Locality in Networks

Jukka Suomela
Helsinki Institute for Information Technology HIIT
Department of Computer Science, University of Helsinki

Foundations of Network Science Workshop
Riga, 7 July 2013

1. Locality and Local Algorithms

– brief introduction

Locality in Networks

- **Basic setting:**
  - nodes act based on **local information only**
  - behaviour of node v = function of information available in $O(1)$-radius neighbourhood of v
- **Question:**
  - what tasks can be solved?

Constant-Radius Neighbourhood

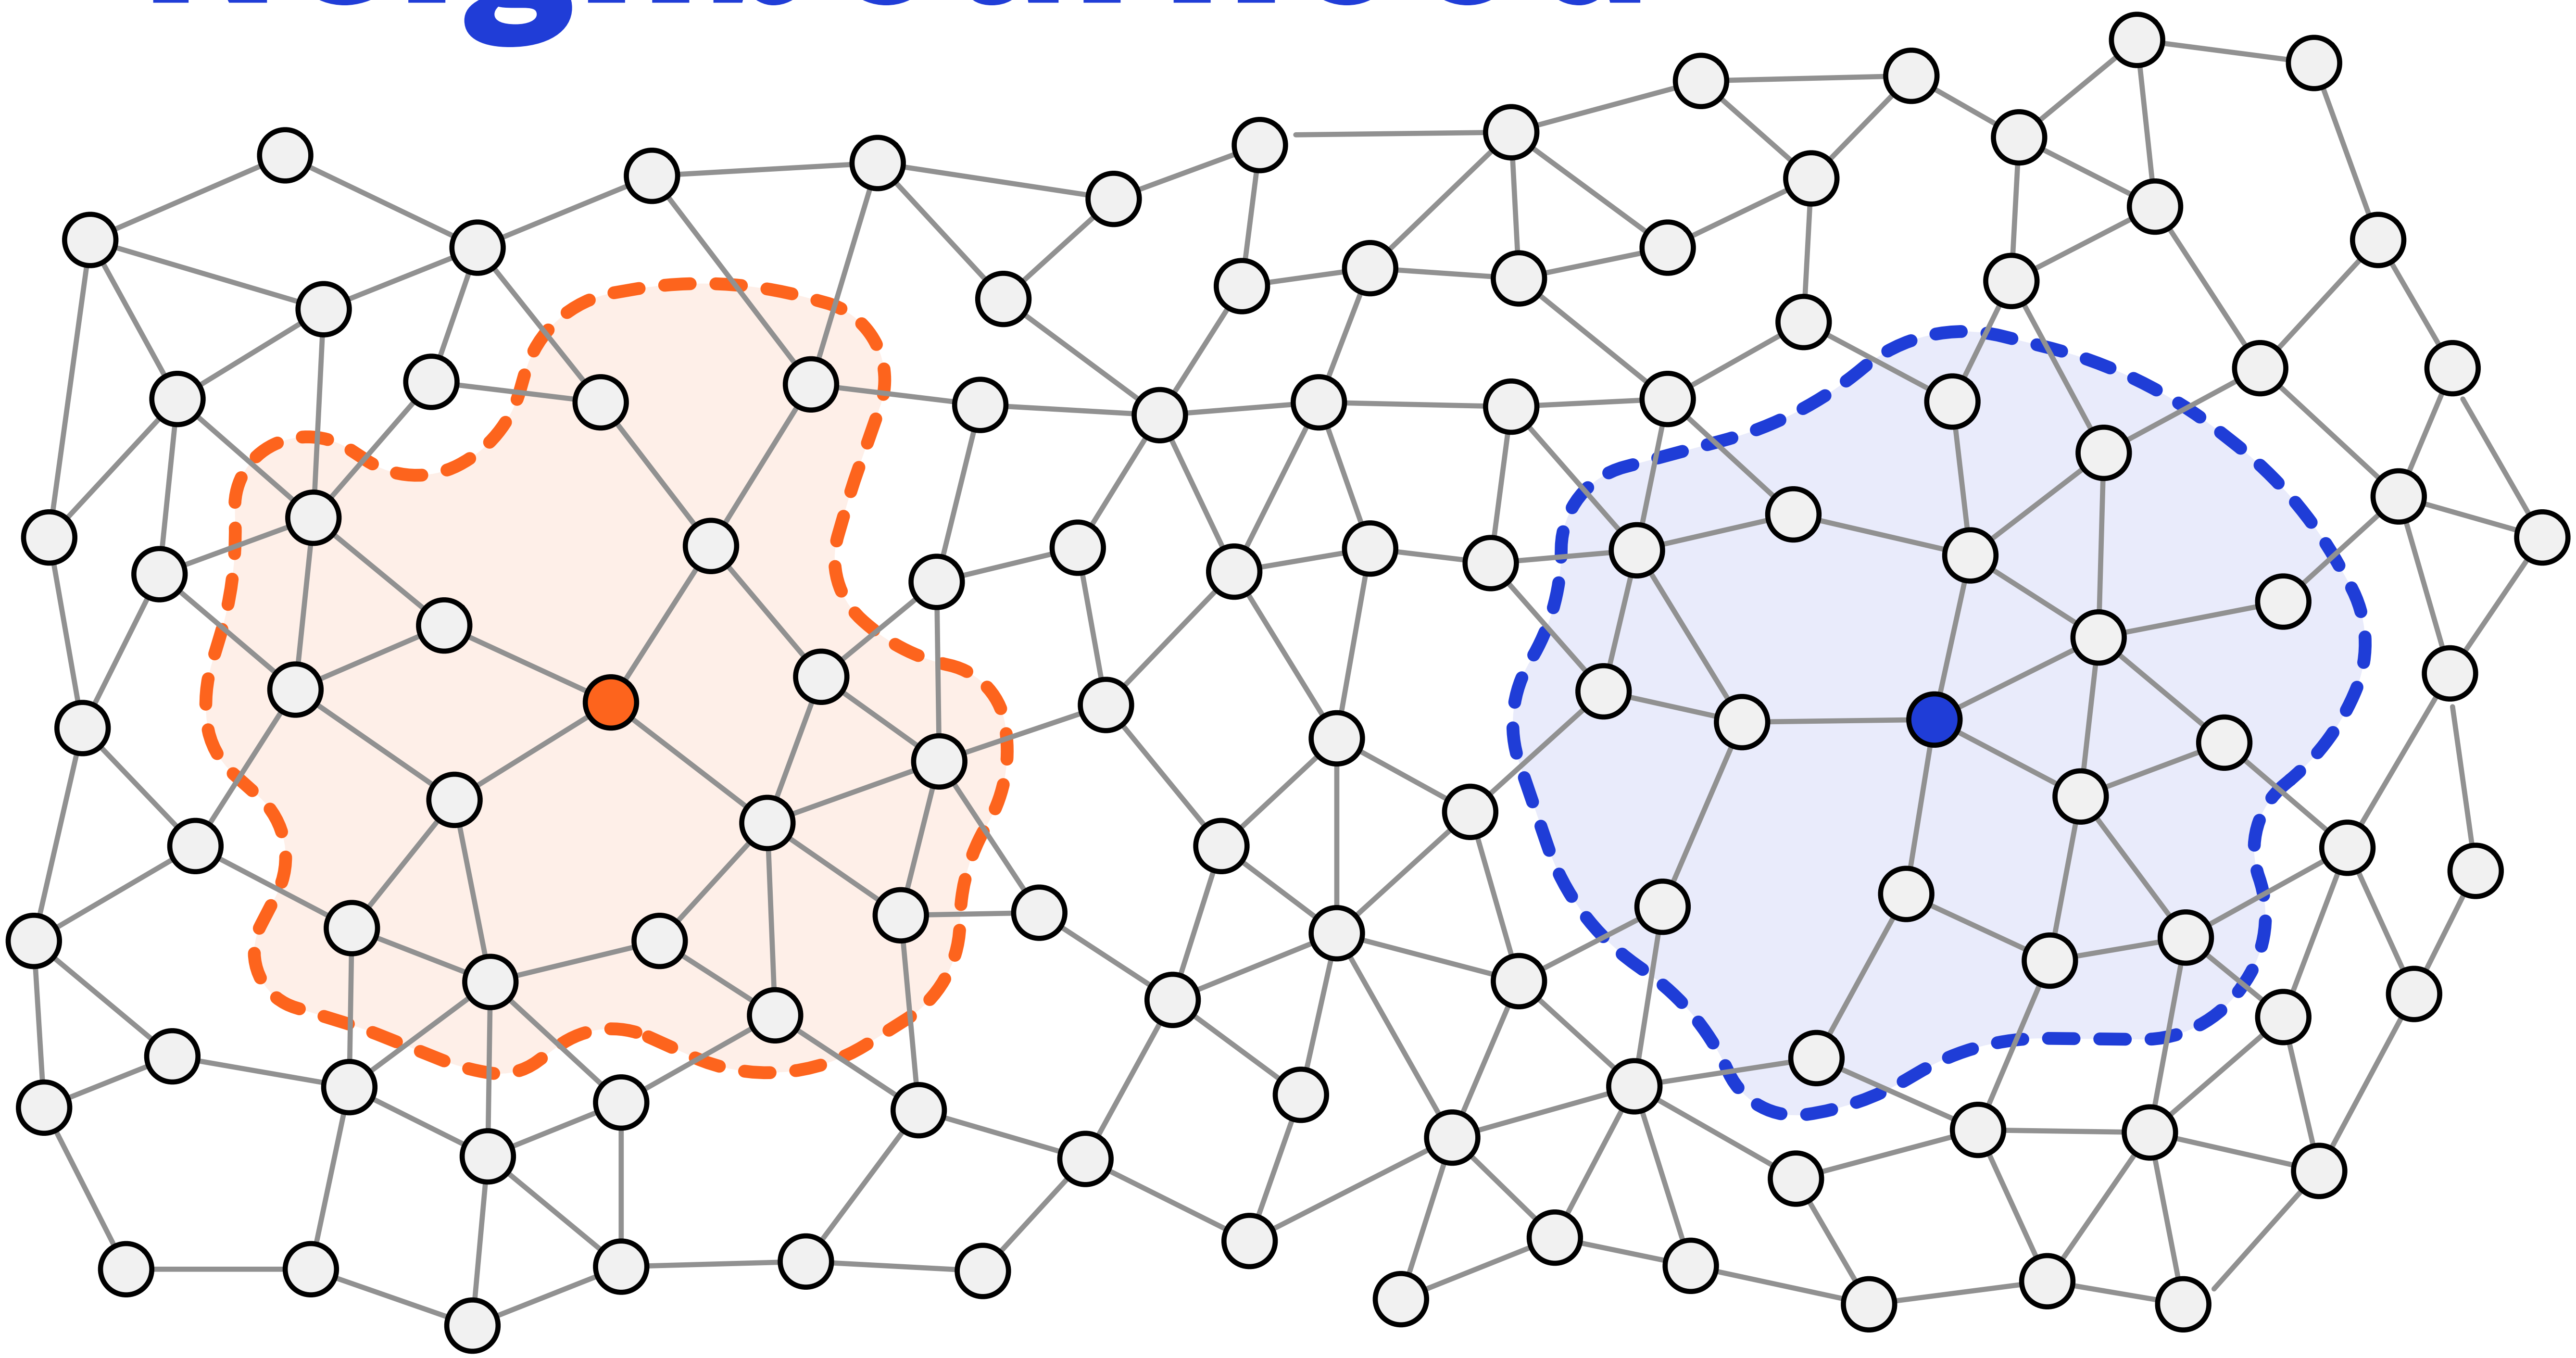

Example:
Matching in Networks

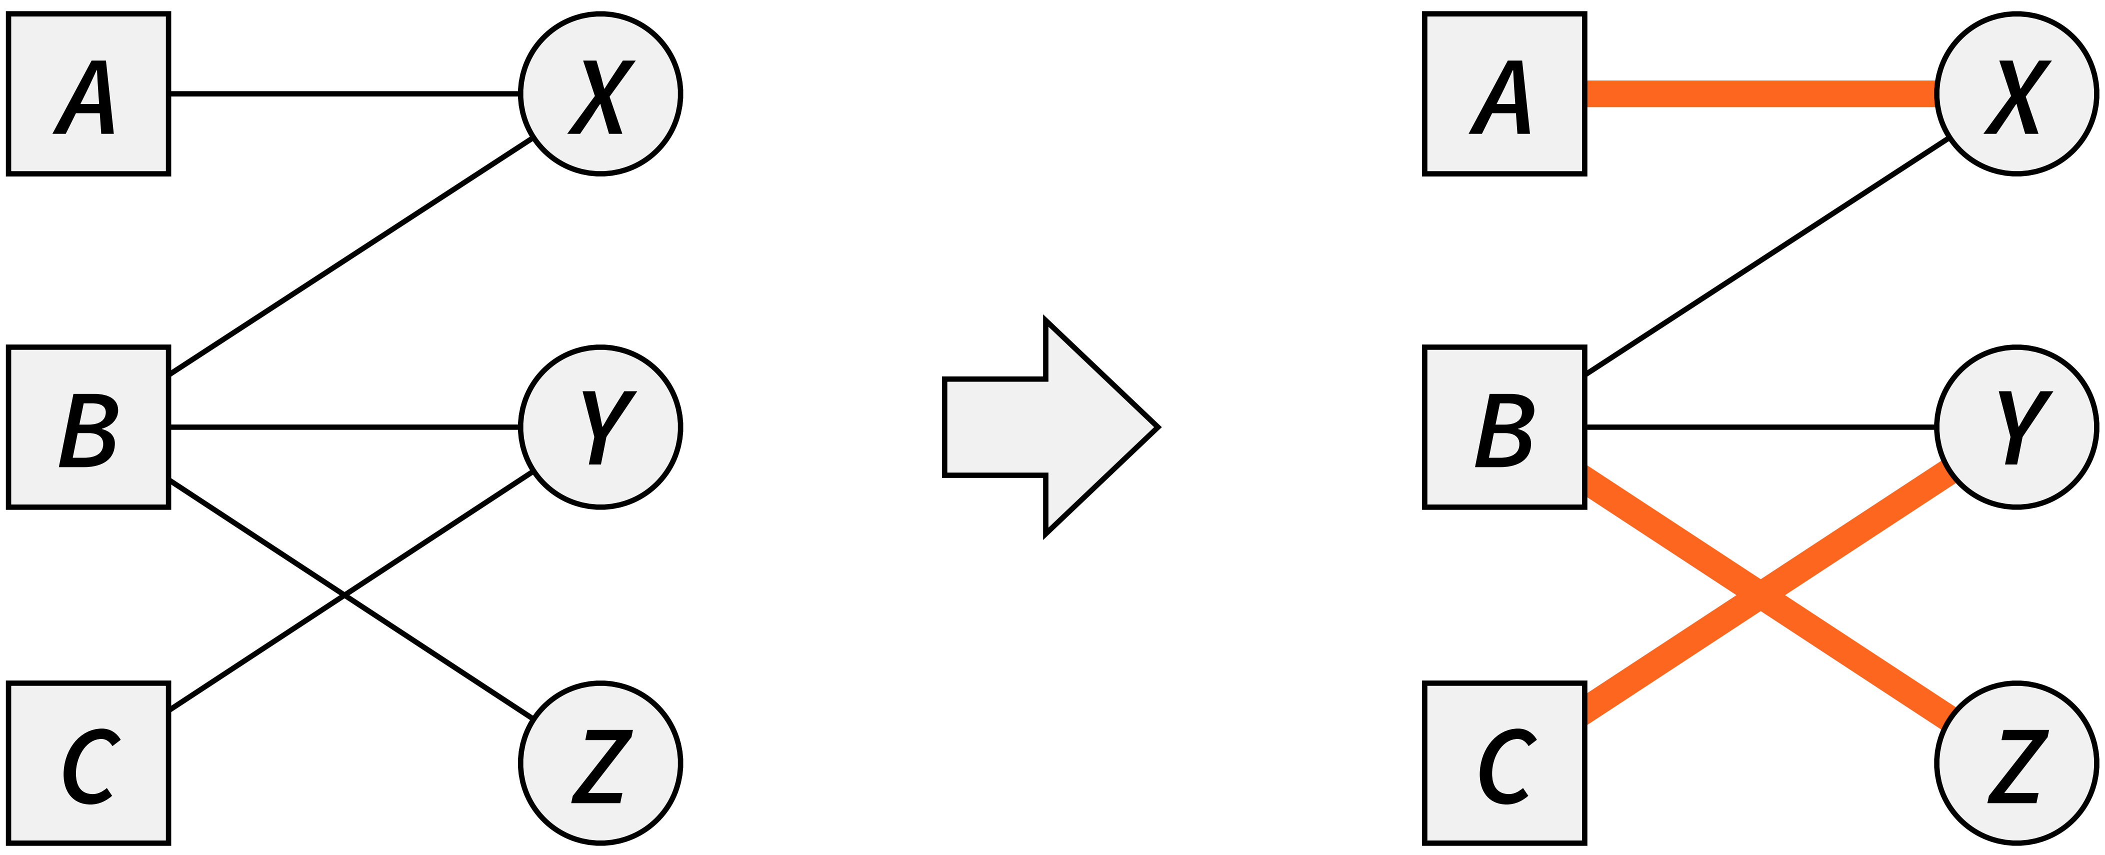

Example:
Matching in Networks

- **Job markets:** open positions and workers
- **Economics:** buyers and sellers
- **Social networks:** marriages
- **Computer networks:** resource allocation

Local Algorithms for Matching in Networks

- **Local perspective:**
  - each player decides with whom to pair based on its local neighbourhood
