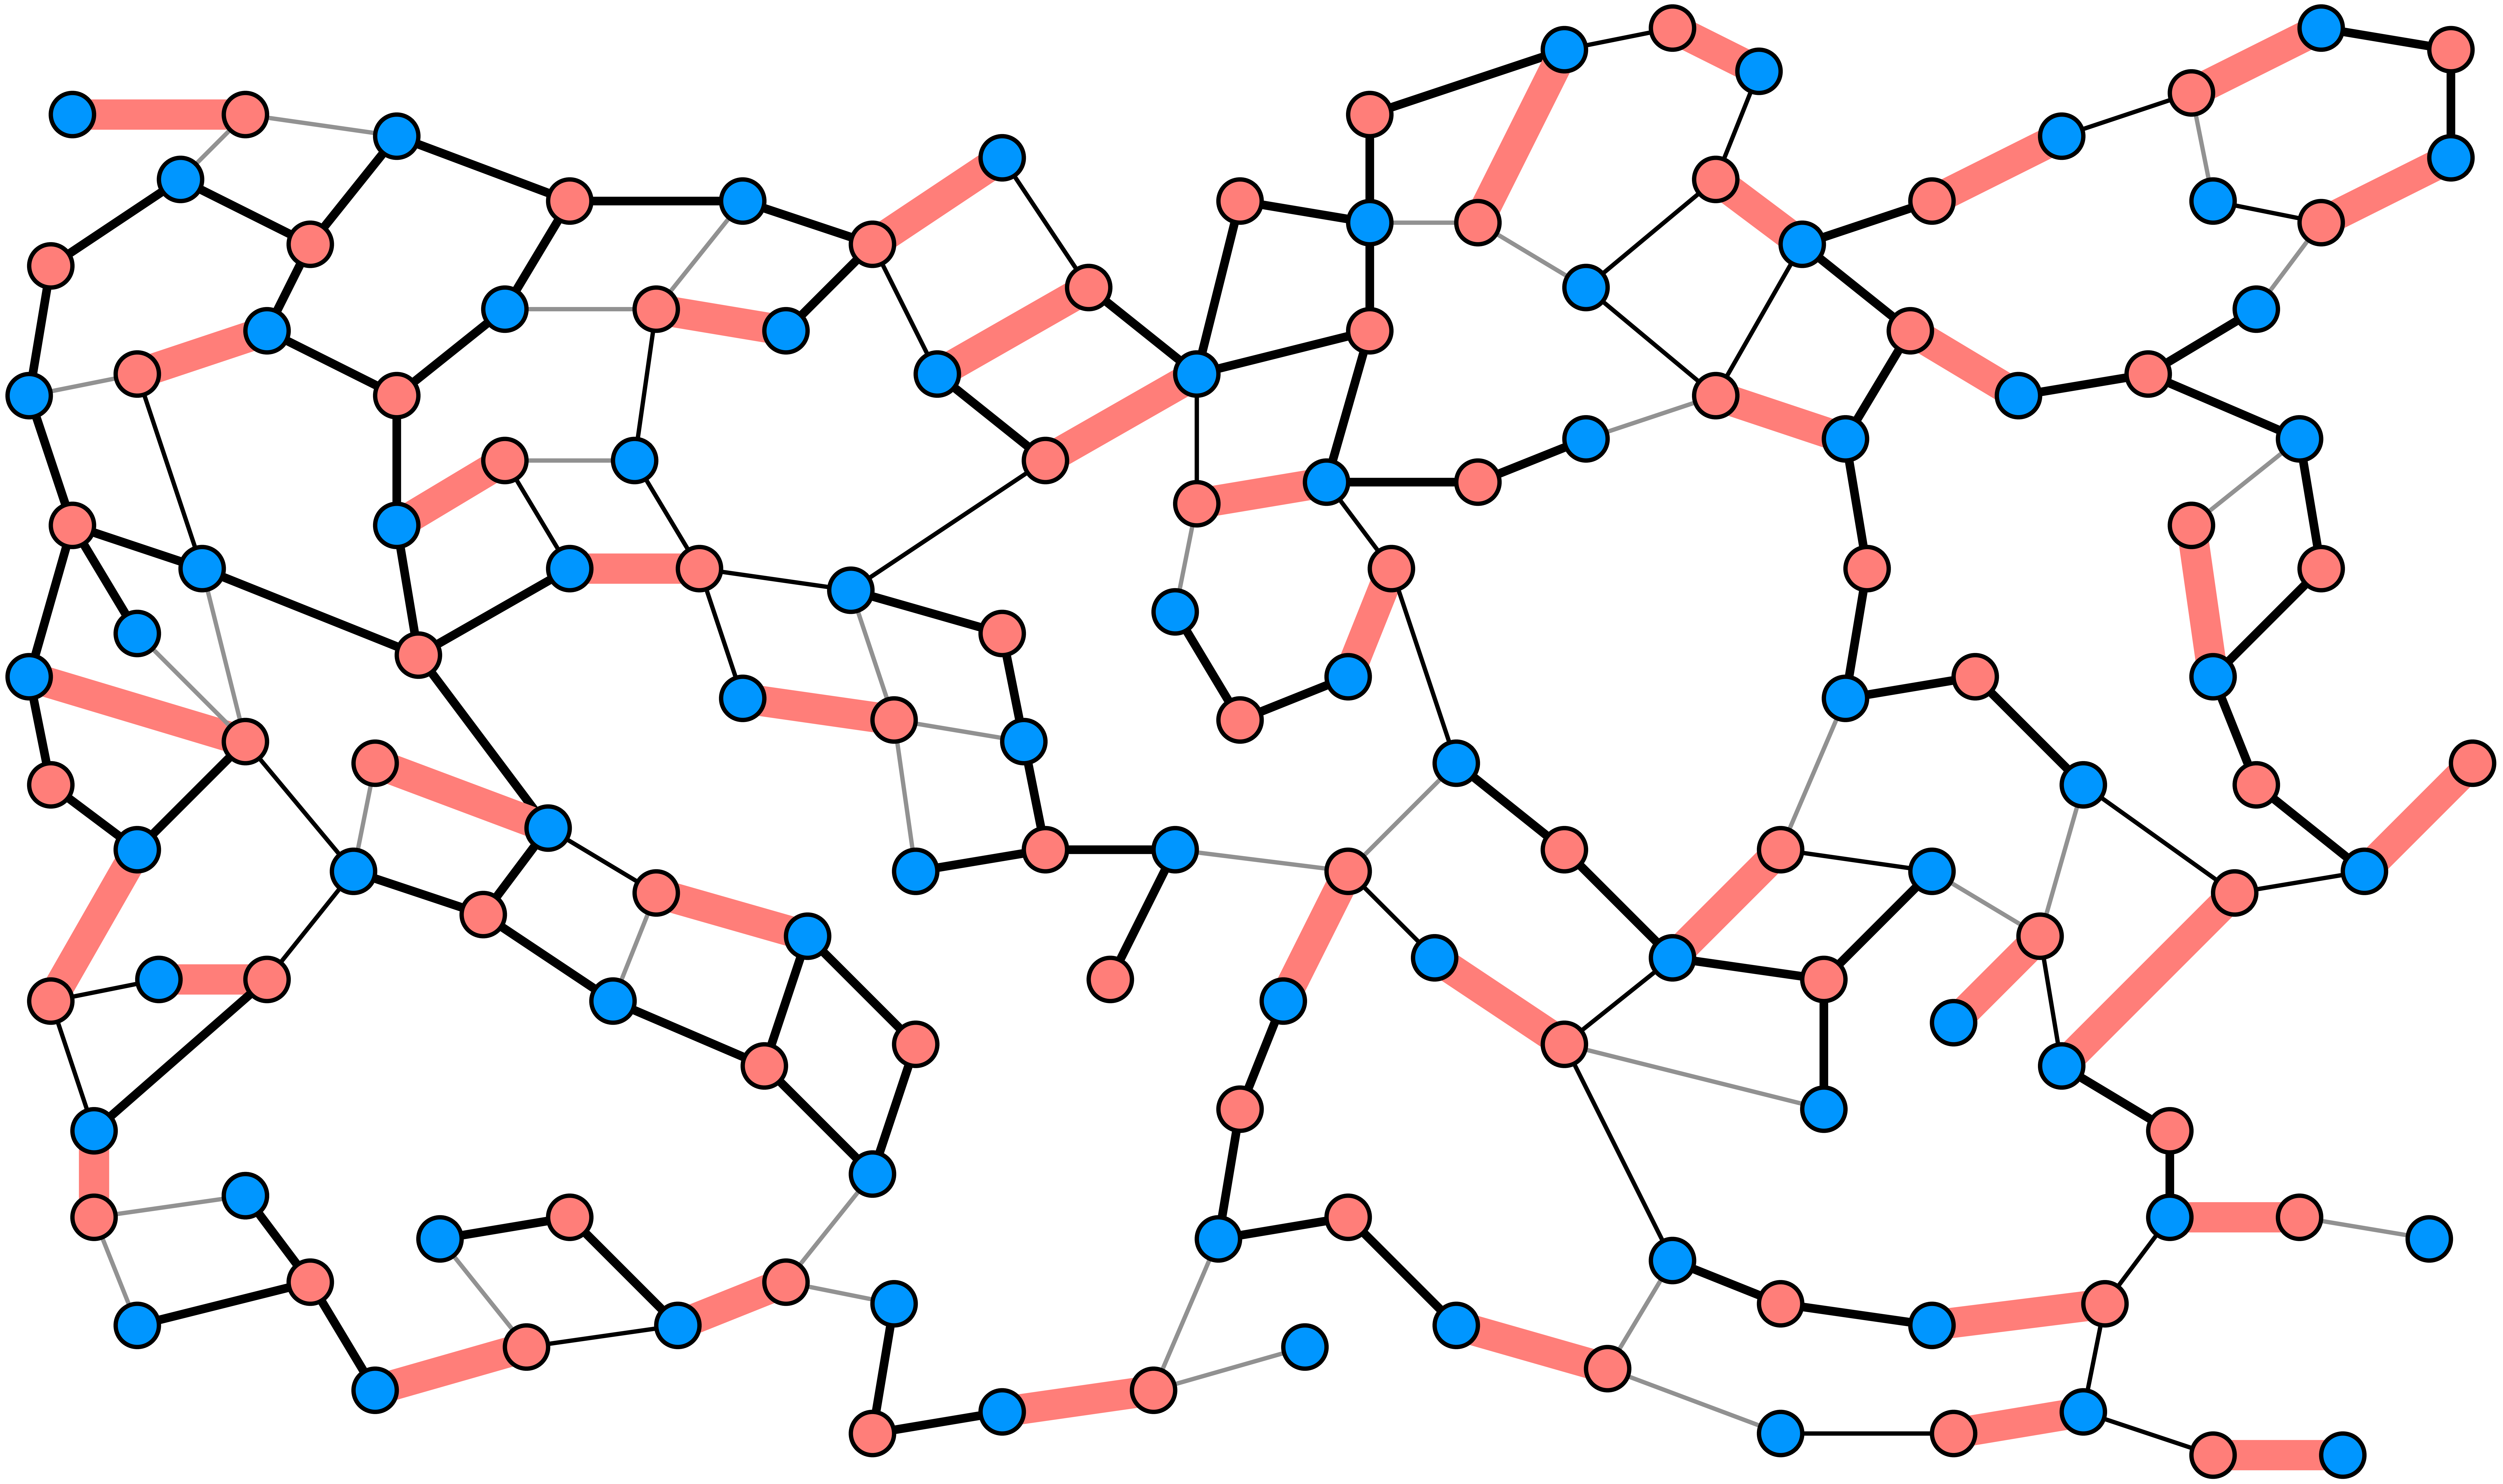

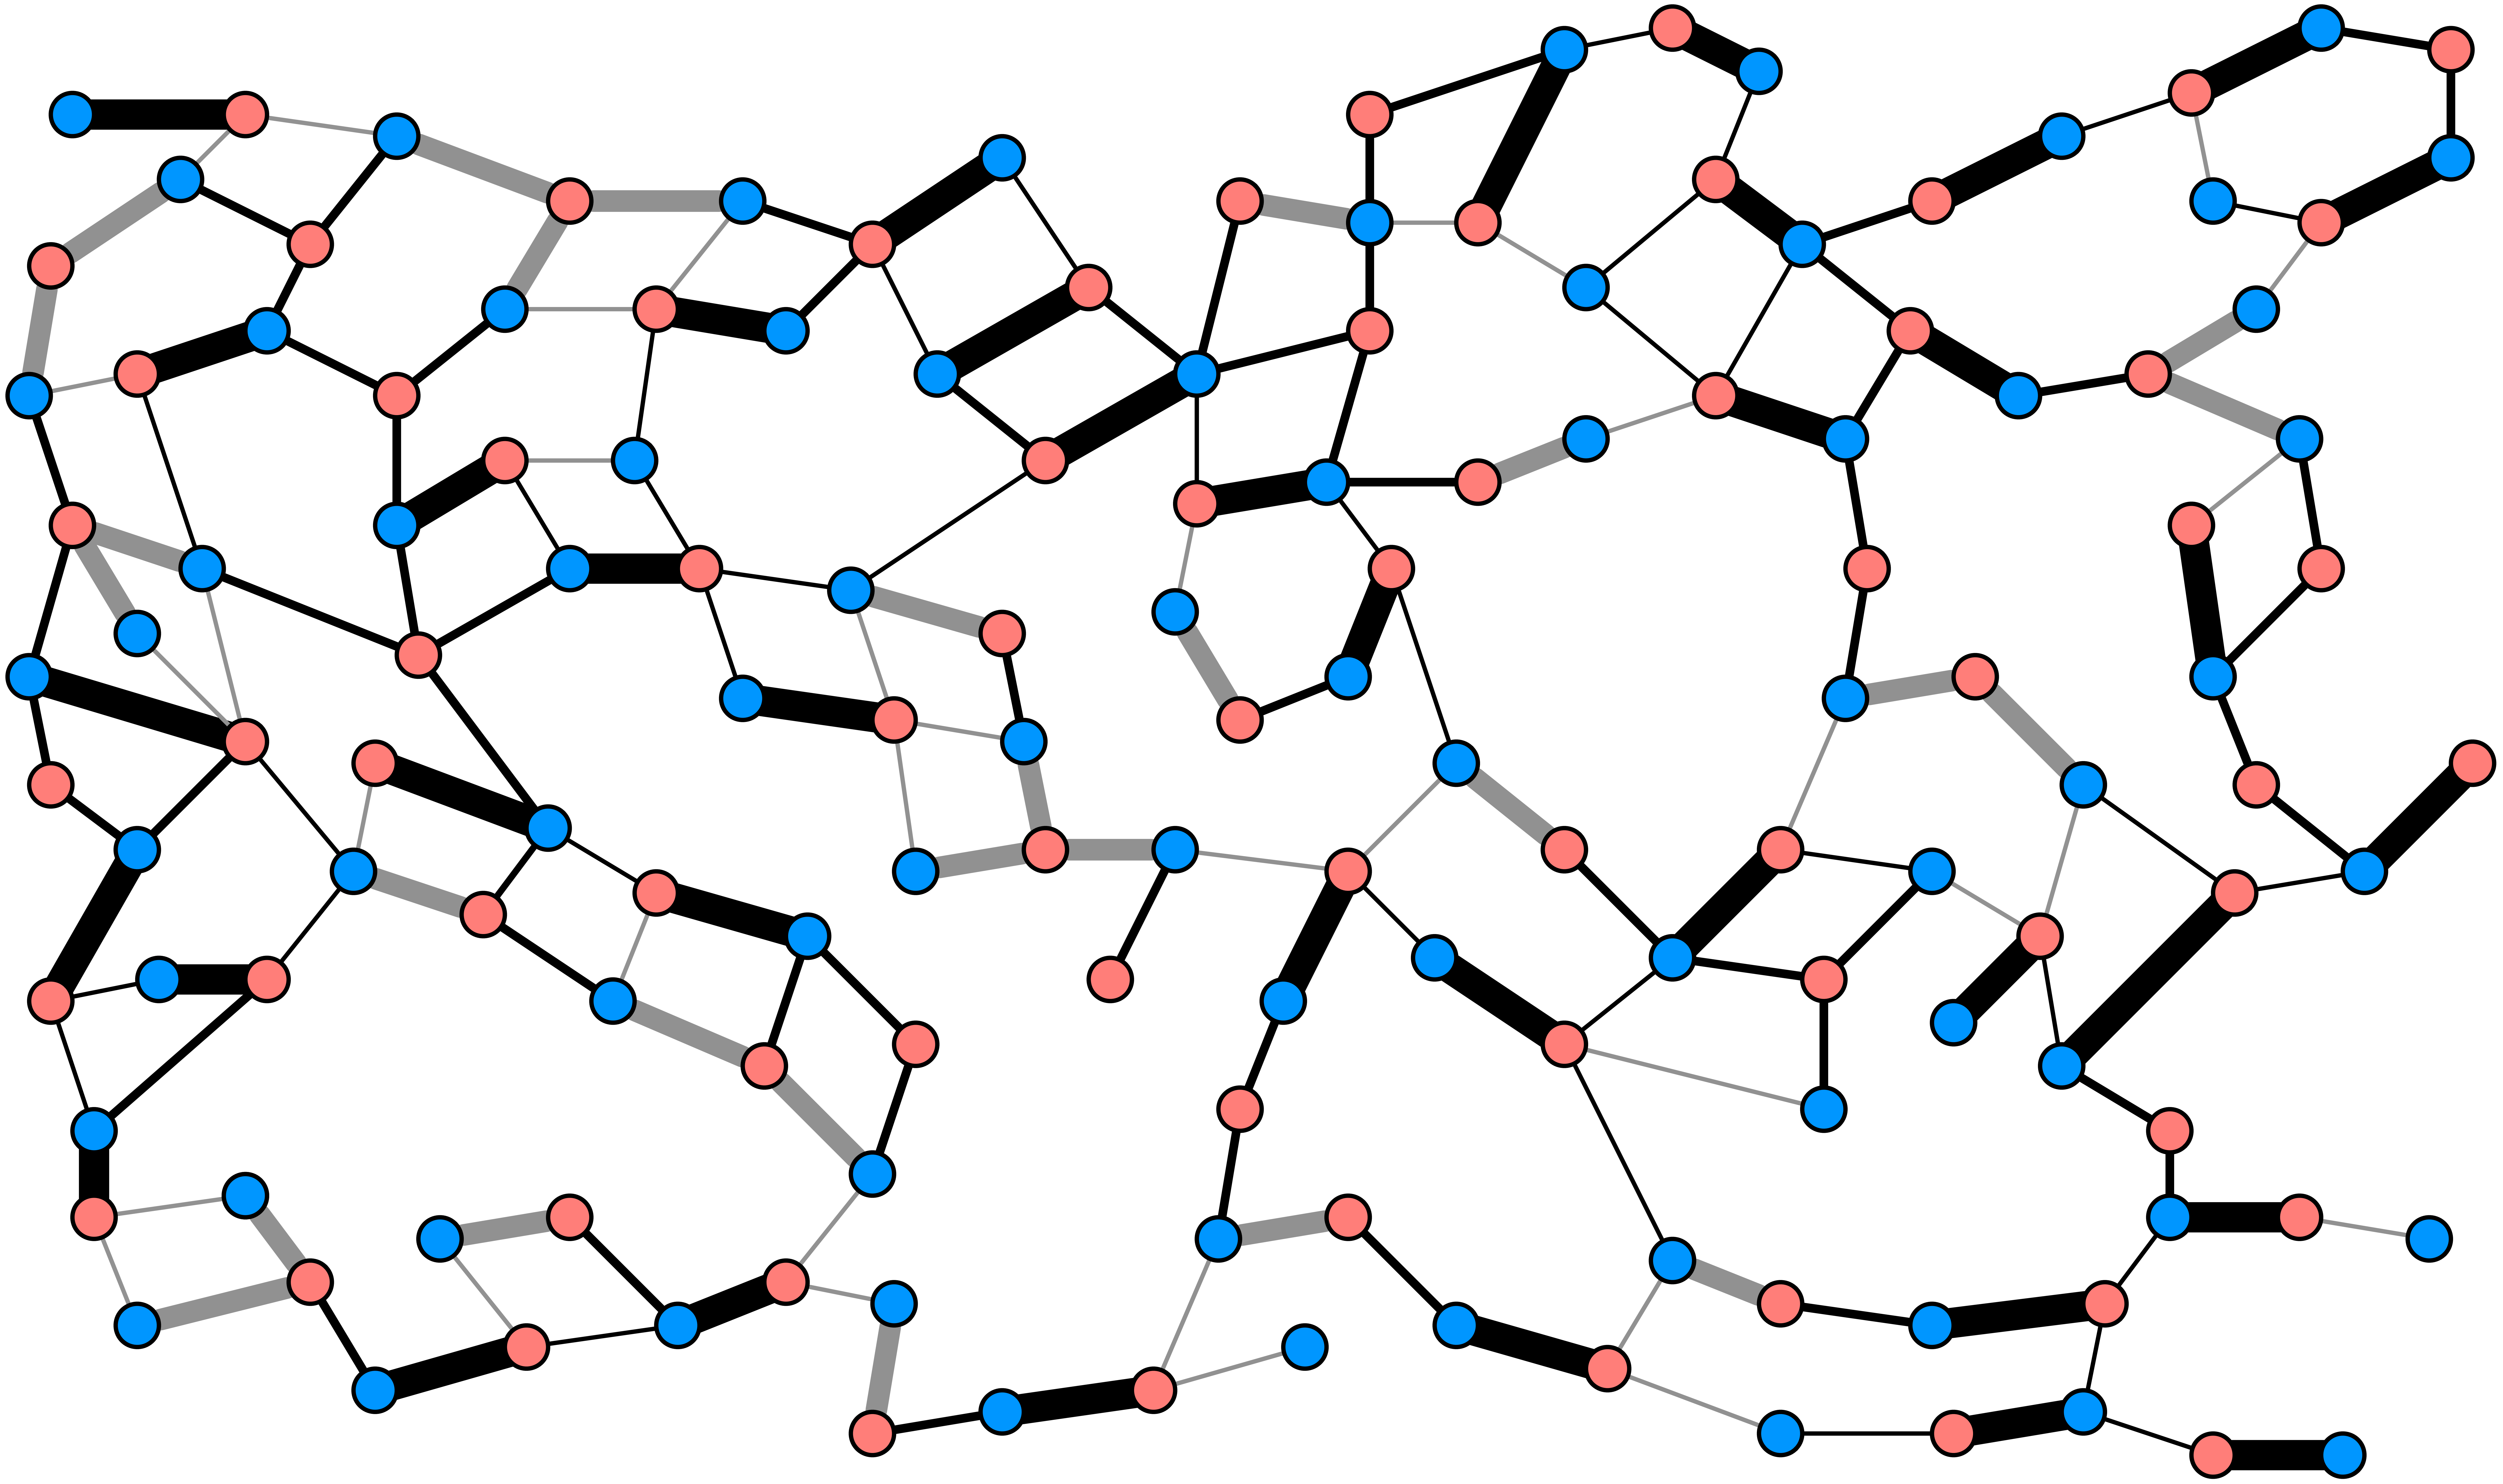

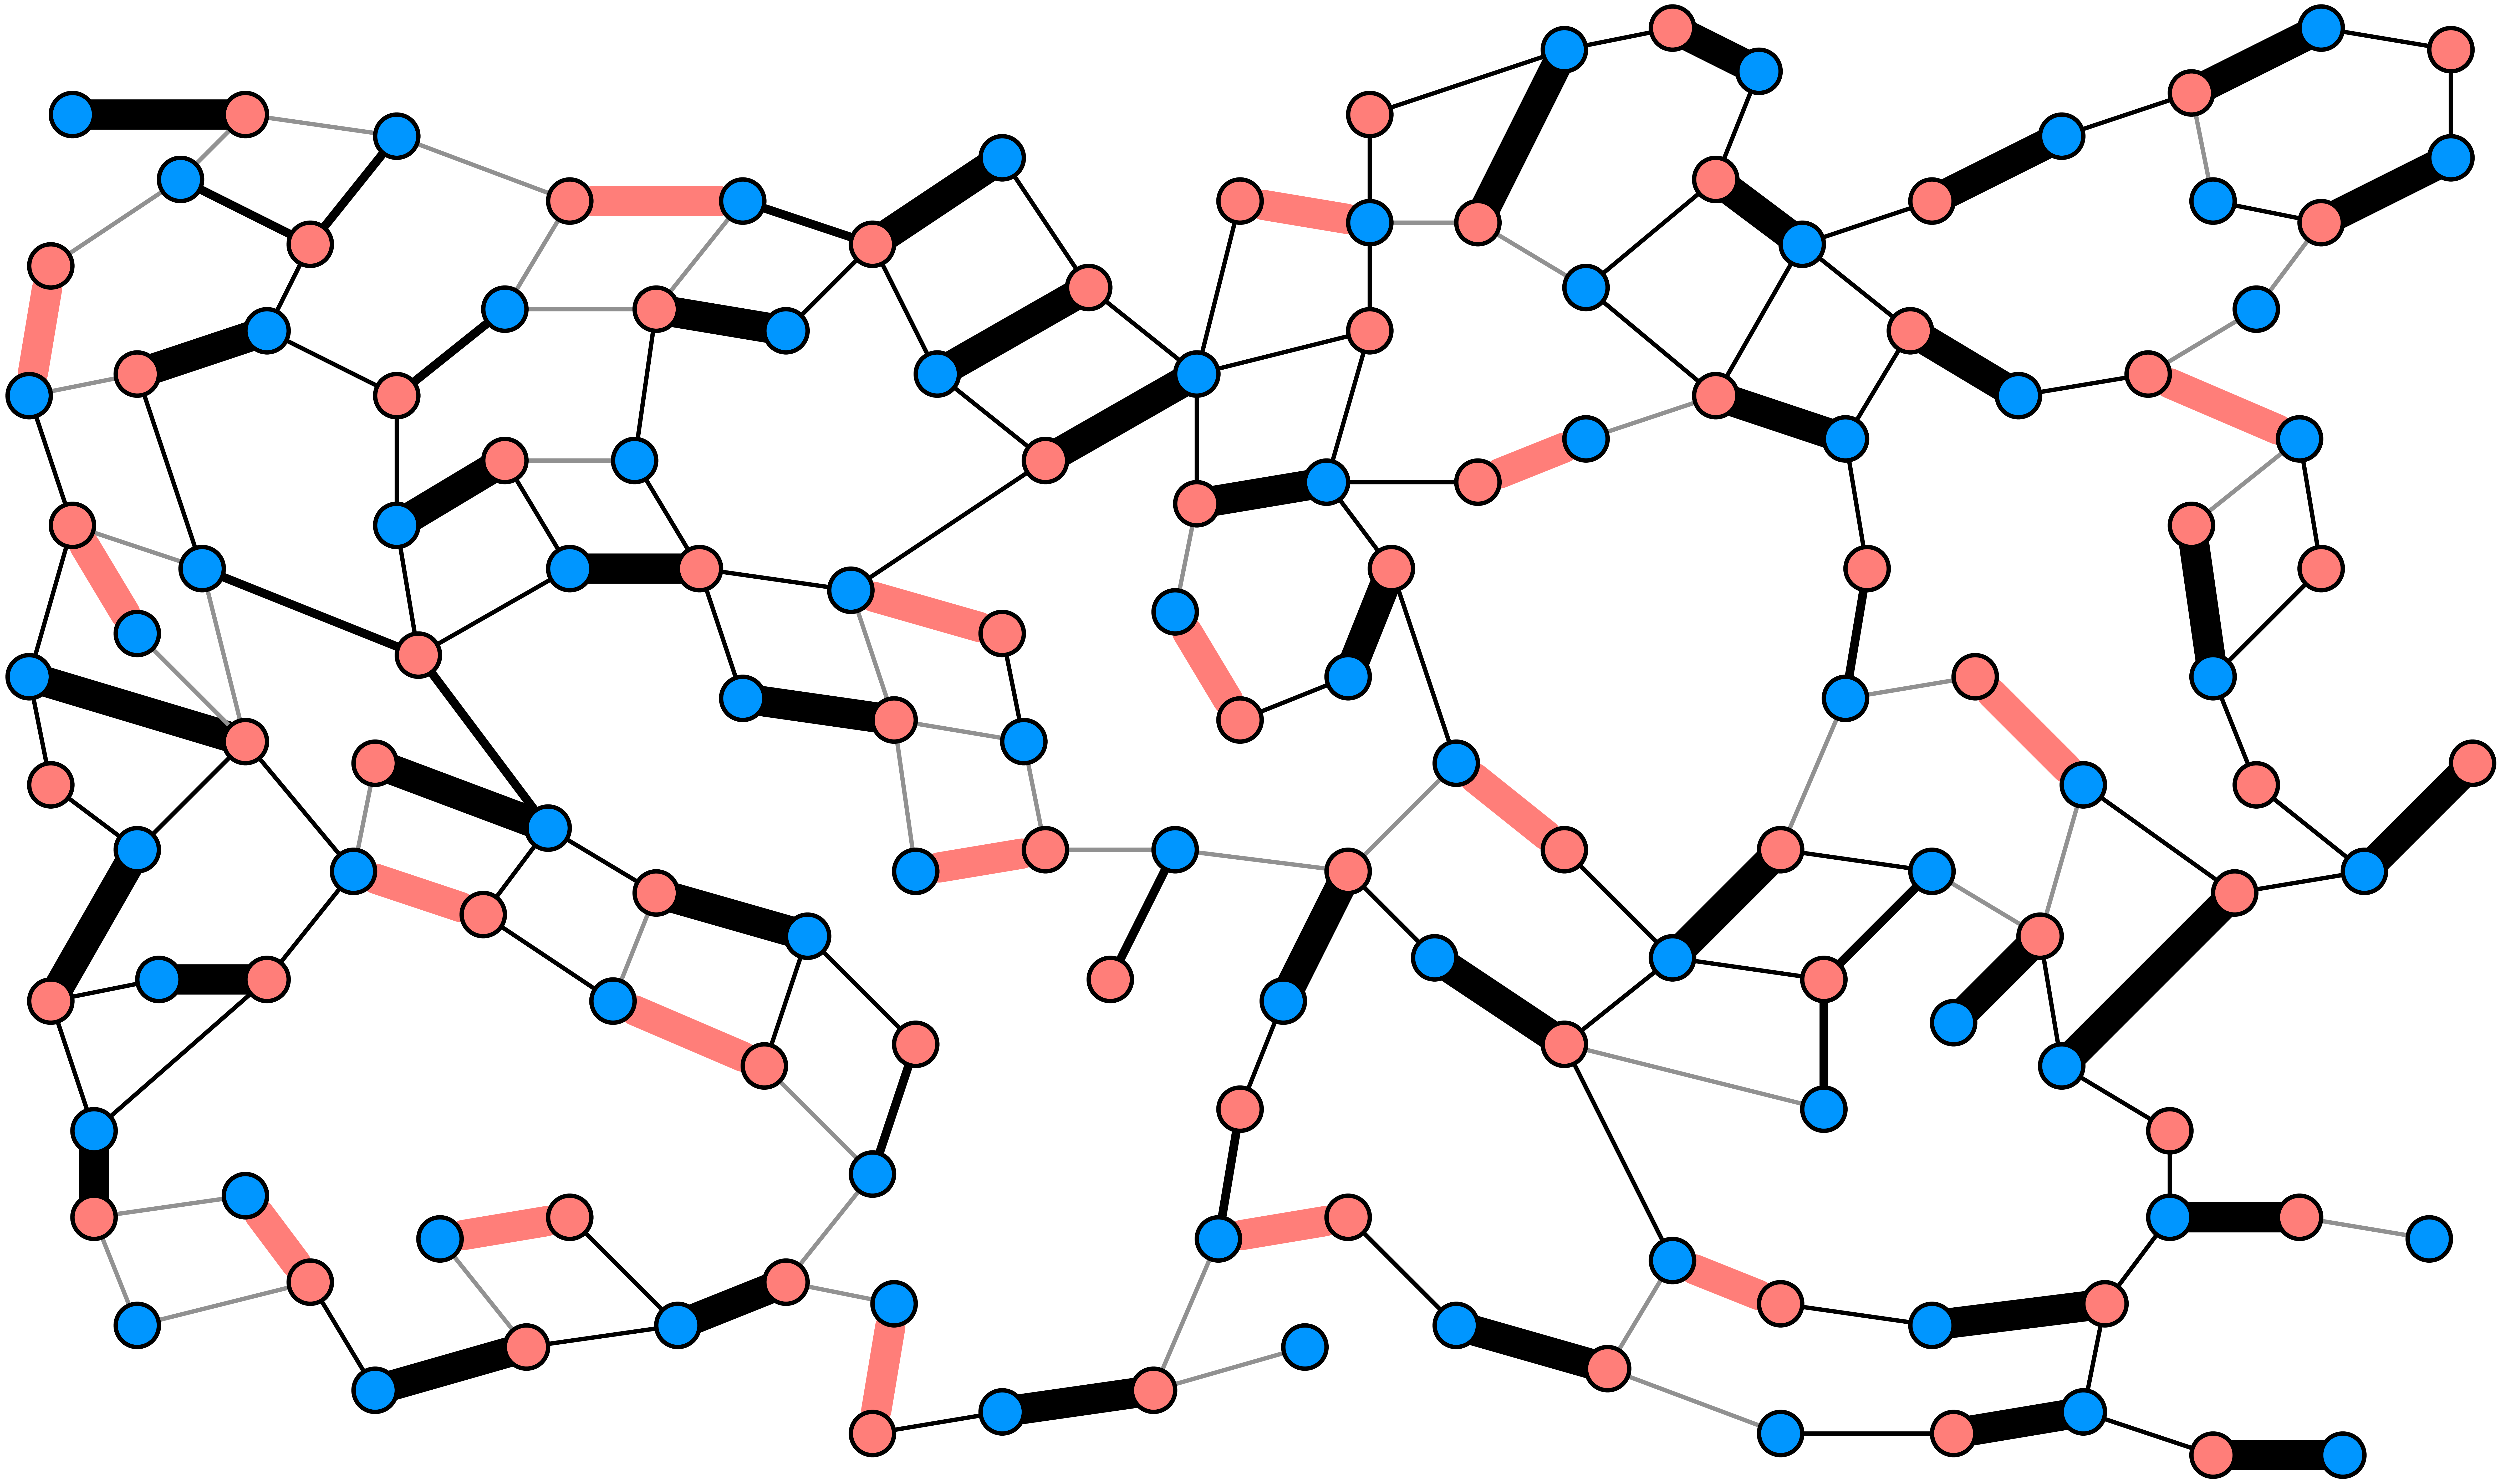

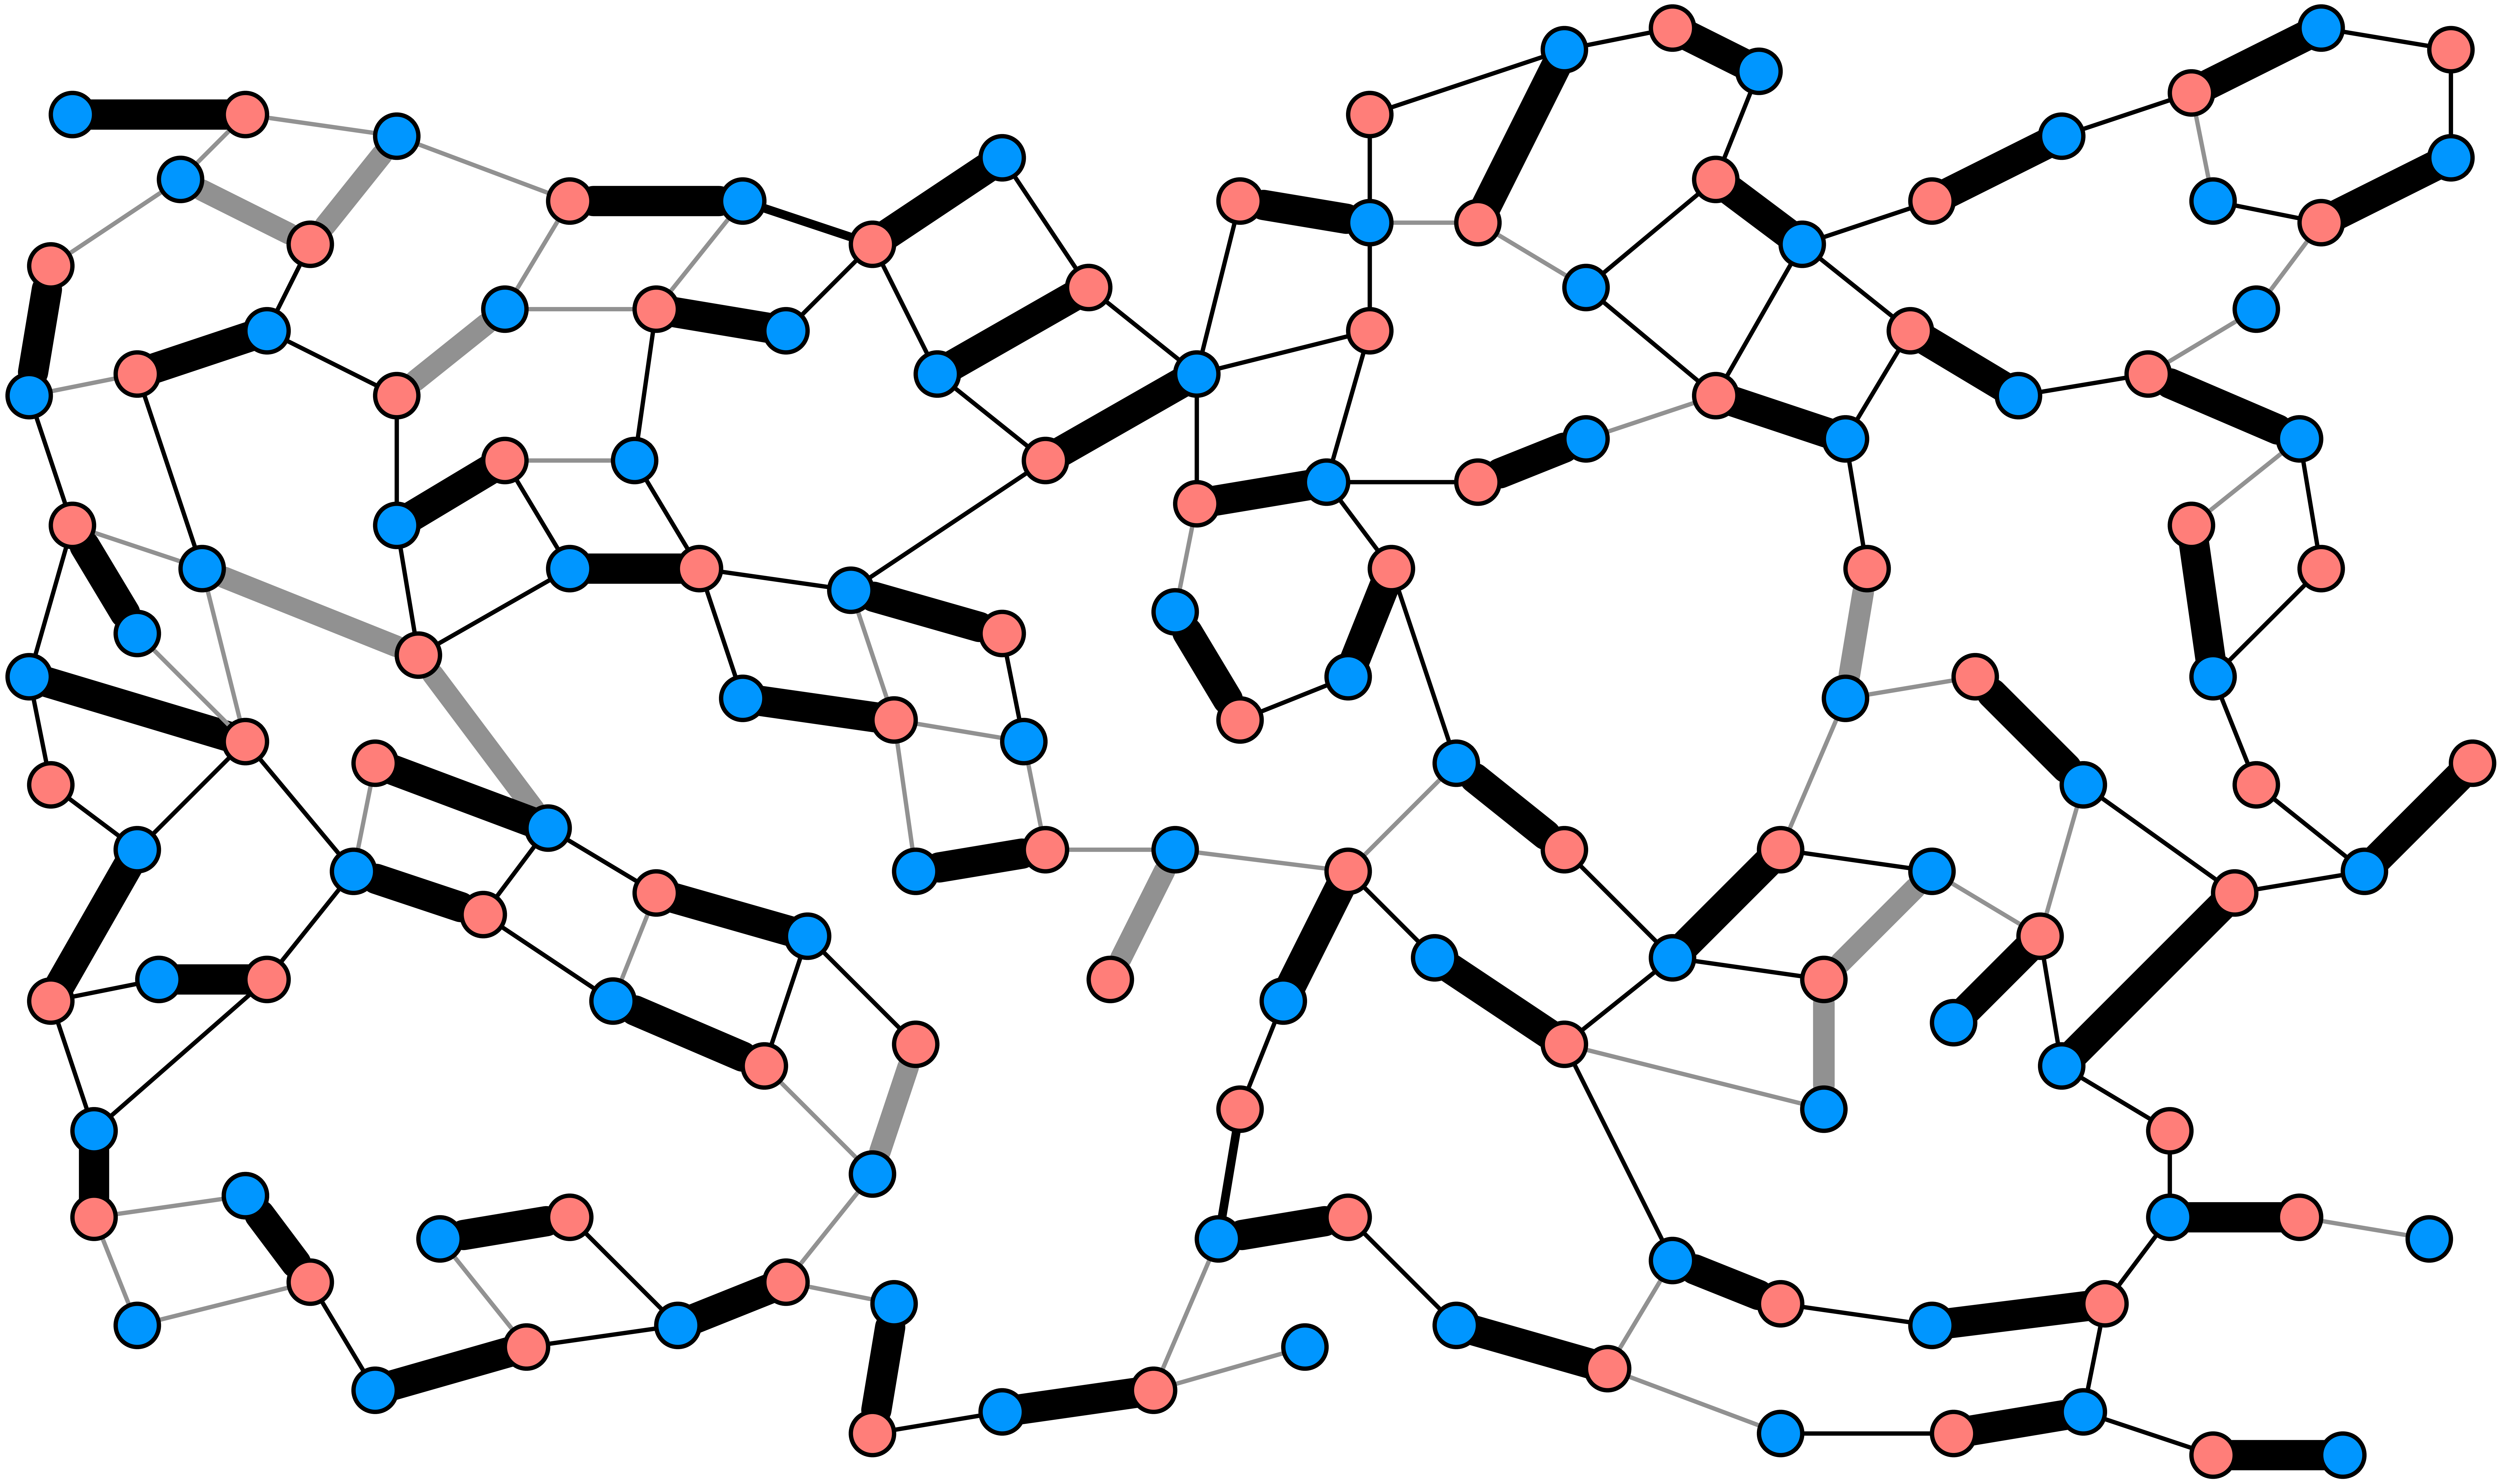