- **Global perspective:**
  - globally consistent solution, good solution (e.g., large matching)

Maximum Matchings

- **Largest possible number of pairs**

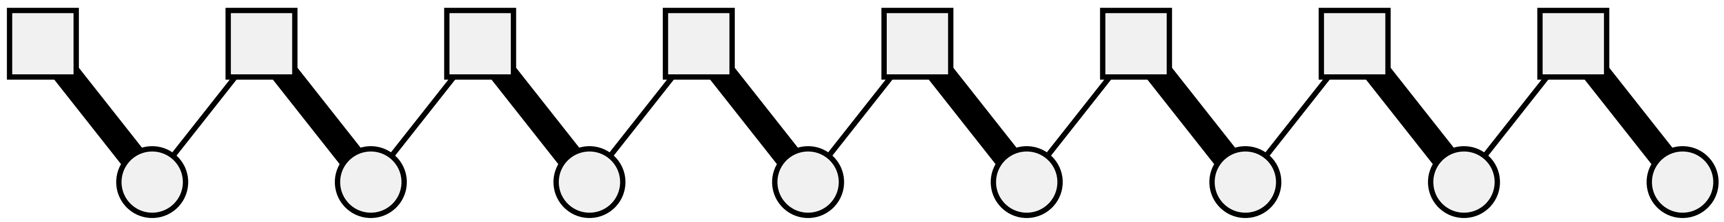

Maximum Matchings

- **No local algorithm — simple proof:**

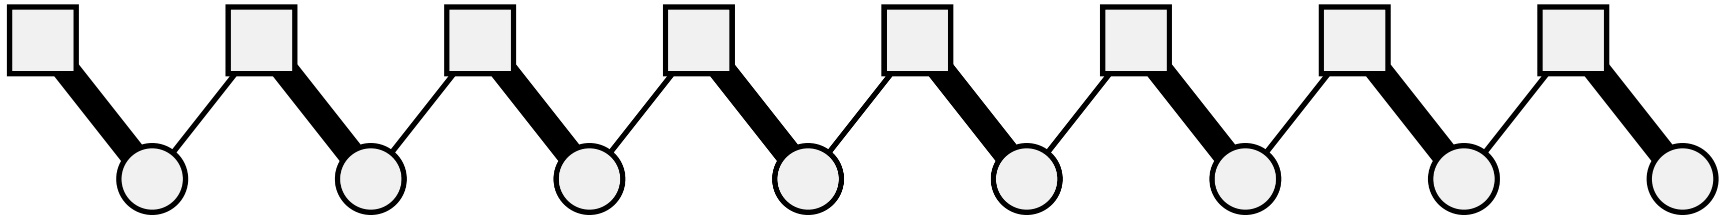

vs.

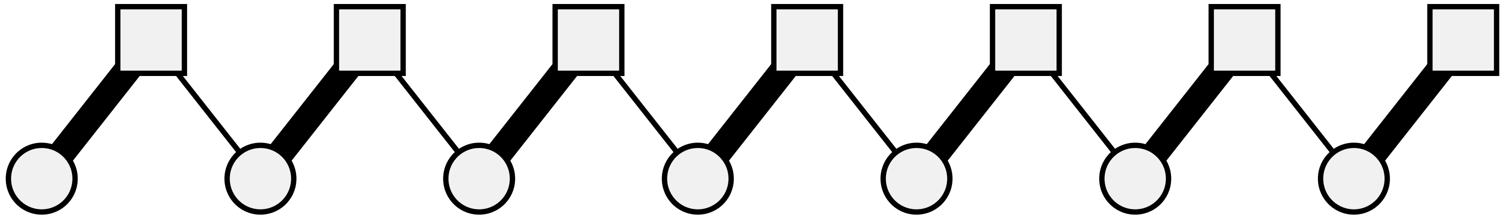

Maximum Matchings

- **Same neighbourhood, different output**

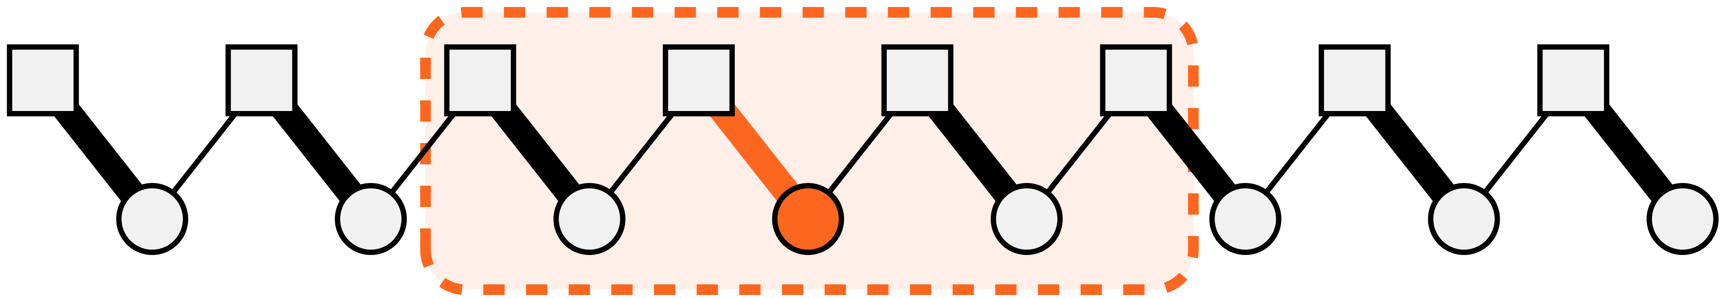

vs.

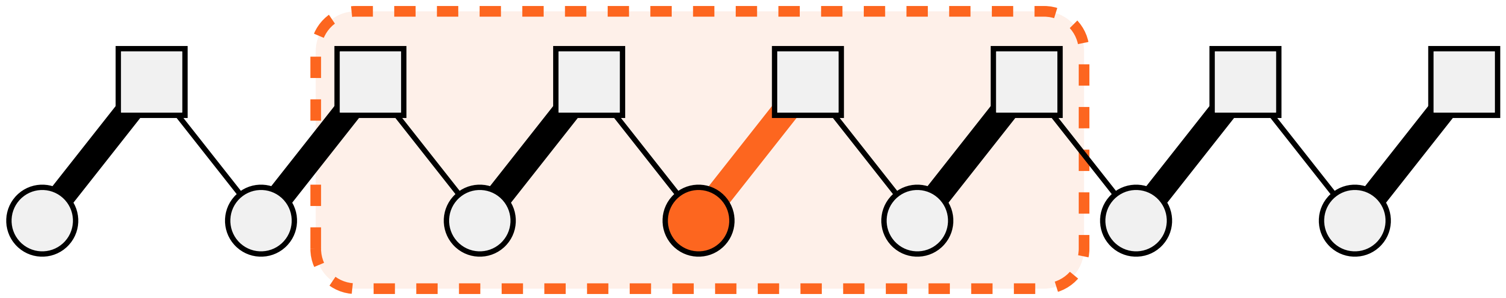

Approximations of **Maximum Matchings**

- **No local algorithm for maximum matching**
- **However, we can find arbitrarily good approximations locally**
  - identify & eliminate all **short** augmenting paths, in parallel
  - local, if maximum degree $O(1)$

Stable Matchings

- **No pair of nodes has incentive to change**
  - *X* prefers *B* to *A*, *B* prefers *X* to *Y*

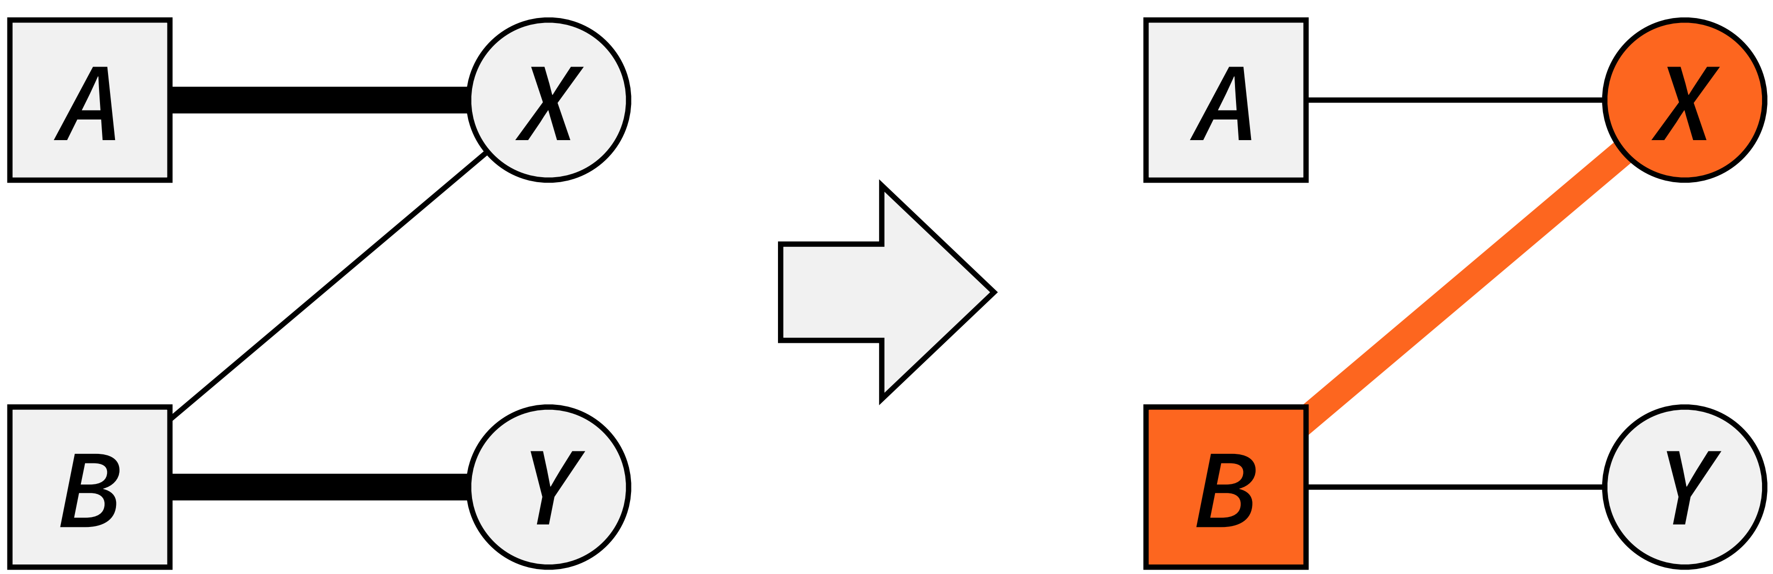

Stable Matchings

- **No pair of nodes has incentive to change**
- **Not possible with local behaviour**
  - long path
  - preferences near endpoints determine what we must do near midpoint

Stable Matchings

- **No pair of nodes has incentive to change**
- **Not possible with local behaviour**
- **Possible if we tolerate a small fraction of unstable edges**
  - simple and natural local algorithm…

Almost Stable Matchings

- **Truncated Gale–Shapley algorithm**
  - currently unmatched “men” propose women in preference order
  - “women” accept the best proposal so far
  - run for $O(1)$ parallel rounds — **local**
- **Few unstable edges** (if low degrees)

Local Algorithms

- **Active subfield of distributed computing**
  - Linial (1992):
    "*Locality in distributed graph algorithms*"
  - Naor & Stockmeyer (1995):
    "*What can be computed locally*"
  - Kuhn, Moscibroda, Wattenhofer (2004):
    "*What cannot be computed locally*"

Local Algorithms for Graph Problems

- **Lots of good approximations — at least in some special cases:**
  - matchings, dominating sets, edge covers, vertex covers, **packing/covering linear programs**, …
- **More details:** "*Survey of local algorithms*" (2013)

2. Network Science Perspective

– reasons to expect locality

– implications

Why Local?

- **Attractive in computer networks**
  - fast, fault-tolerant, robust
  - cheap and simple
  - easy to design, easy to implement
- **What about social networks, markets, biological systems, industrial systems…?**

Why Expect Locality?

- **Privacy, competition, selfishness**
  - why would strangers reveal what they know?
  - why would our competitors do it?
- **Timeliness**
  - distant information is likely outdated,
    so why care about it at all?

Why Expect Locality?

- **Simple and unreliable communication**
  - how to encode lots of data in a mixture of some chemical compounds?
- **Simple entities, limited capabilities**
  - could I keep track of friends of friends of friends?

Implications

- **Distributed systems:**
  - **upper-bound results** are of practical use
  - algorithms that we can implement and run
- **Network science:**
  - **lower-bound results** are of practical use?
  - learn about *possible* behaviour in networks

Locality Lower Bounds: **Predictions**

- **No good matchings in real-world networks**
  - open positions *and* unemployed people
- **No optimal resource allocation**
- **Even if everyone does its best to co-operate!**
  - not price of anarchy but **price of locality**

3. Understanding Locality Lower Bounds

– why are some tasks non-local?