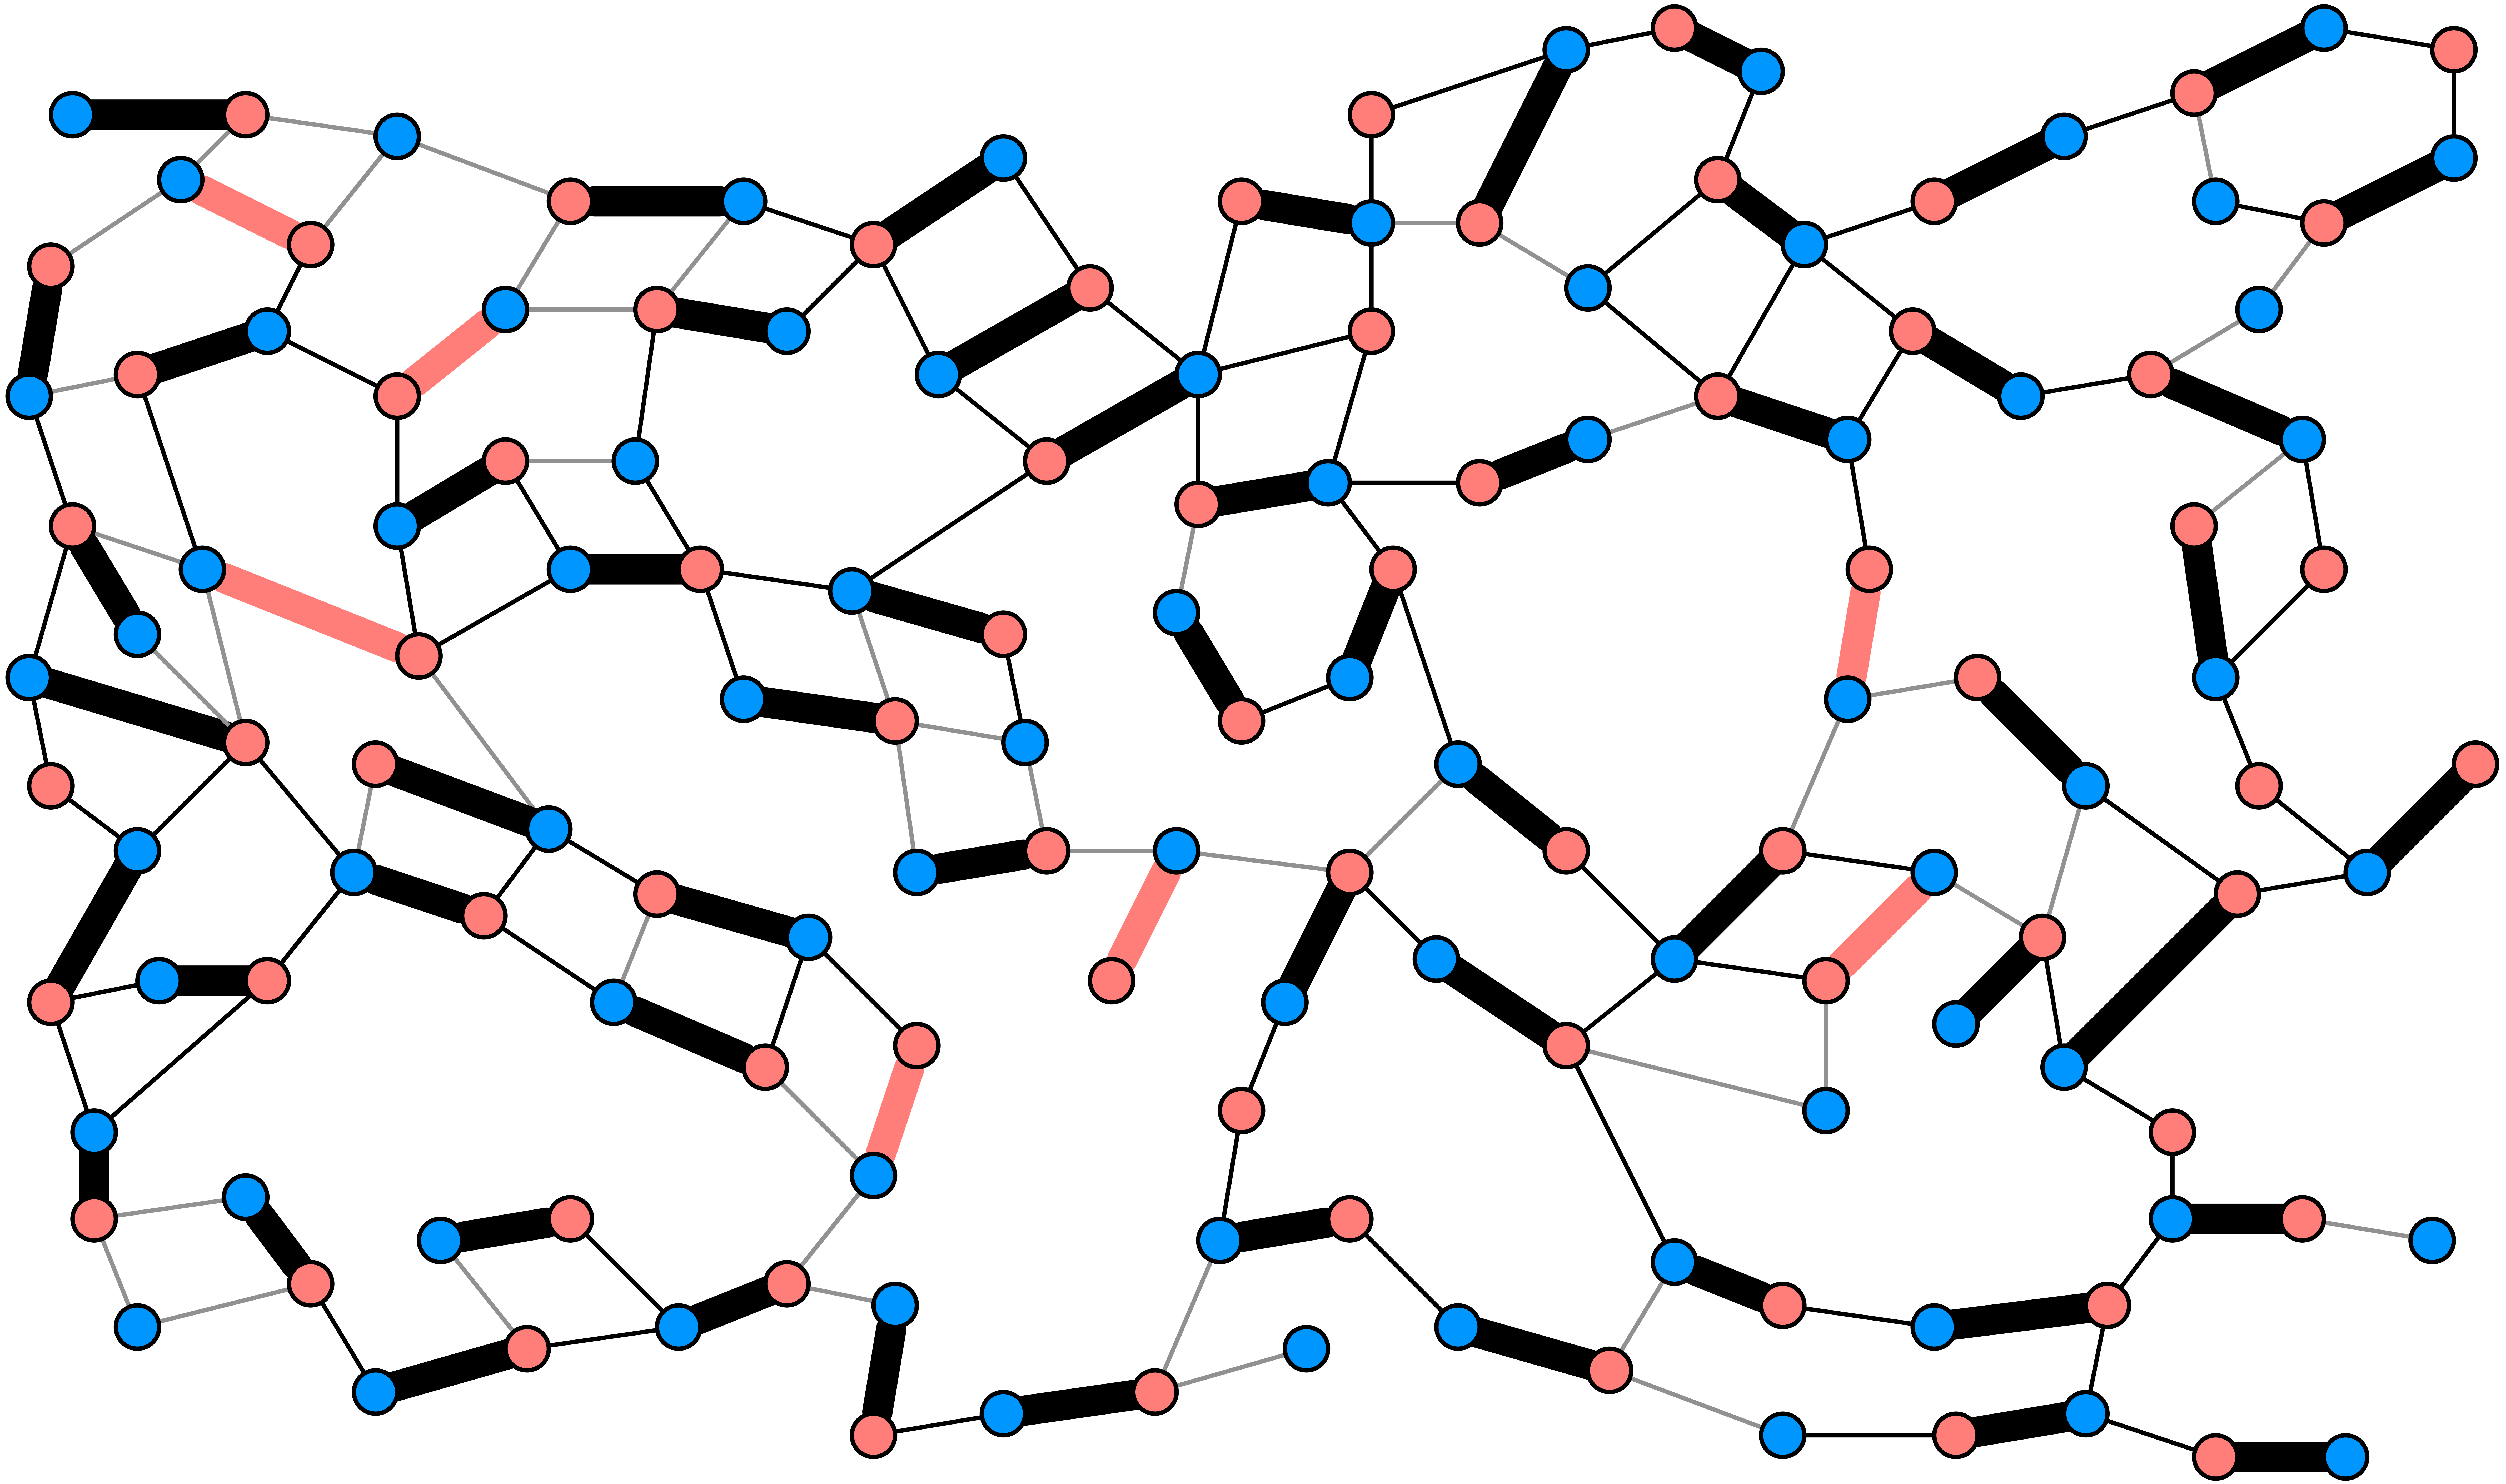

Finding a sinkless orientation

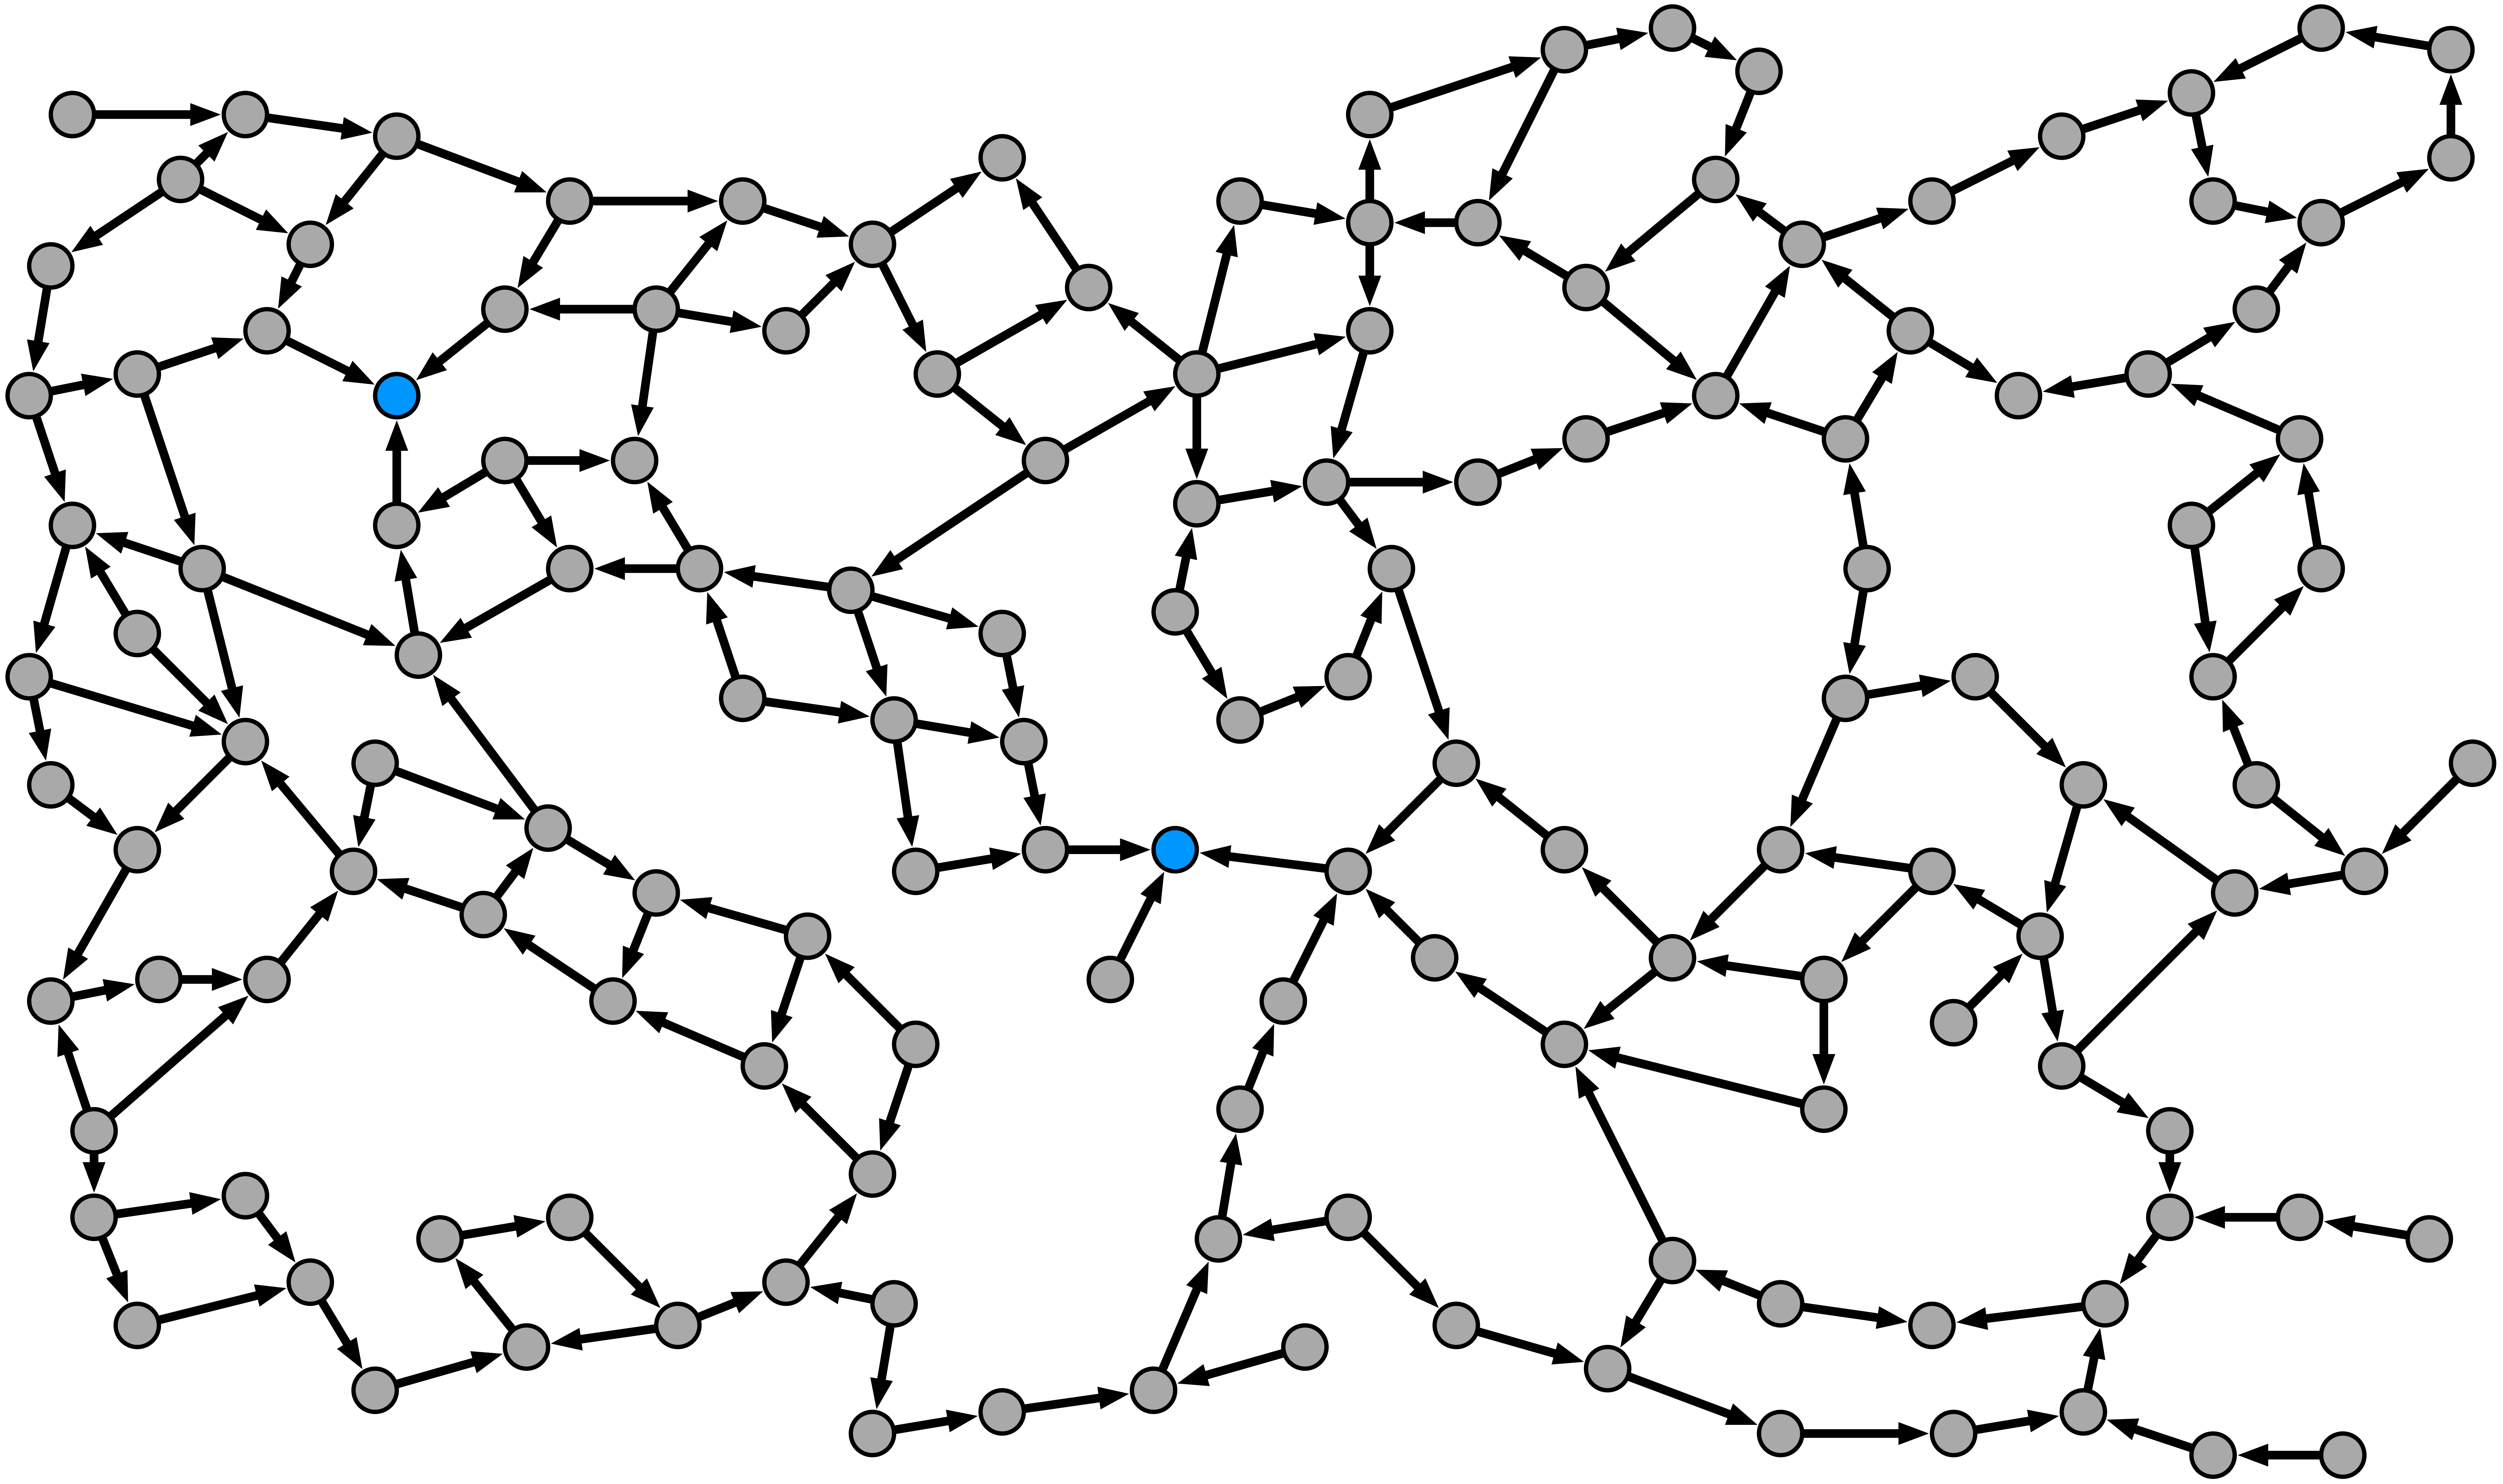

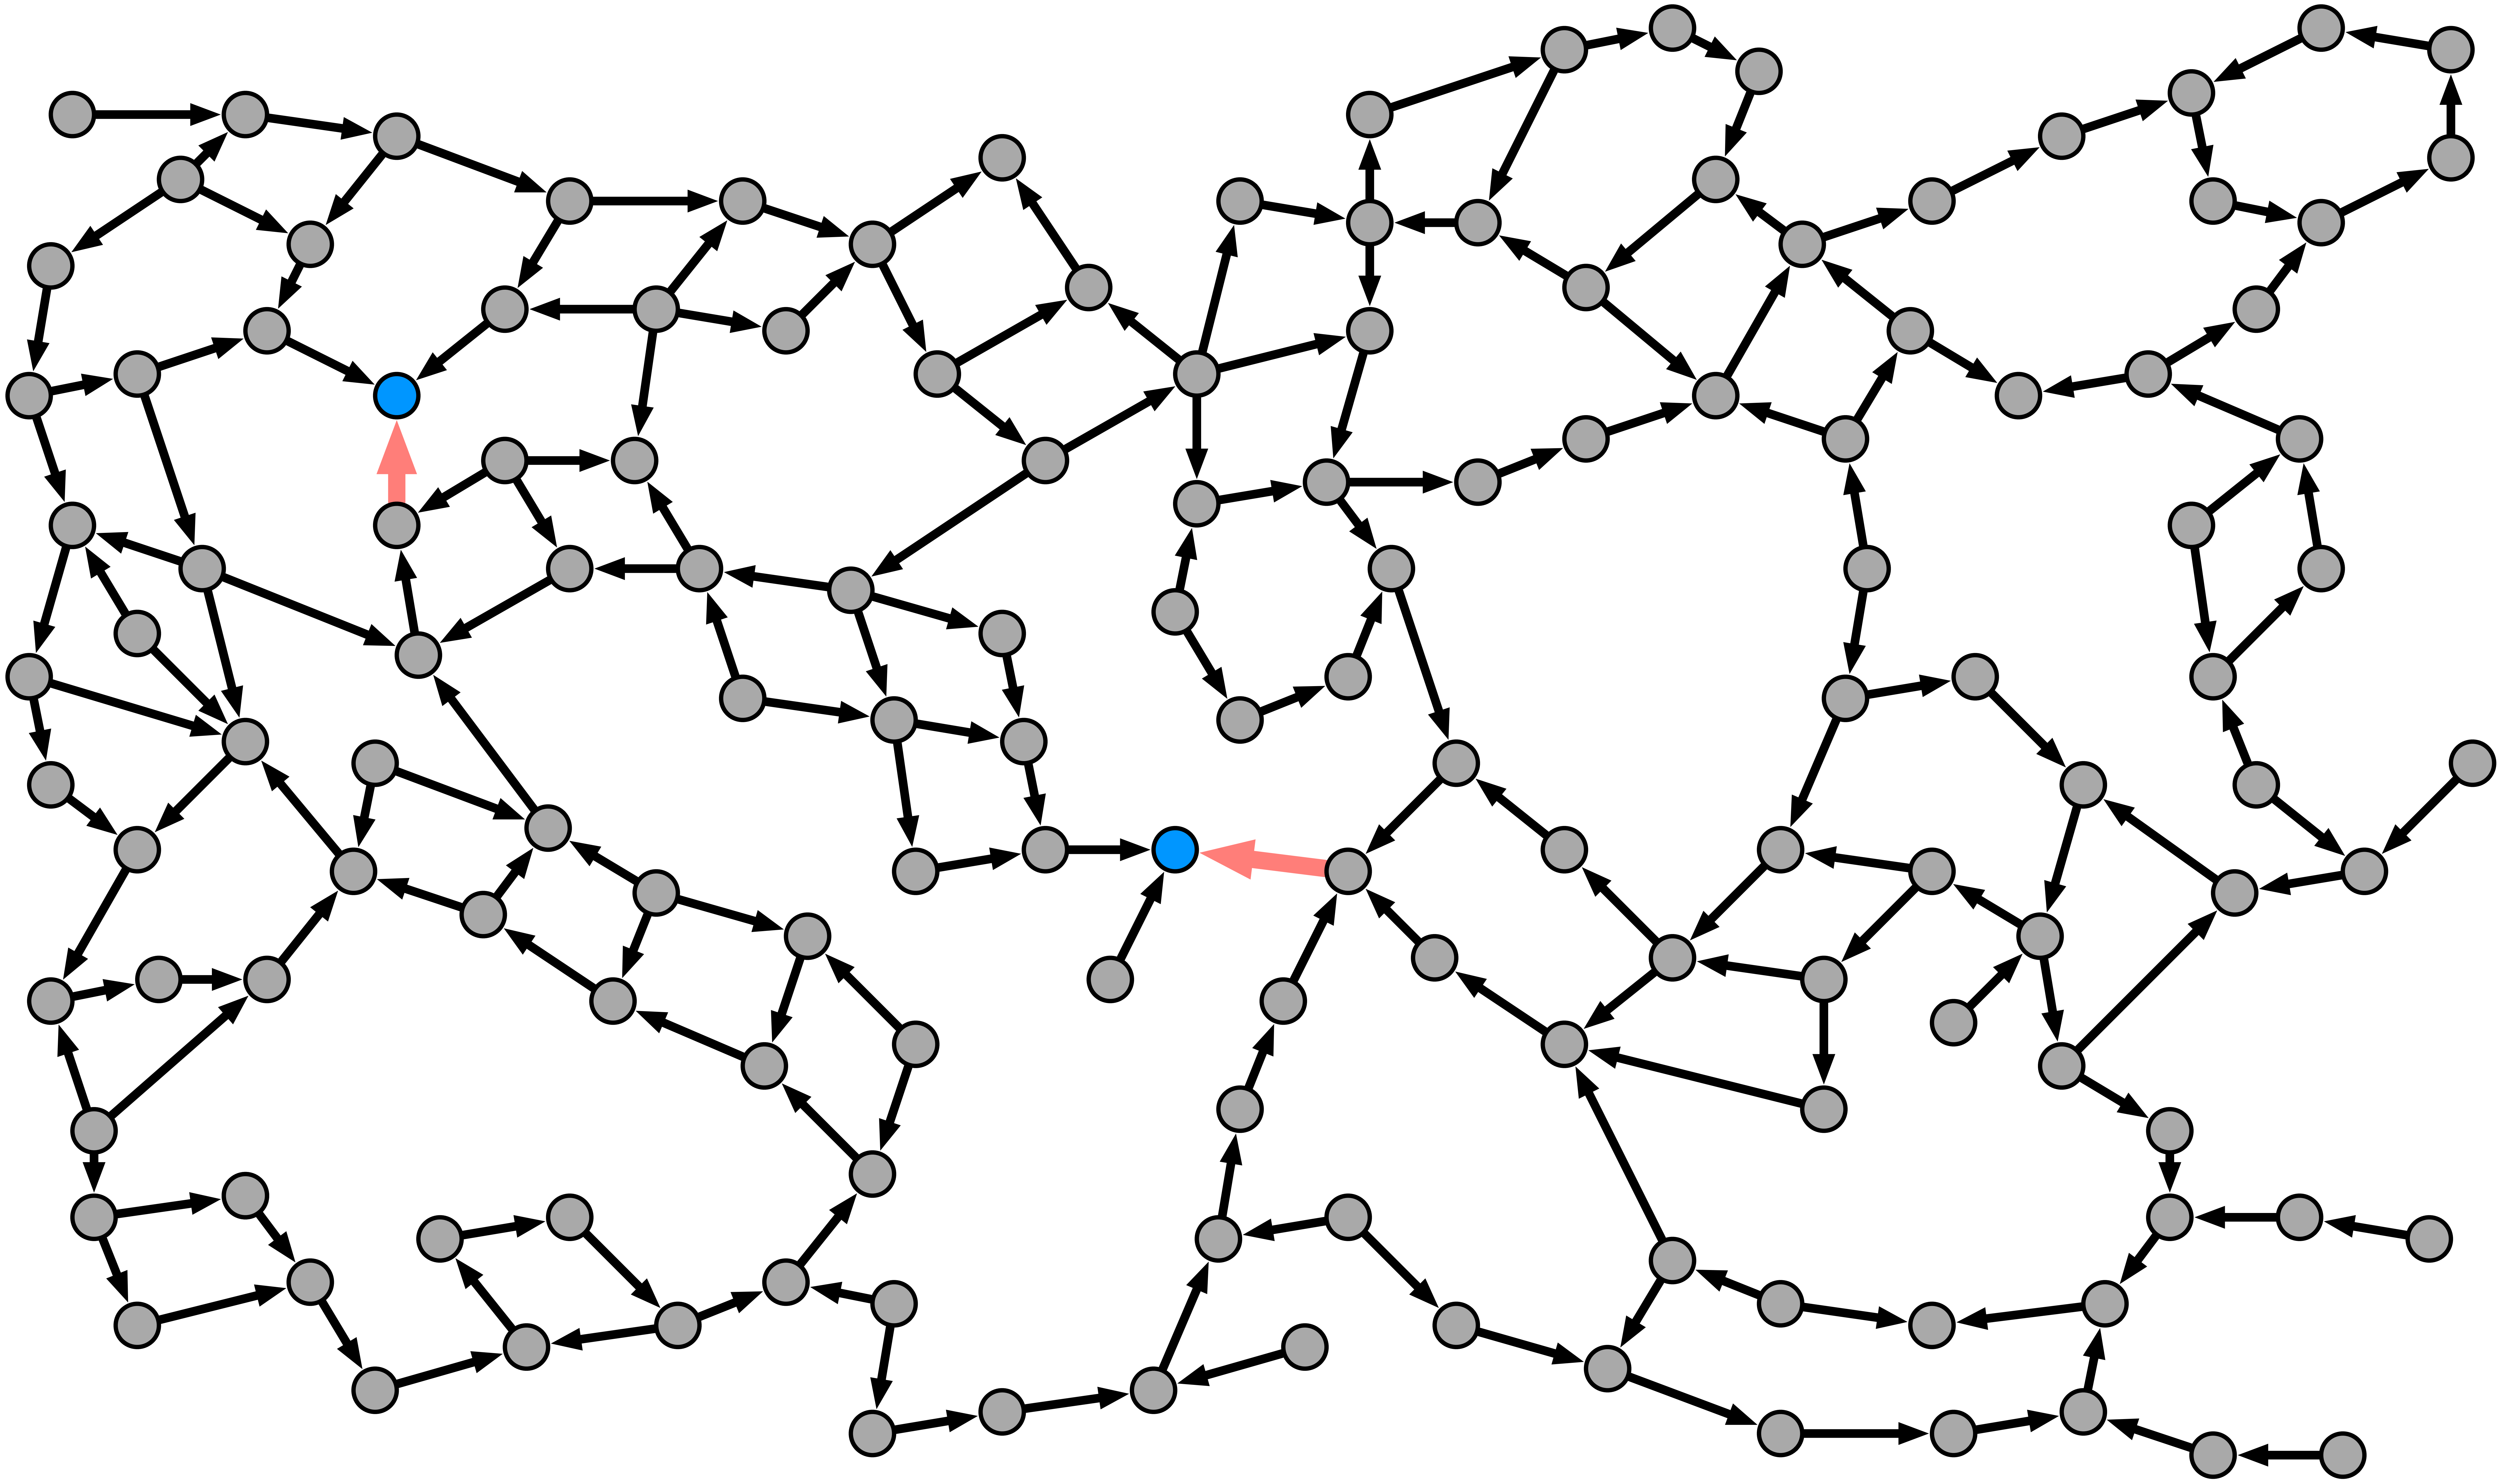

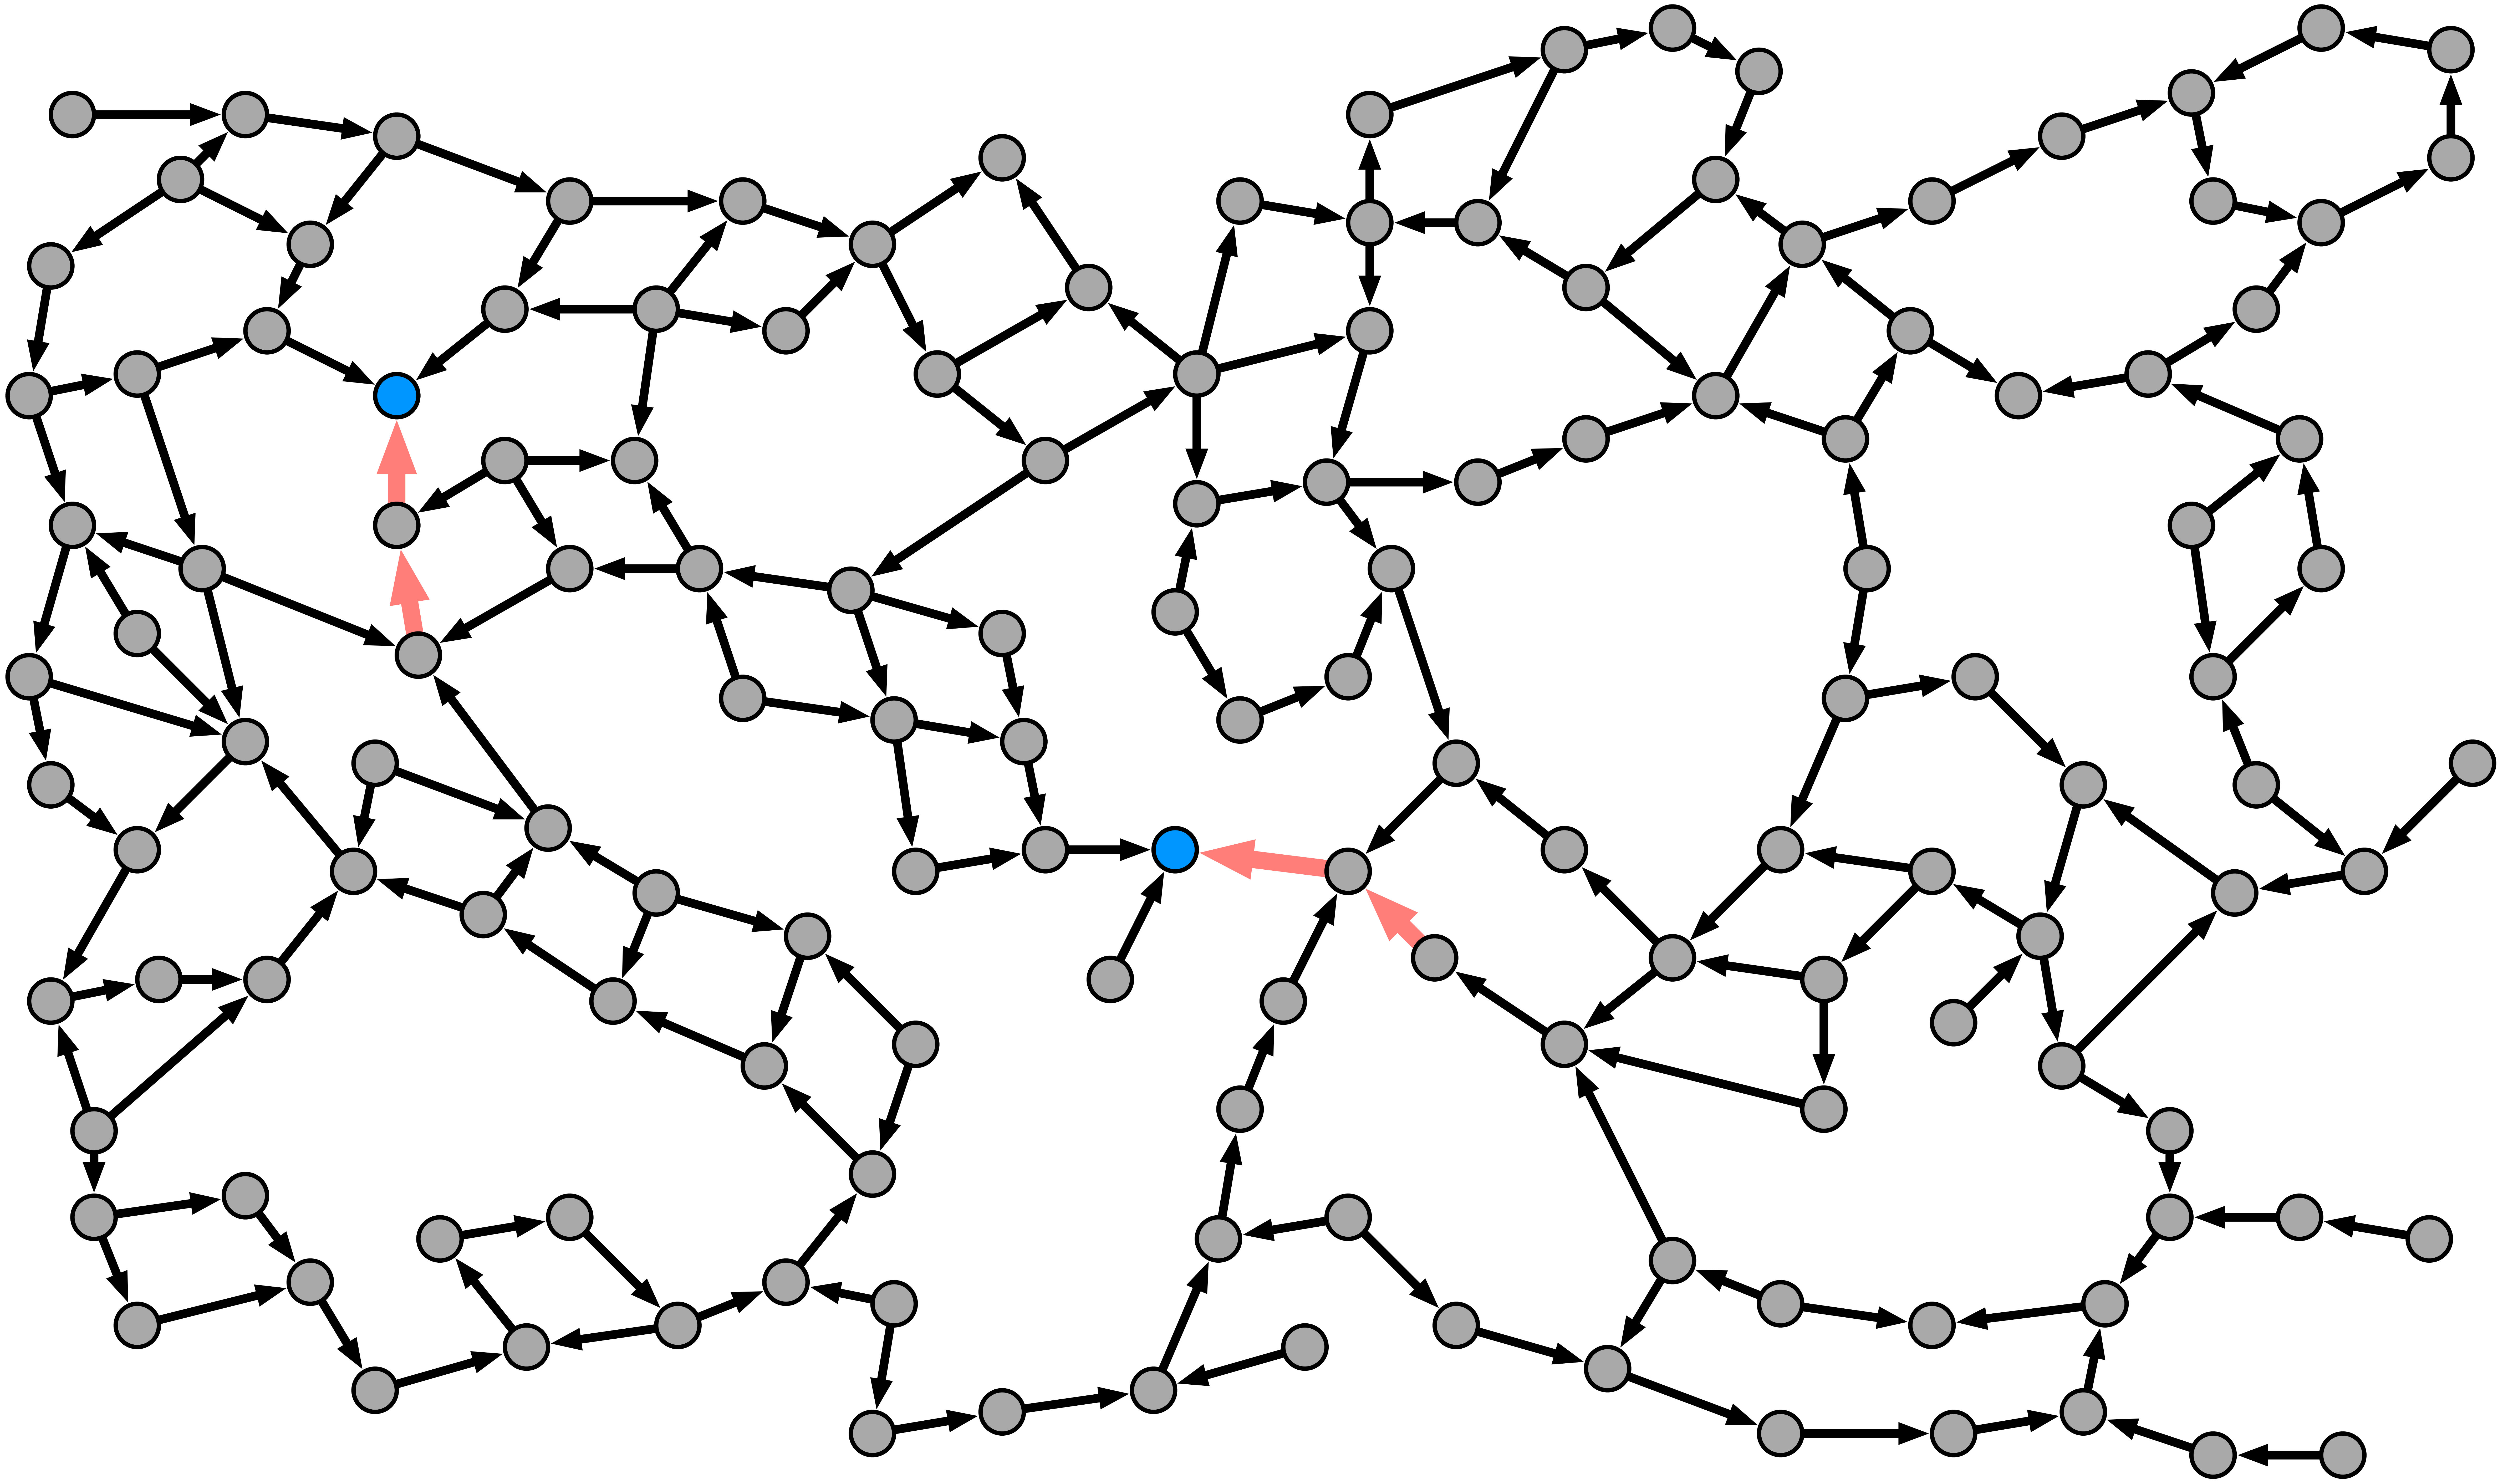

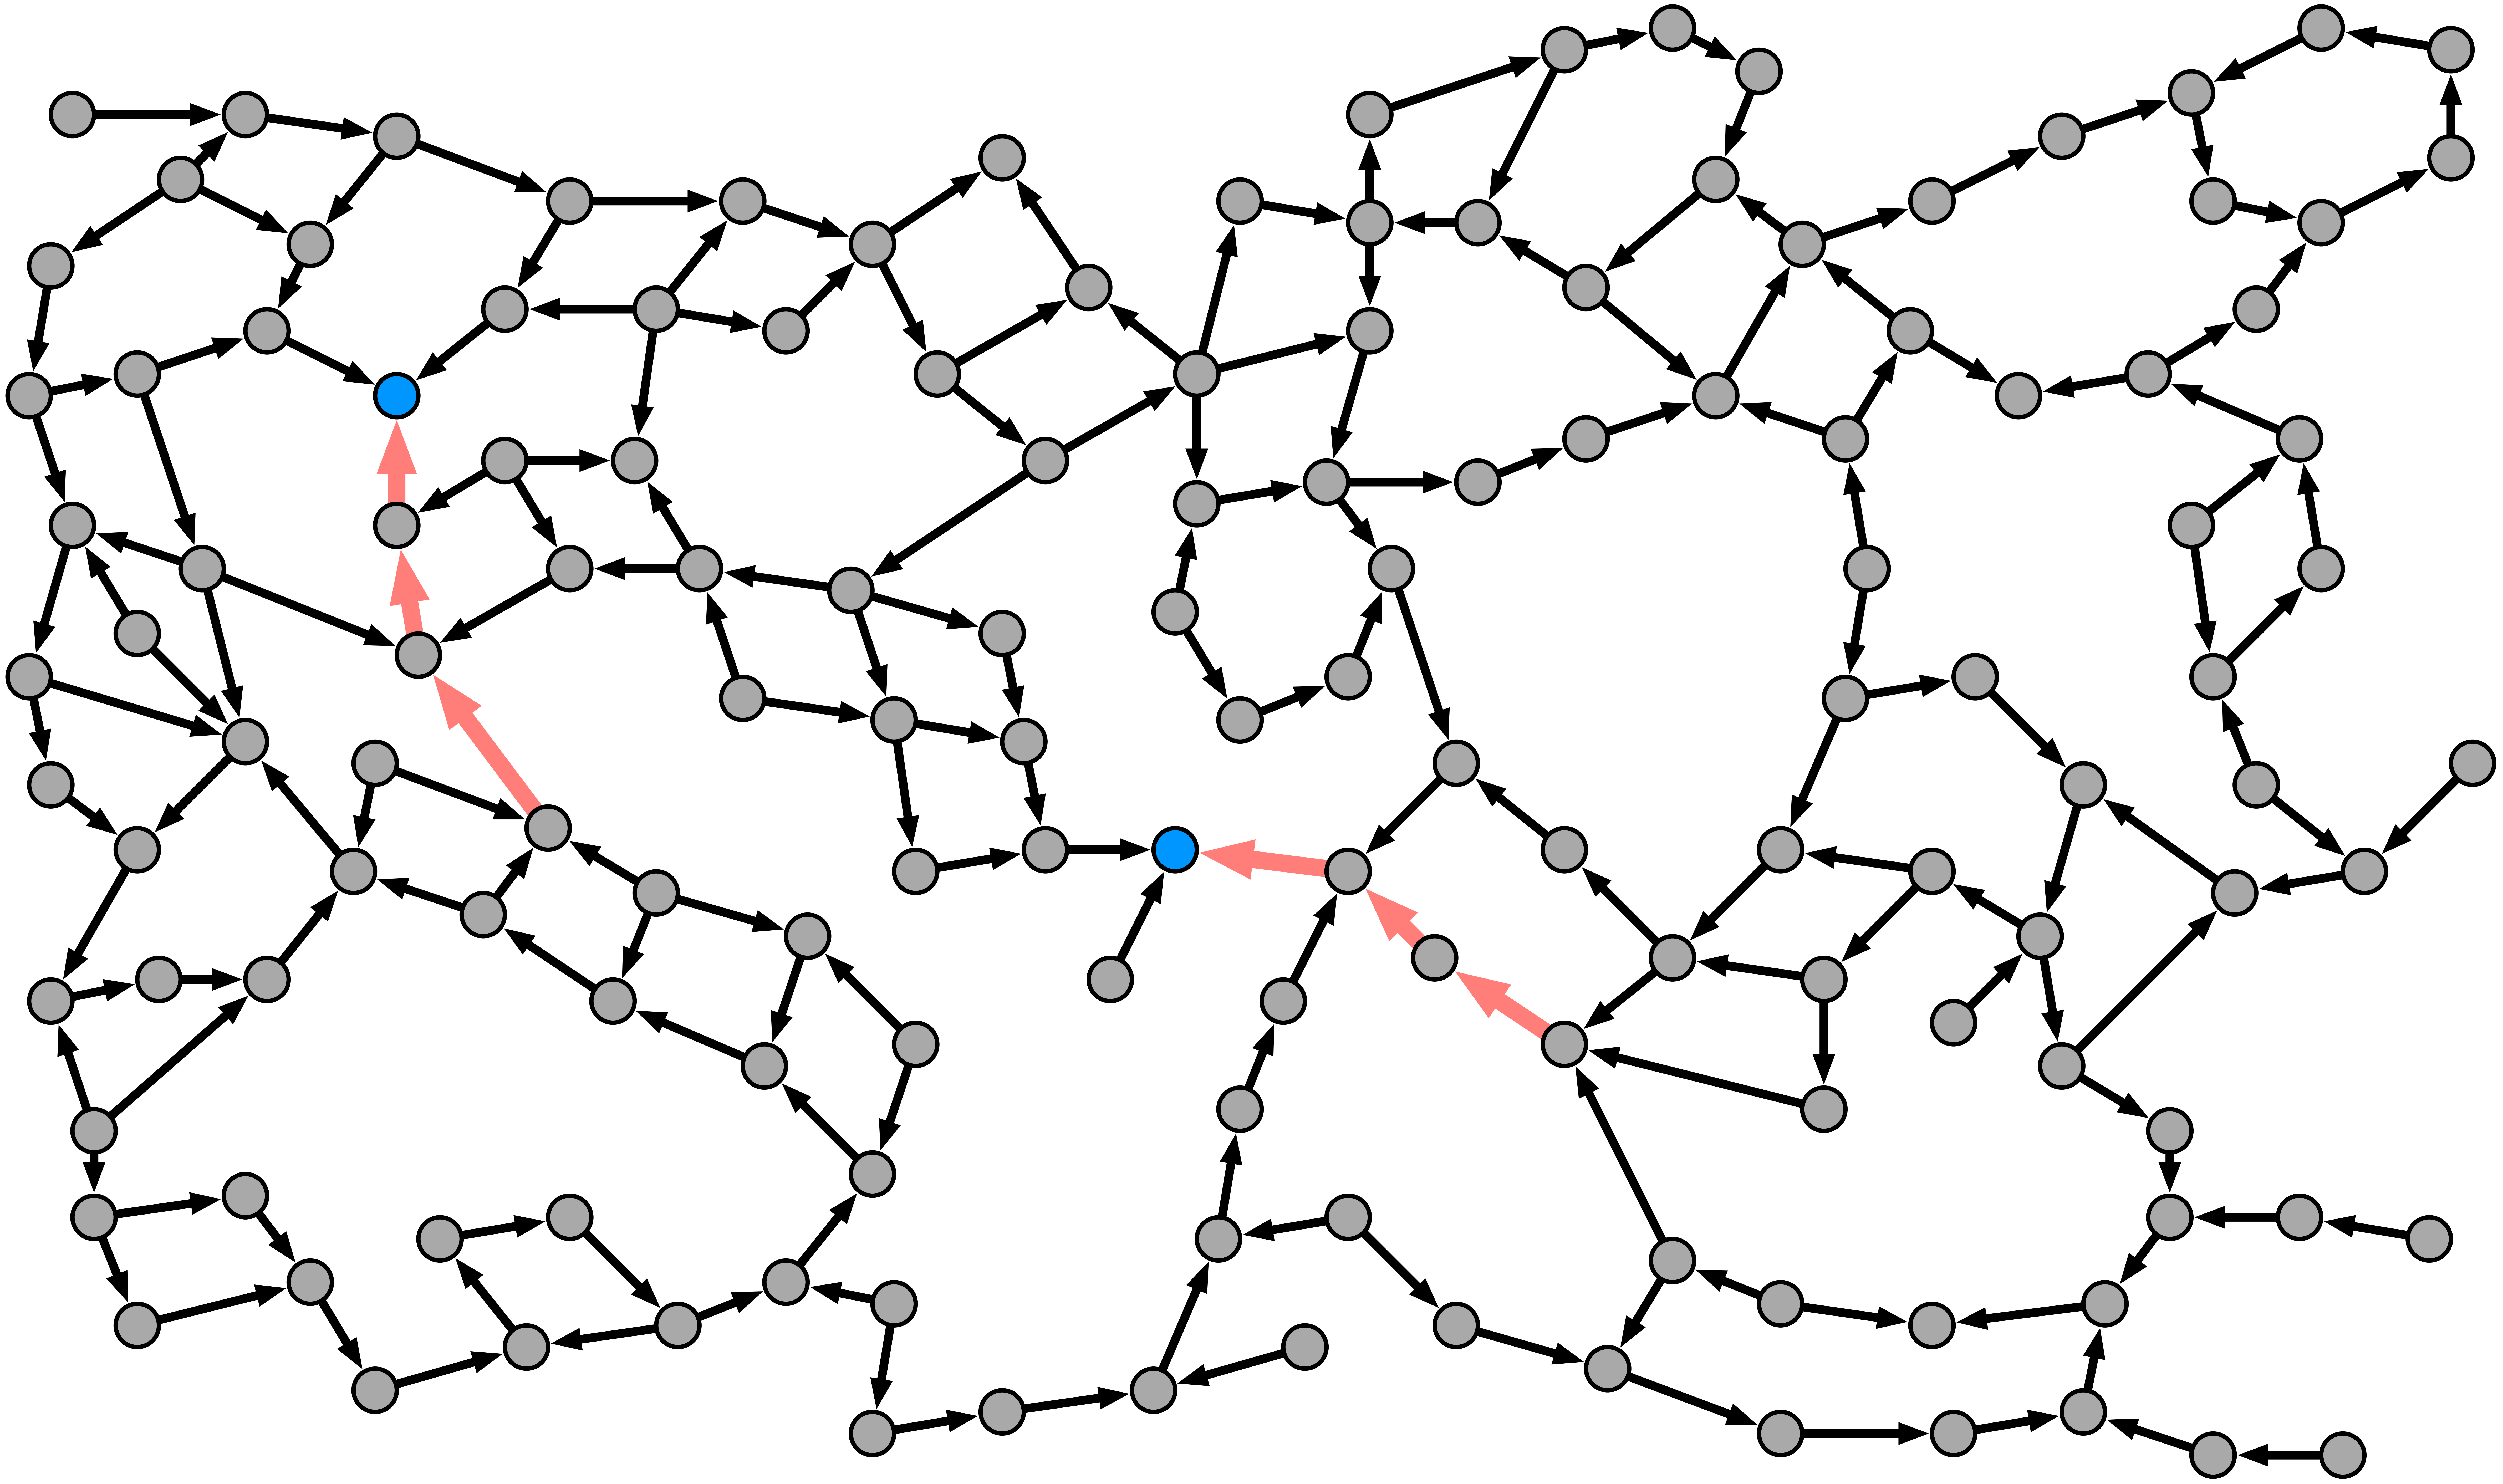

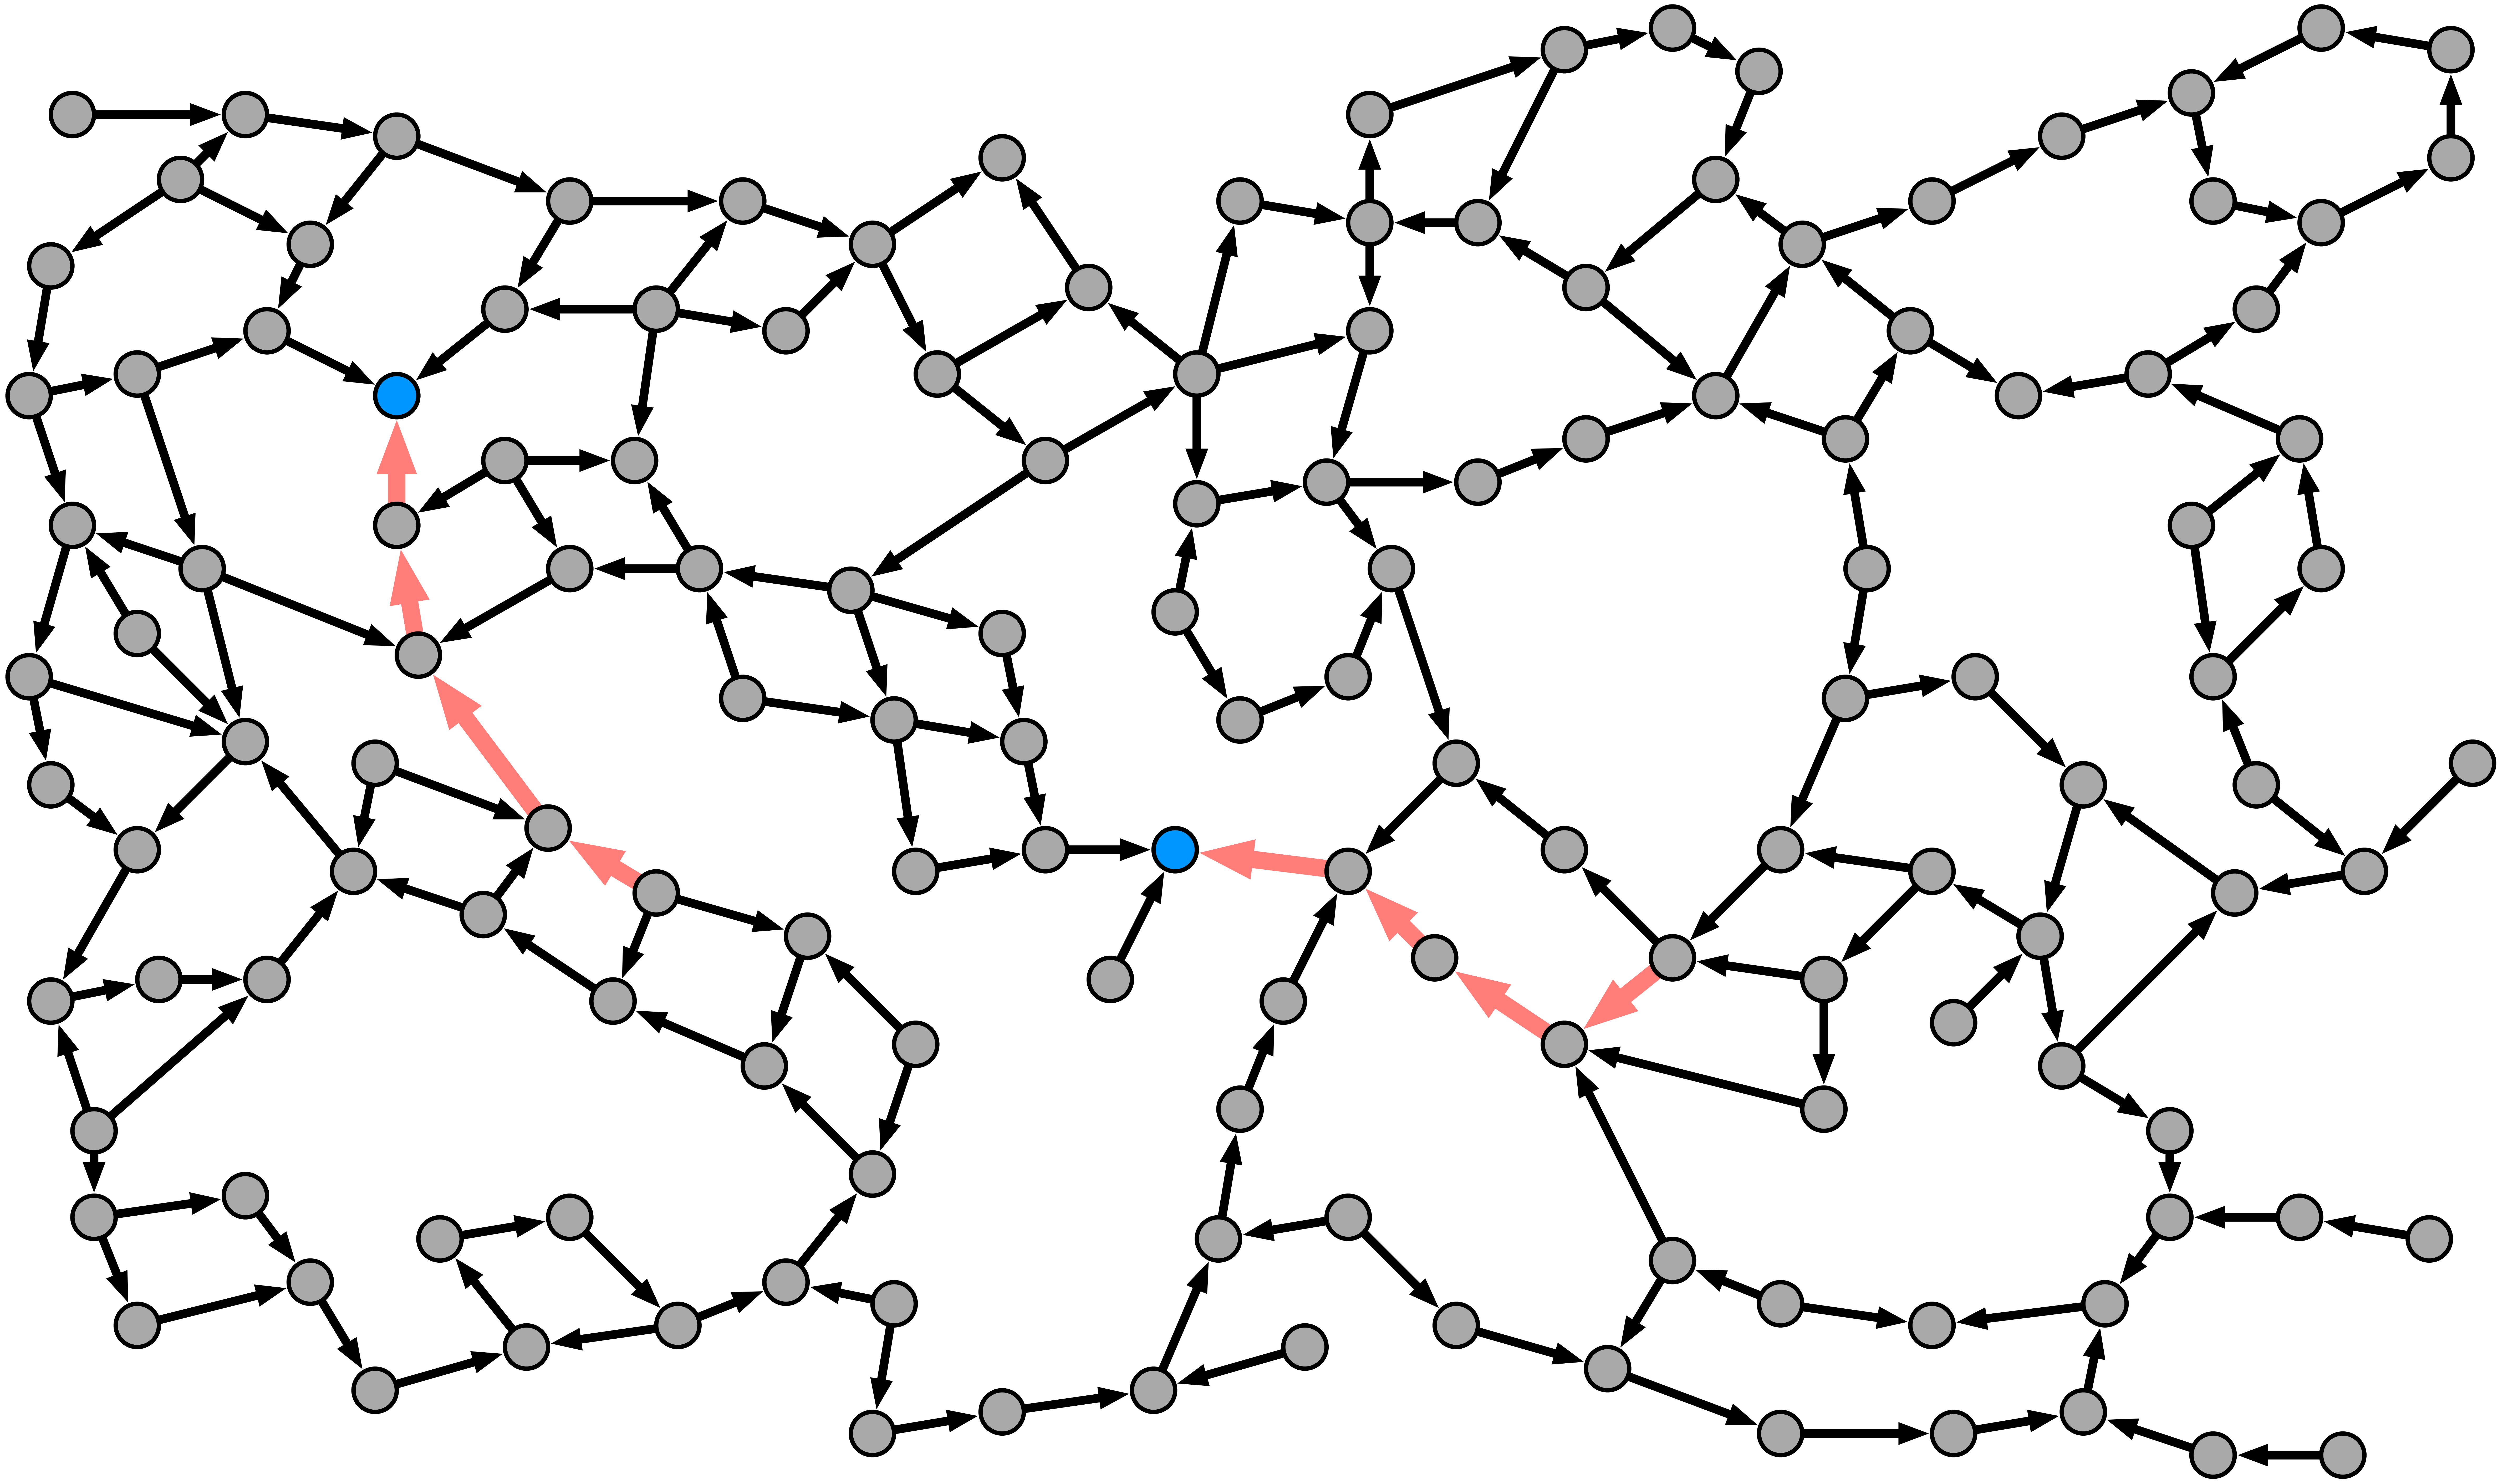

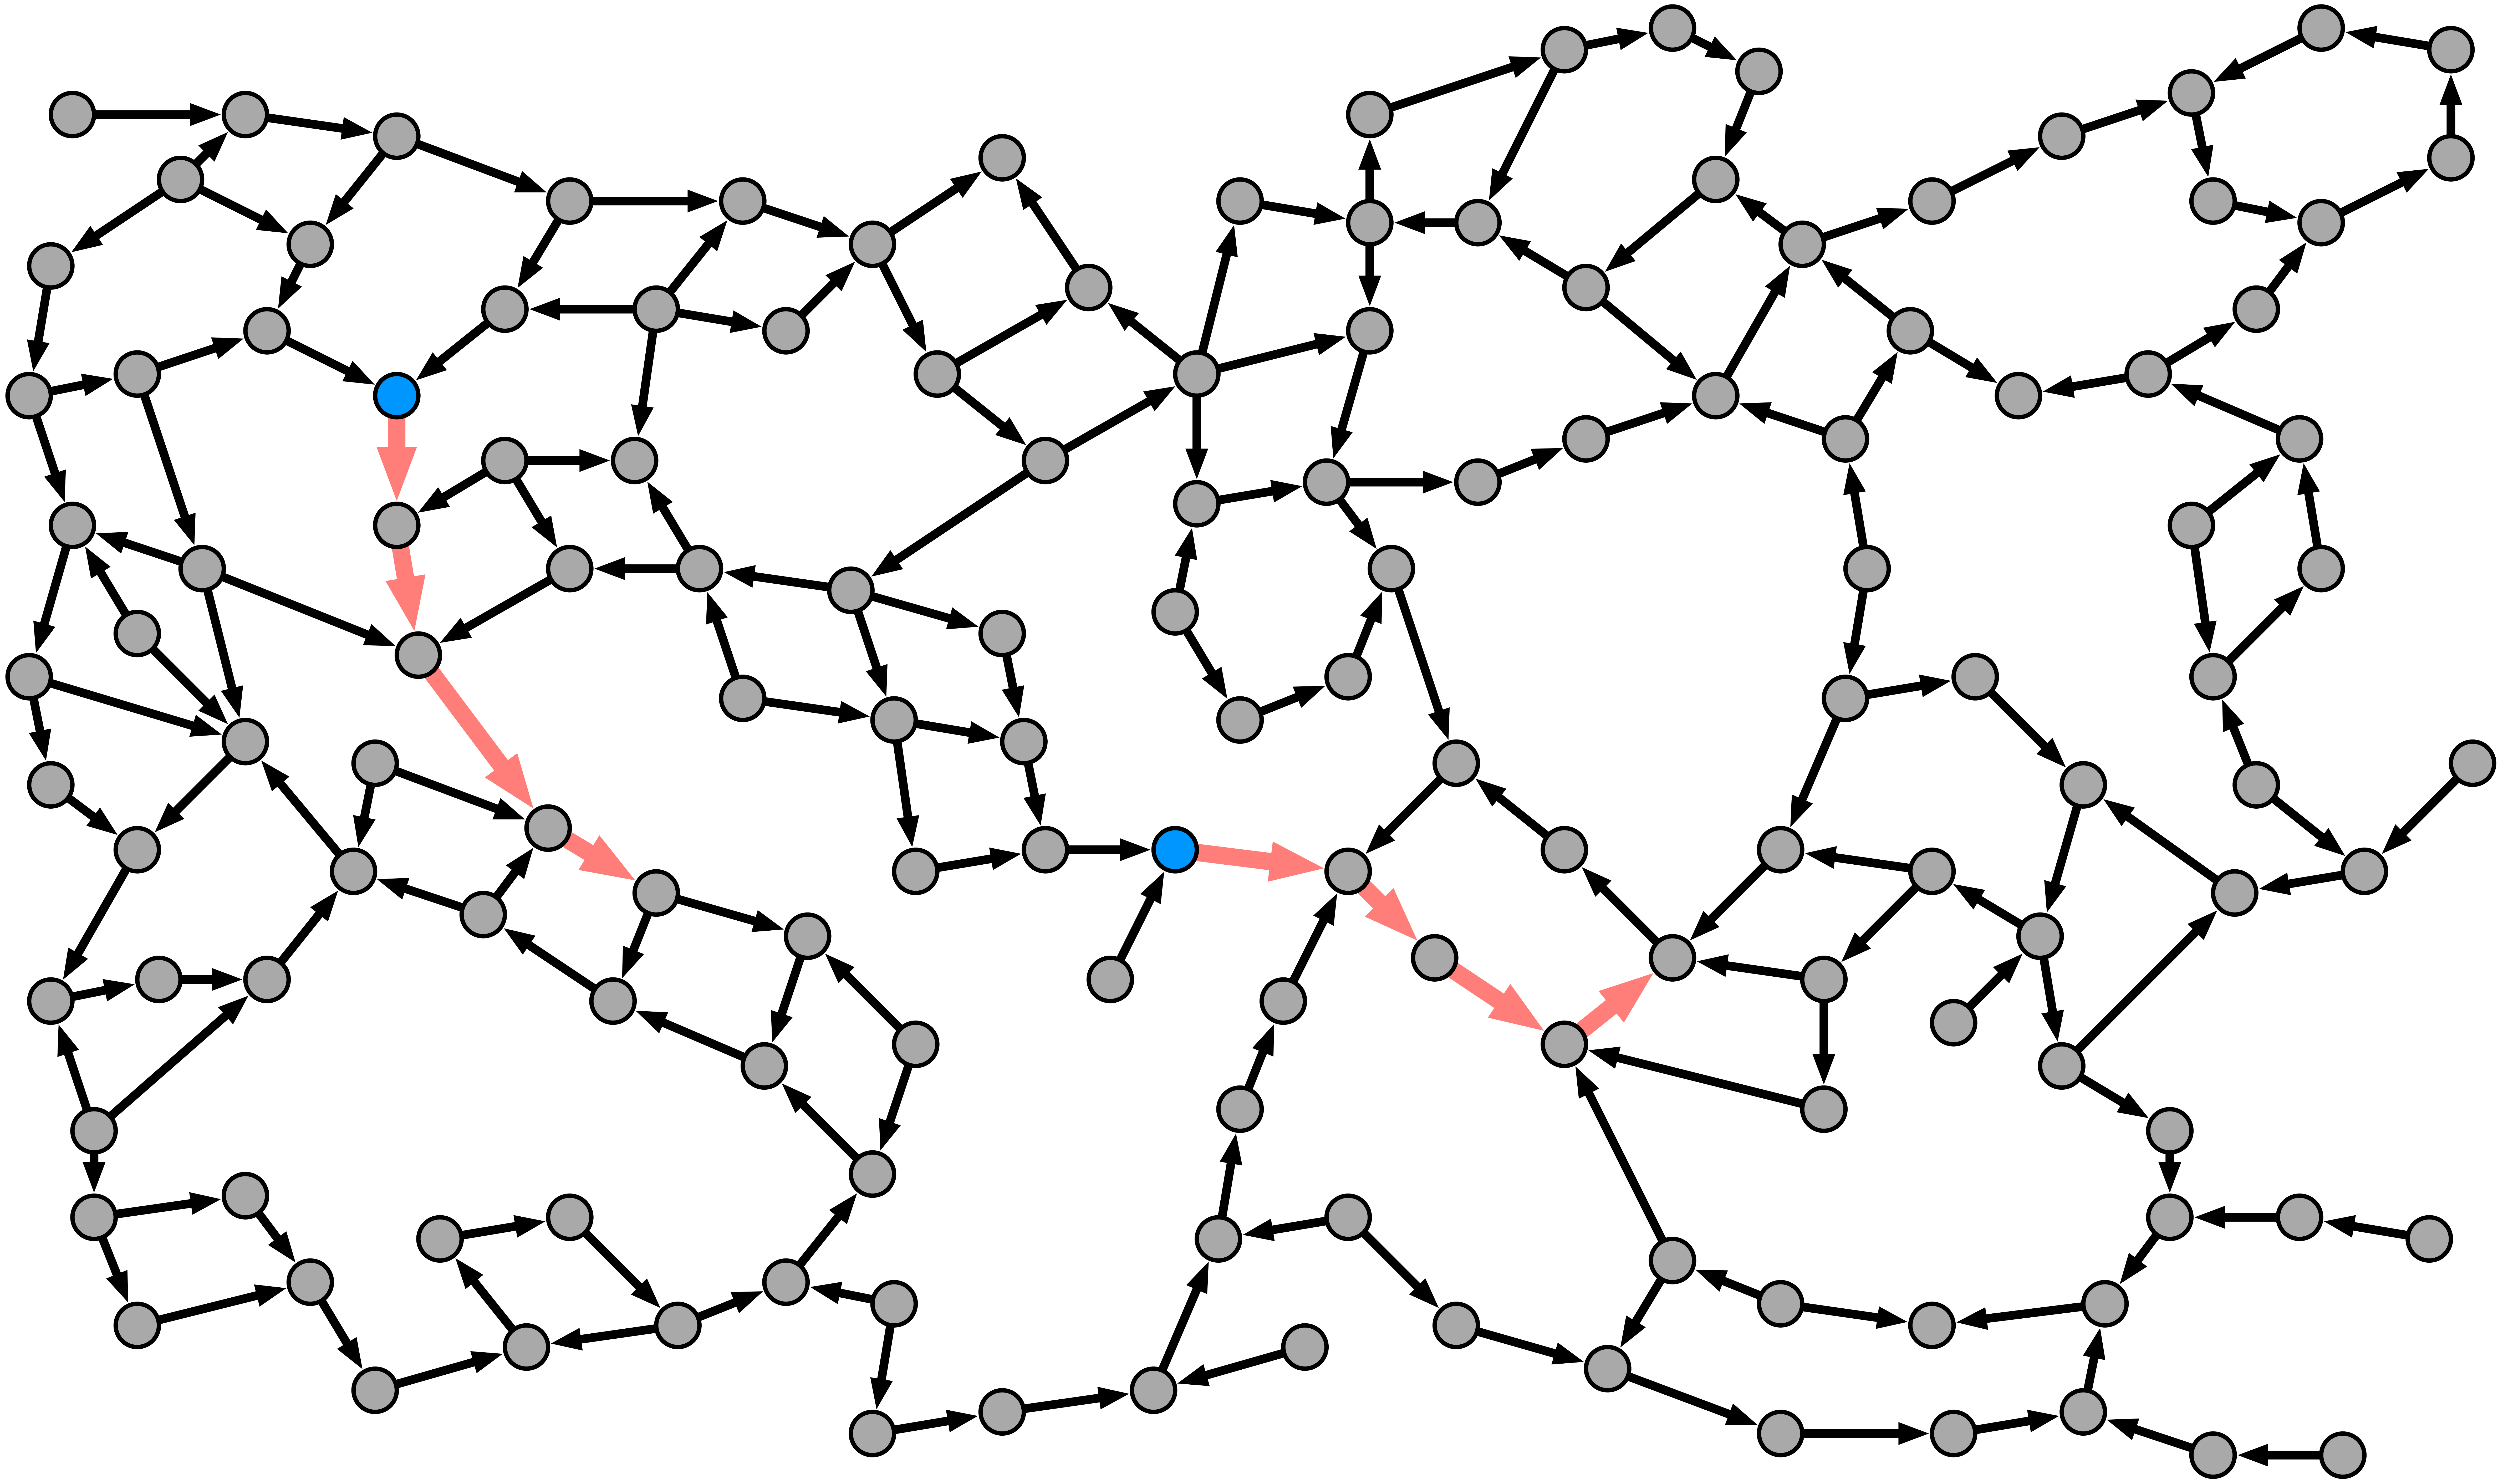

Best possible?

How **local** can distributed algorithms be?

- Proposal algorithm is optimal (FOCS 2019 best paper)

  Balliu, Brandt, **Hirvonen**, **Olivetti**, Rabie, and **Suomela**,

  arXiv:1901.02441

- Sinkless orientation is hard (STOC 2016)

  Brandt, Fischer, **Hirvonen**, **Keller**, Lempiäinen, Rybicki, **Suomela**, and **Uitto**

  arXiv:1511.00900

Speedup simulation

Assume problem $\mathbf{\Pi_0}$**:** can be solved in **T** rounds

↓

There exists problem $\mathbf{\Pi_1}$**:** can be solved in **T-1** rounds

⋮ mechanical transform

↓

There exists problem $\mathbf{\Pi_T}$**:** can be solved in **0** rounds

Contradiction?