Reasons for Non-Locality

- **Common theme:**
  - nodes *u* and *v* have identical local neighbourhoods
  - nodes *u* and *v* should make different decisions

Reasons for Non-Locality

- **Example: maximum matching**

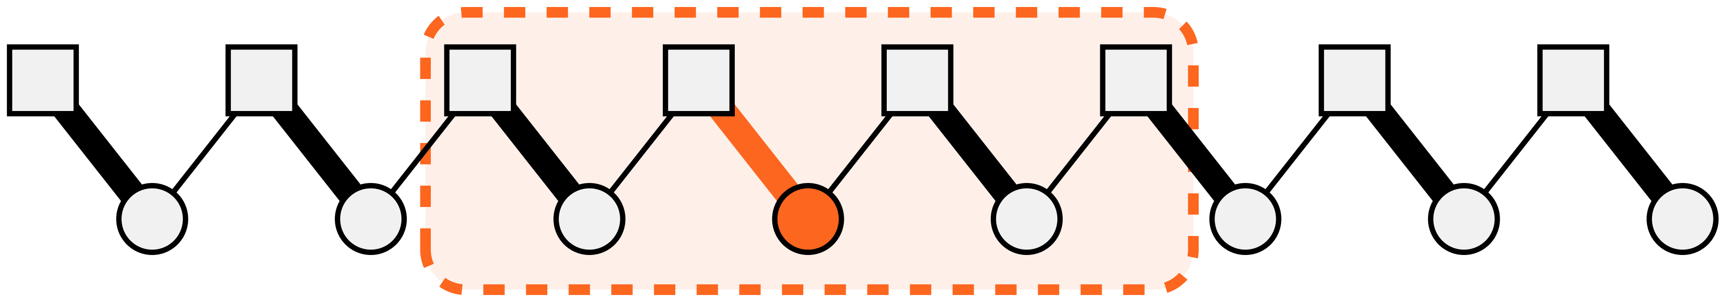

vs.

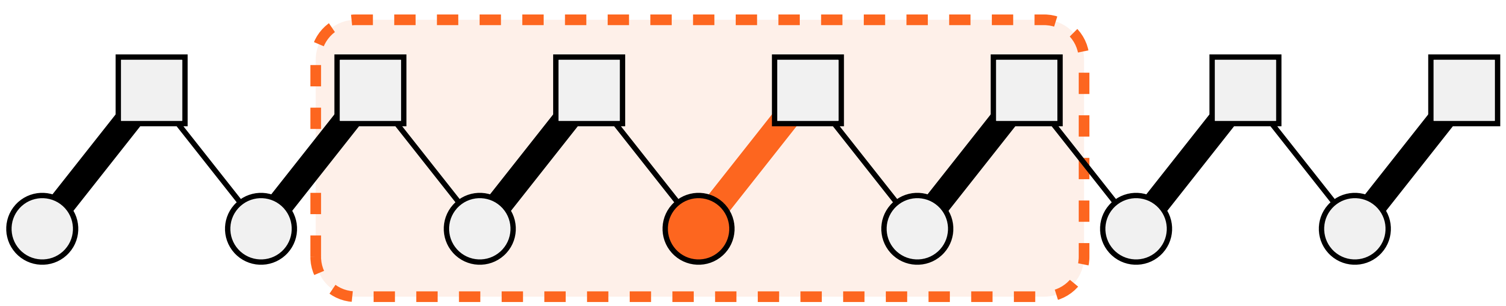

Reasons for Non-Locality

- **Example: graph colouring**

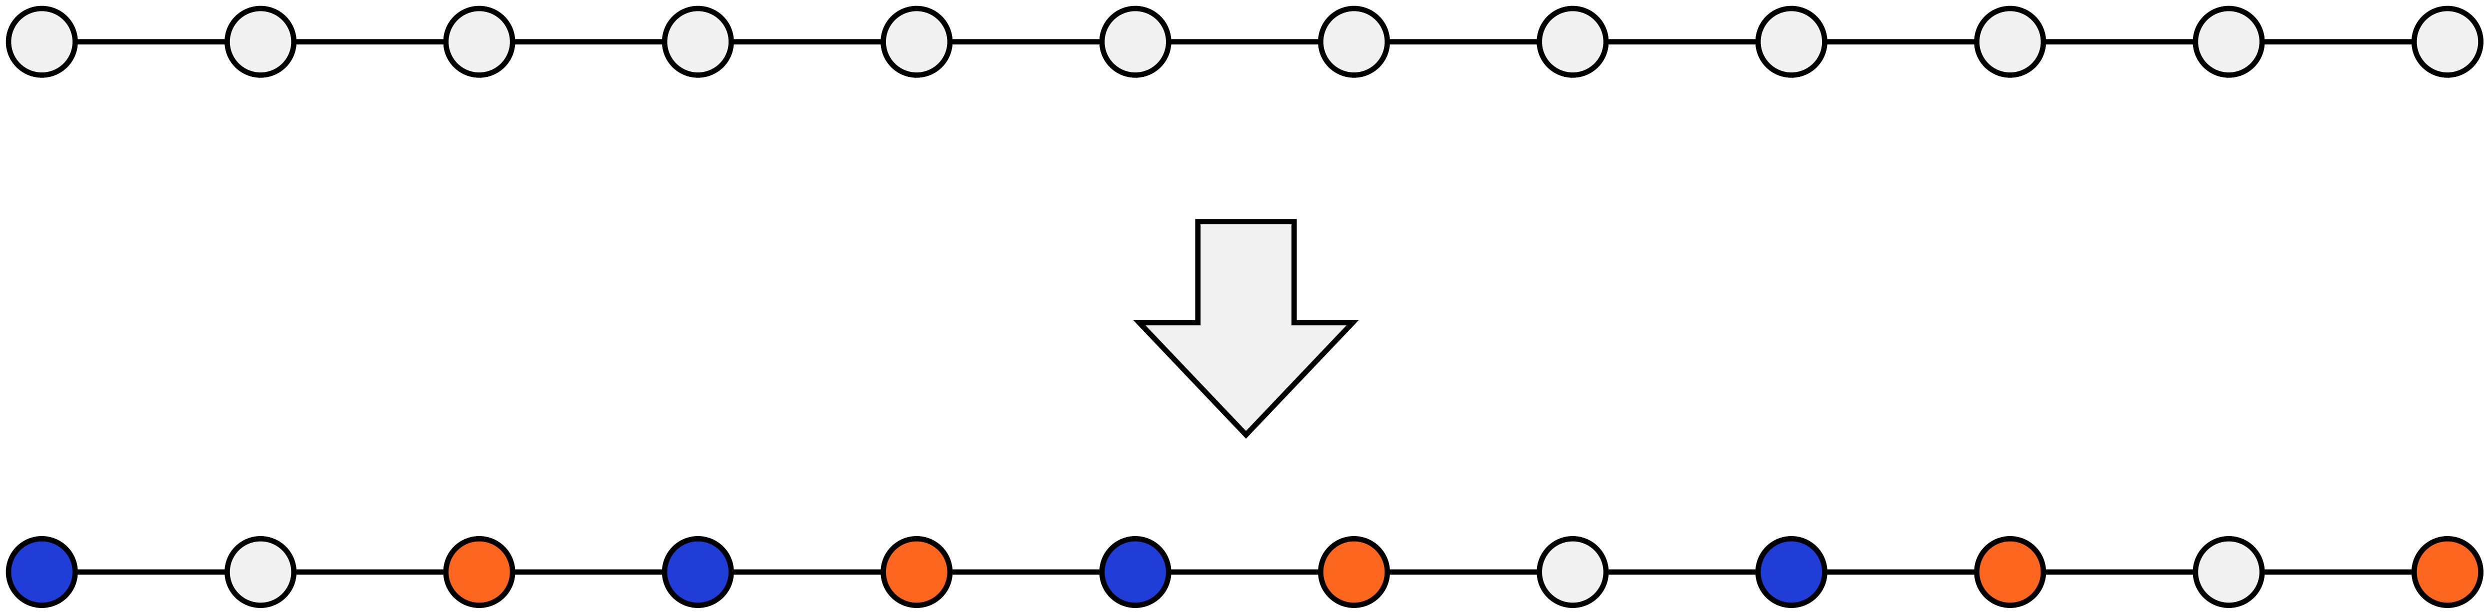

Reasons for Non-Locality

- **Example: graph colouring**

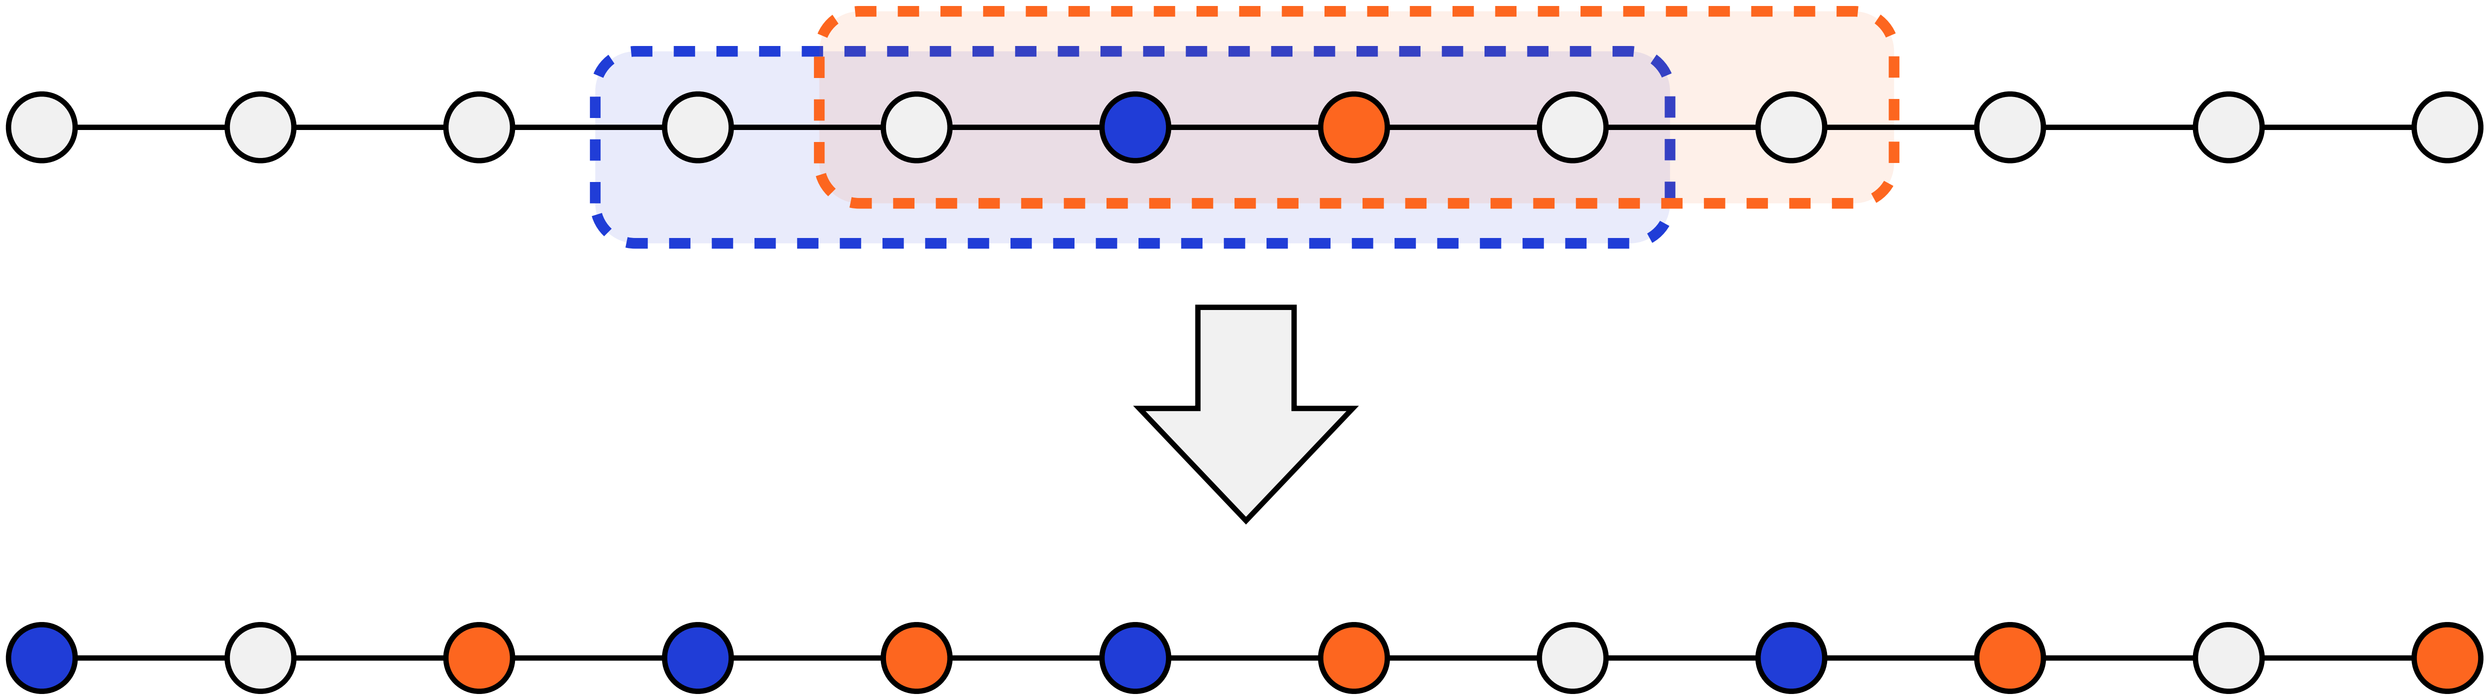

Reasons for Non-Locality

- **Maximum matching:**
  - global optimum needs global information
- **Graph colouring:**
  - extra information needed to break symmetry
- **But there are also less obvious reasons…**

Example:
Large Cuts

- **Label nodes with orange/blue**
- **Cut edge: endpoints with different colours**

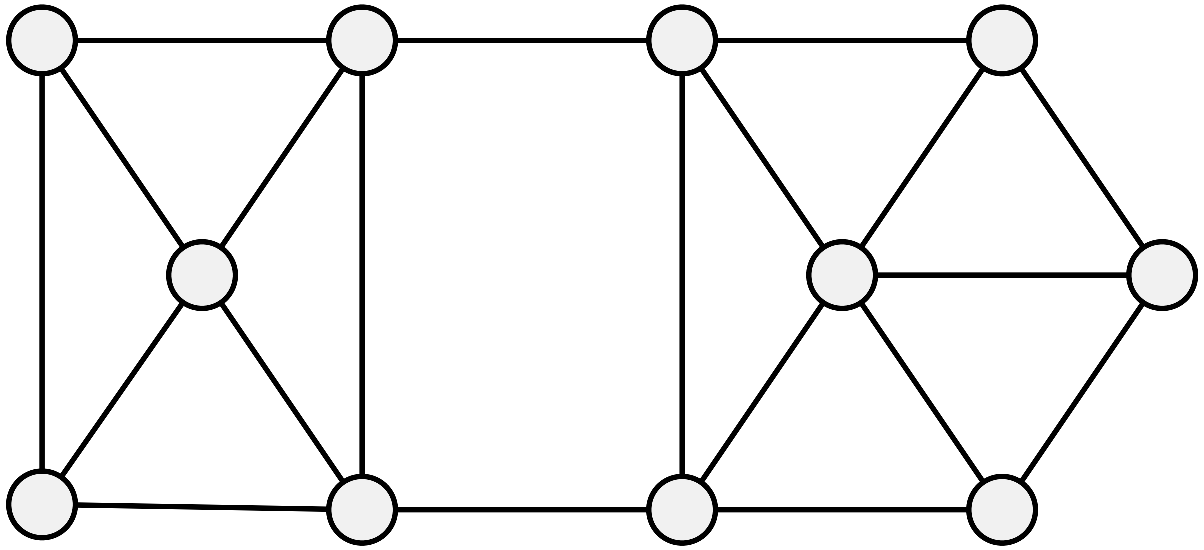

Example:
Large Cuts

- **Label nodes with orange/blue**
- **Cut edge: endpoints with different colours**