2016: state of the art

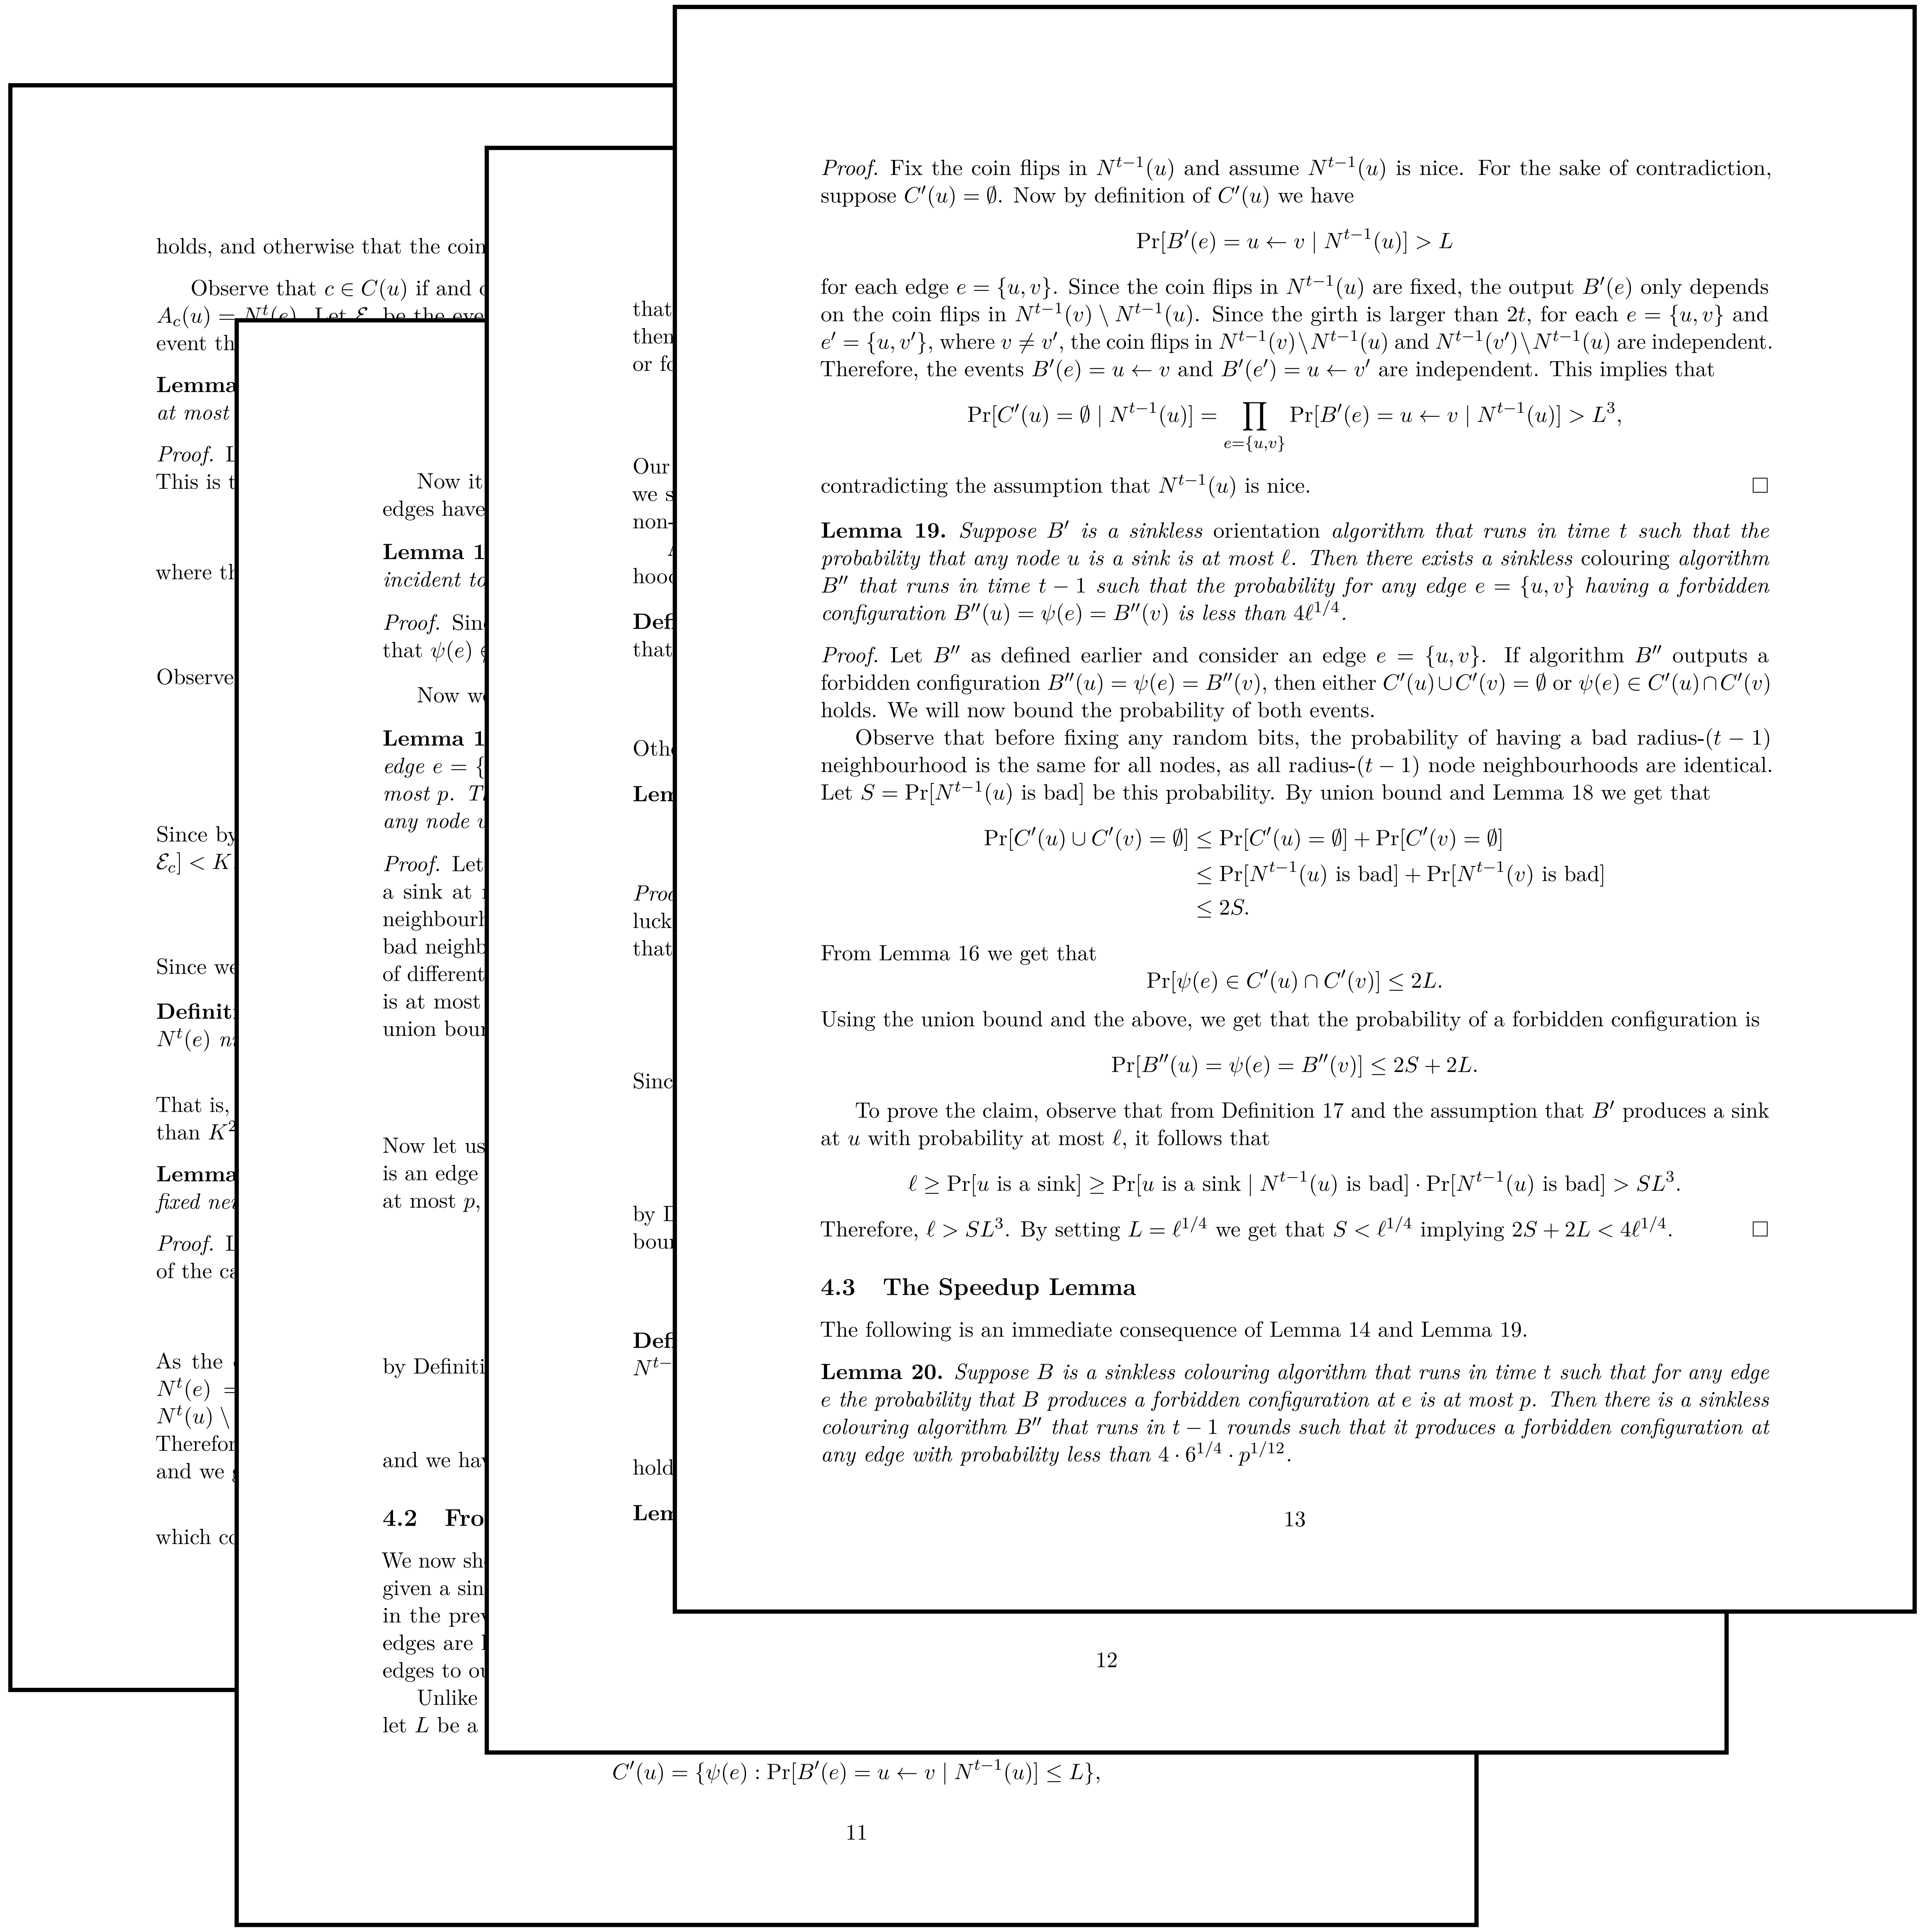

Proof. Fix the coin flips in $N^{t-1}(u)$ and assume $N^{t-1}(u)$ is nice. For the sake of contradiction, suppose $C'(u) = \emptyset$. Now by definition of $C'(u)$ we have

$$\Pr[B'(e) = u \leftarrow v \mid N^{t-1}(u)] > L$$

for each edge $e = \{u, v\}$. Since the coin flips in $N^{t-1}(u)$ are fixed, the output $B'(e)$ only depends on the coin flips in $N^{t-1}(v) \setminus N^{t-1}(u)$. Since the girth is larger than $2t$, for each $e = \{u, v\}$ and $e' = \{u, v'\}$, where $v \neq v'$, the coin flips in $N^{t-1}(v) \setminus N^{t-1}(u)$ and $N^{t-1}(v') \setminus N^{t-1}(u)$ are independent. Therefore, the events $B'(e) = u \leftarrow v$ and $B'(e') = u \leftarrow v'$ are independent. This implies that

$$\Pr[C'(u) = \emptyset \mid N^{t-1}(u)] = \prod_{e=\{u,v\}} \Pr[B'(e) = u \leftarrow v \mid N^{t-1}(u)] > L^3,$$

contradicting the assumption that $N^{t-1}(u)$ is nice. □

Lemma 19. *Suppose B' is a sinkless* orientation *algorithm that runs in time t such that the probability that any node u is a sink is at most ℓ. Then there exists a sinkless* colouring *algorithm B'' that runs in time $t-1$ such that the probability for any edge $e = \{u, v\}$ having a forbidden configuration $B''(u) = \psi(e) = B''(v)$ is less than $4\ell^{1/4}$.*

Proof. Let B'' as defined earlier and consider an edge $e = \{u, v\}$. If algorithm B'' outputs a forbidden configuration $B''(u) = \psi(e) = B''(v)$, then either $C'(u) \cup C'(v) = \emptyset$ or $\psi(e) \in C'(u) \cap C'(v)$ holds. We will now bound the probability of both events.

Observe that before fixing any random bits, the probability of having a bad radius-$(t-1)$ neighbourhood is the same for all nodes, as all radius-$(t-1)$ node neighbourhoods are identical. Let $S = \Pr[N^{t-1}(u) \text{ is bad}]$ be this probability. By union bound and Lemma 18 we get that

$$\begin{aligned} \Pr[C'(u) \cup C'(v) = \emptyset] &\leq \Pr[C'(u) = \emptyset] + \Pr[C'(v) = \emptyset] \\ &\leq \Pr[N^{t-1}(u) \text{ is bad}] + \Pr[N^{t-1}(v) \text{ is bad}] \\ &\leq 2S. \end{aligned}$$

From Lemma 16 we get that

$$\Pr[\psi(e) \in C'(u) \cap C'(v)] \leq 2L.$$

Using the union bound and the above, we get that the probability of a forbidden configuration is

$$\Pr[B''(u) = \psi(e) = B''(v)] \leq 2S + 2L.$$

To prove the claim, observe that from Definition 17 and the assumption that B' produces a sink at u with probability at most ℓ, it follows that

$$\ell \geq \Pr[u \text{ is a sink}] \geq \Pr[u \text{ is a sink} \mid N^{t-1}(u) \text{ is bad}] \cdot \Pr[N^{t-1}(u) \text{ is bad}] > SL^3.$$

Therefore, $\ell > SL^3$. By setting $L = \ell^{1/4}$ we get that $S < \ell^{1/4}$ implying $2S + 2L < 4\ell^{1/4}$. □

4.3 The Speedup Lemma

The following is an immediate consequence of Lemma 14 and Lemma 19.

Lemma 20. *Suppose B is a sinkless colouring algorithm that runs in time t such that for any edge e the probability that B produces a forbidden configuration at e is at most p. Then there is a sinkless colouring algorithm B'' that runs in $t-1$ rounds such that it produces a forbidden configuration at any edge with probability less than $4 \cdot 6^{1/4} \cdot p^{1/12}$.*

13

2016: state of the art

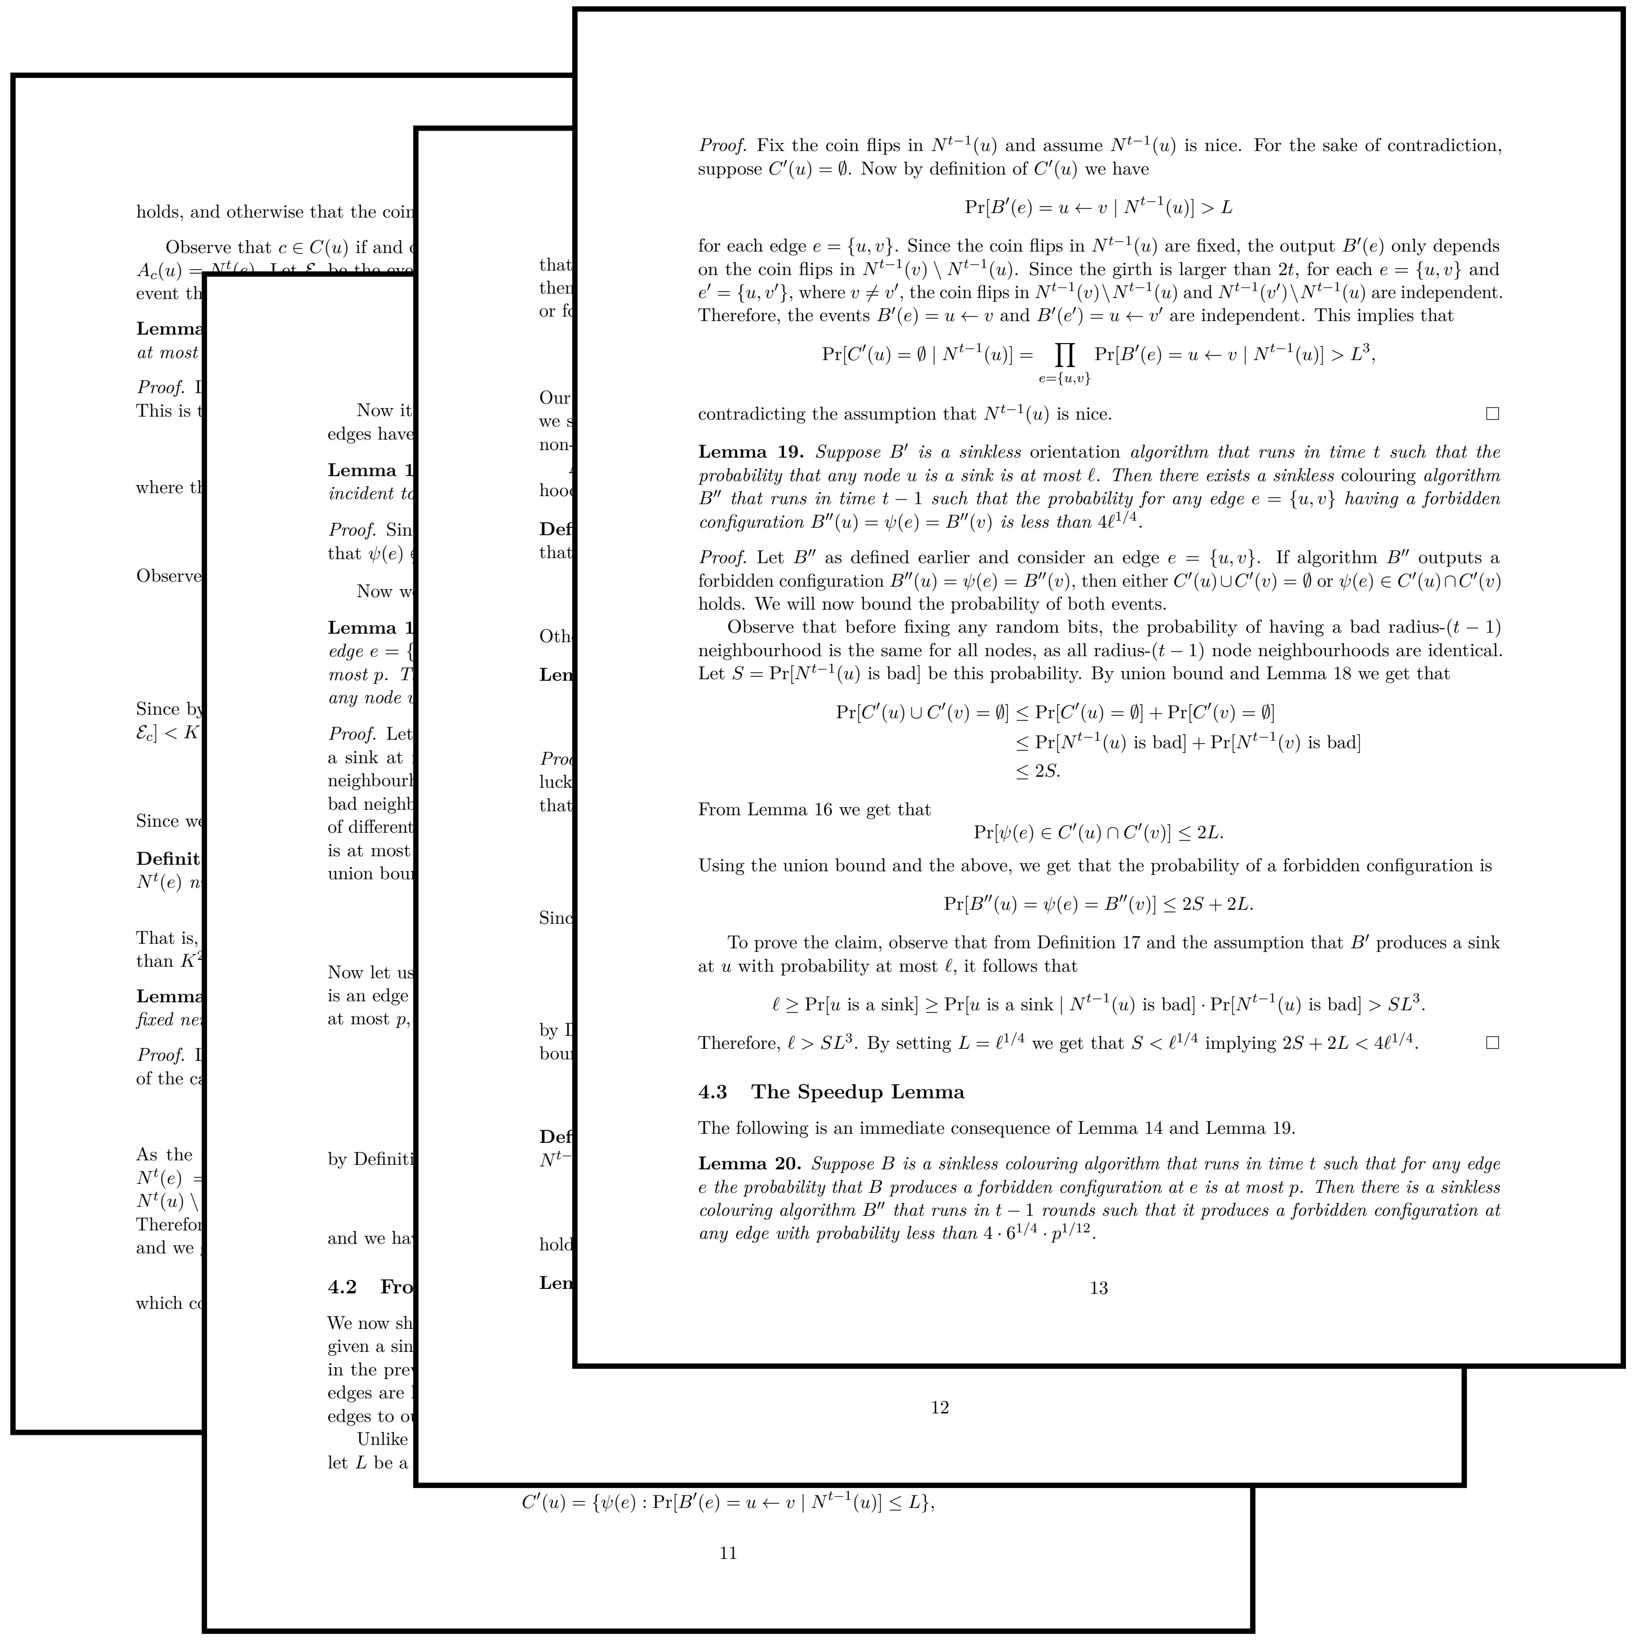

Proof. Fix the coin flips in $N^{t-1}(u)$ and assume $N^{t-1}(u)$ is nice. For the sake of contradiction, suppose $C'(u) = \emptyset$. Now by definition of $C'(u)$ we have

$$\Pr[B'(e) = u \leftarrow v \mid N^{t-1}(u)] > L$$

for each edge $e = \{u, v\}$. Since the coin flips in $N^{t-1}(u)$ are fixed, the output $B'(e)$ only depends on the coin flips in $N^{t-1}(v) \setminus N^{t-1}(u)$. Since the girth is larger than $2t$, for each $e = \{u, v\}$ and $e' = \{u, v'\}$, where $v \neq v'$, the coin flips in $N^{t-1}(v) \setminus N^{t-1}(u)$ and $N^{t-1}(v') \setminus N^{t-1}(u)$ are independent. Therefore, the events $B'(e) = u \leftarrow v$ and $B'(e') = u \leftarrow v'$ are independent. This implies that

$$\Pr[C'(u) = \emptyset \mid N^{t-1}(u)] = \prod_{e=\{u,v\}} \Pr[B'(e) = u \leftarrow v \mid N^{t-1}(u)] > L^3,$$

contradicting the assumption that $N^{t-1}(u)$ is nice. □

Lemma 19. *Suppose B' is a sinkless* orientation *algorithm that runs in time t such that the probability that any node u is a sink is at most ℓ. Then there exists a sinkless* colouring *algorithm B'' that runs in time $t-1$ such that the probability for any edge $e = \{u, v\}$ having a forbidden configuration $B''(u) = \psi(e) = B''(v)$ is less than $4\ell^{1/4}$.*

Proof. Let B'' as defined earlier and consider an edge $e = \{u, v\}$. If algorithm B'' outputs a forbidden configuration $B''(u) = \psi(e) = B''(v)$, then either $C'(u) \cup C'(v) = \emptyset$ or $\psi(e) \in C'(u) \cap C'(v)$ holds. We will now bound the probability of both events.

Observe that before fixing any random bits, the probability of having a bad radius-$(t-1)$ neighbourhood is the same for all nodes, as all radius-$(t-1)$ node neighbourhoods are identical. Let $S = \Pr[N^{t-1}(u) \text{ is bad}]$ be this probability. By union bound and Lemma 18 we get that

$$\begin{aligned} \Pr[C'(u) \cup C'(v) = \emptyset] &\leq \Pr[C'(u) = \emptyset] + \Pr[C'(v) = \emptyset] \\ &\leq \Pr[N^{t-1}(u) \text{ is bad}] + \Pr[N^{t-1}(v) \text{ is bad}] \\ &\leq 2S. \end{aligned}$$

From Lemma 16 we get that

$$\Pr[\psi(e) \in C'(u) \cap C'(v)] \leq 2L.$$

Using the union bound and the above, we get that the probability of a forbidden configuration is

$$\Pr[B''(u) = \psi(e) = B''(v)] \leq 2S + 2L.$$

To prove the claim, observe that from Definition 17 and the assumption that B' produces a sink at u with probability at most ℓ, it follows that

$$\ell \geq \Pr[u \text{ is a sink}] \geq \Pr[u \text{ is a sink} \mid N^{t-1}(u) \text{ is bad}] \cdot \Pr[N^{t-1}(u) \text{ is bad}] > SL^3.$$

Therefore, $\ell > SL^3$. By setting $L = \ell^{1/4}$ we get that $S < \ell^{1/4}$ implying $2S + 2L < 4\ell^{1/4}$. □

4.3 The Speedup Lemma

The following is an immediate consequence of Lemma 14 and Lemma 19.

Lemma 20. *Suppose B is a sinkless colouring algorithm that runs in time t such that for any edge e the probability that B produces a forbidden configuration at e is at most p. Then there is a sinkless colouring algorithm B'' that runs in $t-1$ rounds such that it produces a forbidden configuration at any edge with probability less than $4 \cdot 6^{1/4} \cdot p^{1/12}$.*

13

2019: state of the art

O O O O
O O O I
O O I I
O I I I

I O

speedup

intermediate complexity, **Ω(log n)**