Bad solution:

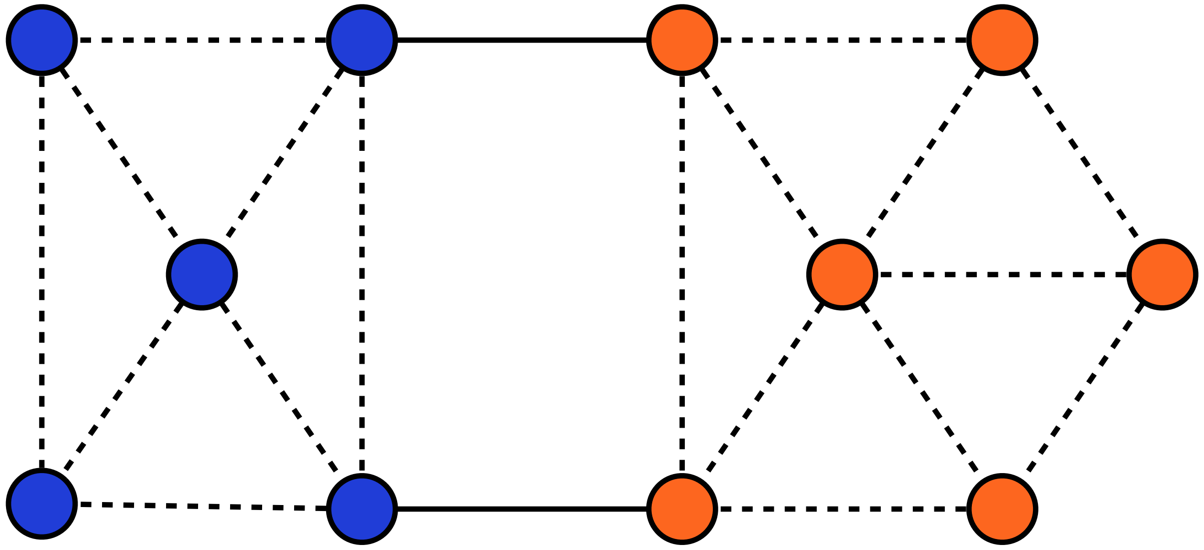

Example:
Large Cuts

- **Label nodes with orange/blue**
- **Cut edge: endpoints with different colours**

Good solution:

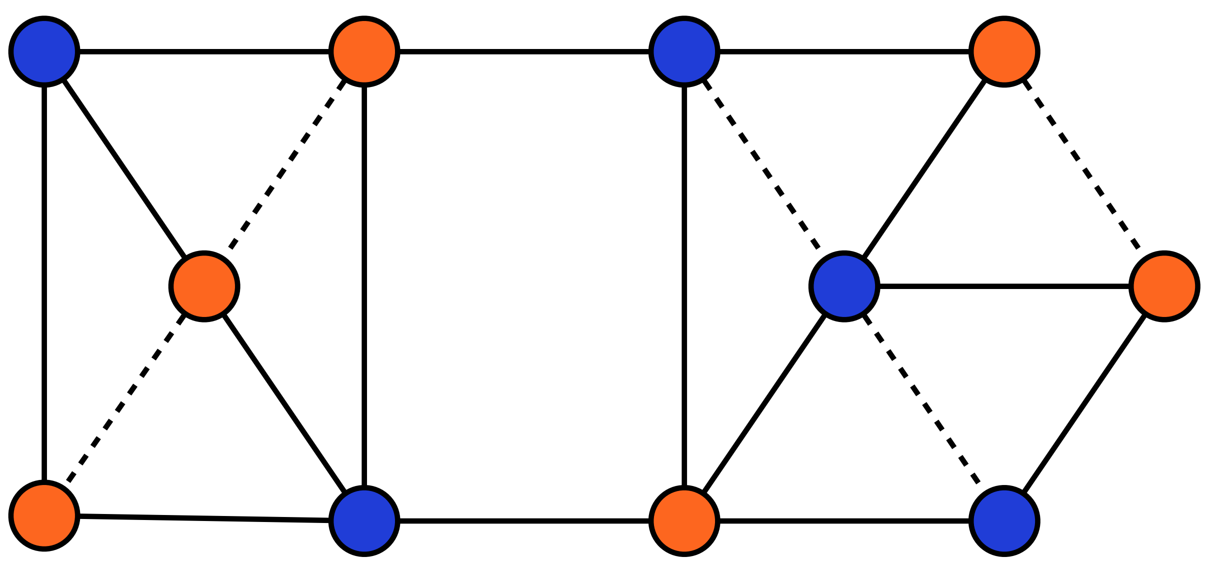

Example: **Large Cuts**

- **Simple local rule: flip coins to pick labels**
  - in expectation 1/2 of all edges are good
  - trivial 1/2-approximation
- **Can we do better?**
  - what if we looked further?
  - what if we used more random bits?

Example:
Large Cuts

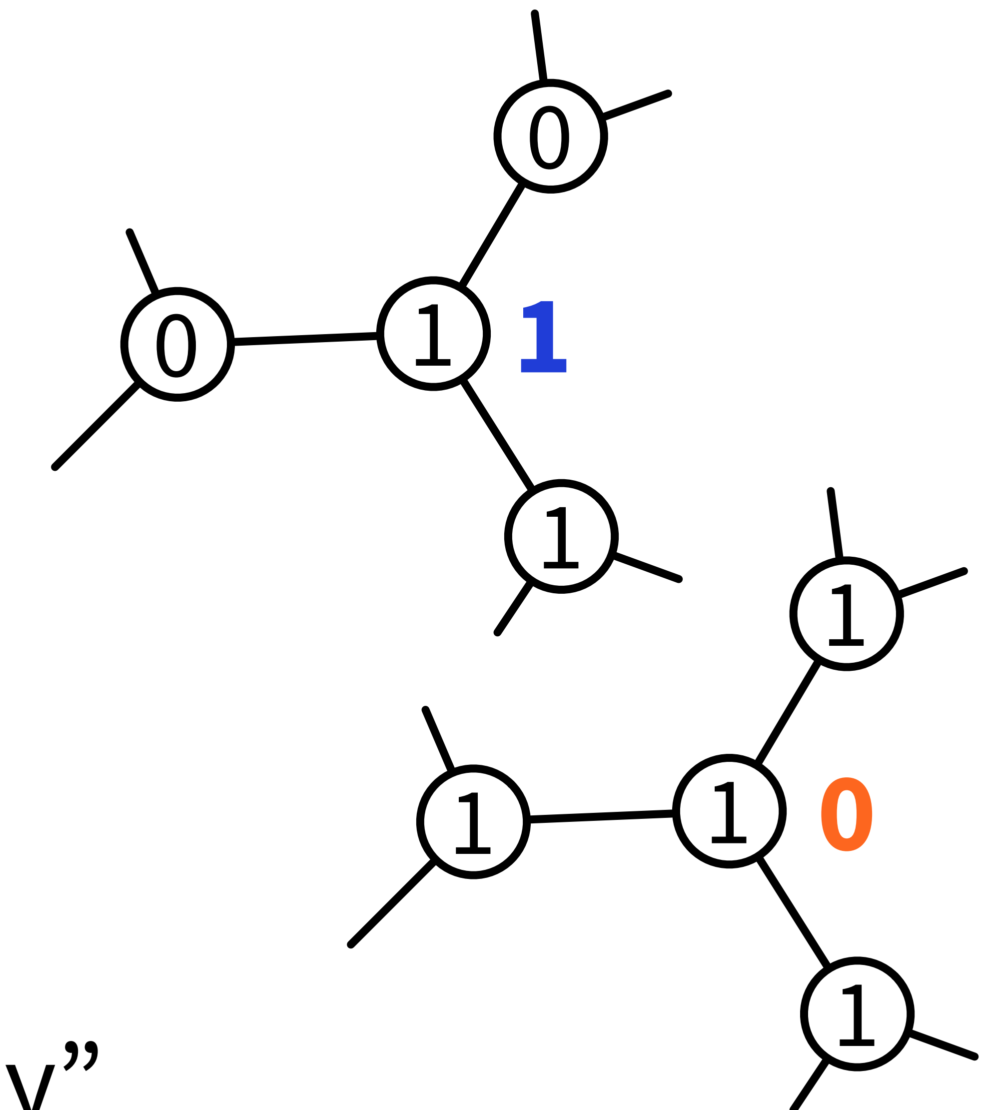

- **We can do *slightly* better:**
  - flip coins
  - **change mind** if "too many" neighbours with the same random bit
  - *d*-regular triangle-free graphs: $1/2 + \Theta\left(1/\sqrt{d}\right)$
- **Best possible approximation ratio — why?**

Lower Bound: **Large Cuts**

- **Networks X and Y look locally identical:**
  - X has large cuts, Y does not have large cuts
- **Local algorithm A: same behaviour in X and Y**
  - must produce small cuts in Y
  - therefore produces small cuts in X, too
  - poor approximation ratio in X

Lower Bound:
Large Cuts

- ***Y* = non-bipartite Ramanujan graphs**
  - high girth — looks locally like a tree
  - no large cuts (**spectral properties**)
- ***X* = bipartite double cover of *Y***
  - looks locally identical to *Y*
  - has a large cut (**bipartite**)

Bipartite Double Cover

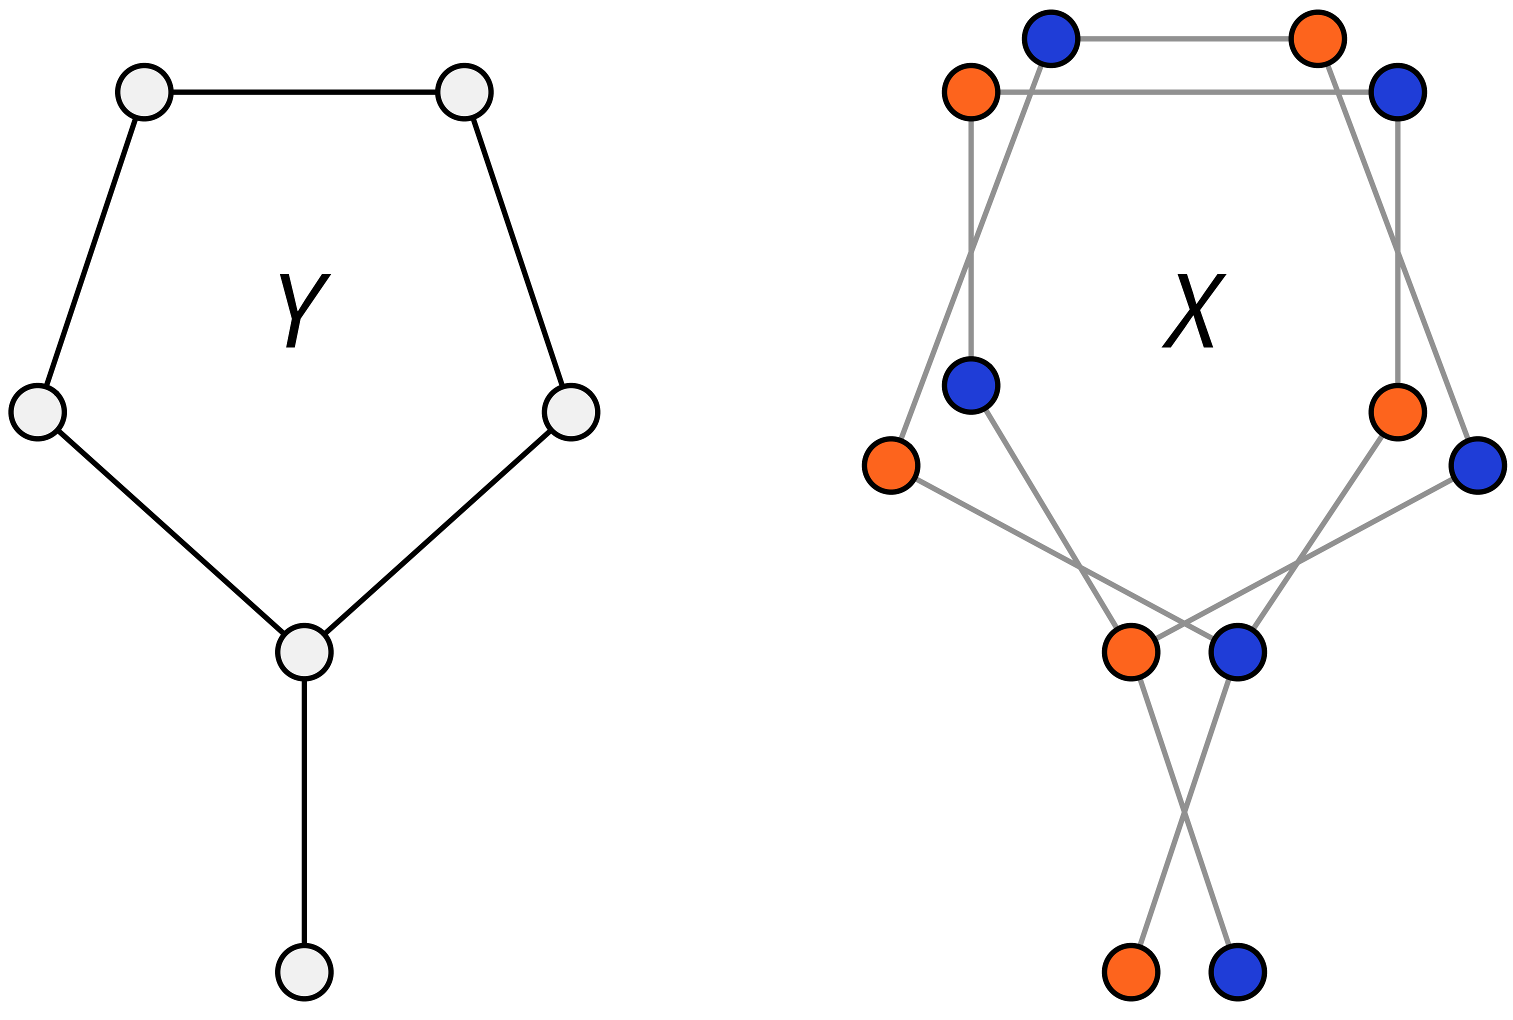

Identical Local Neighbourhoods

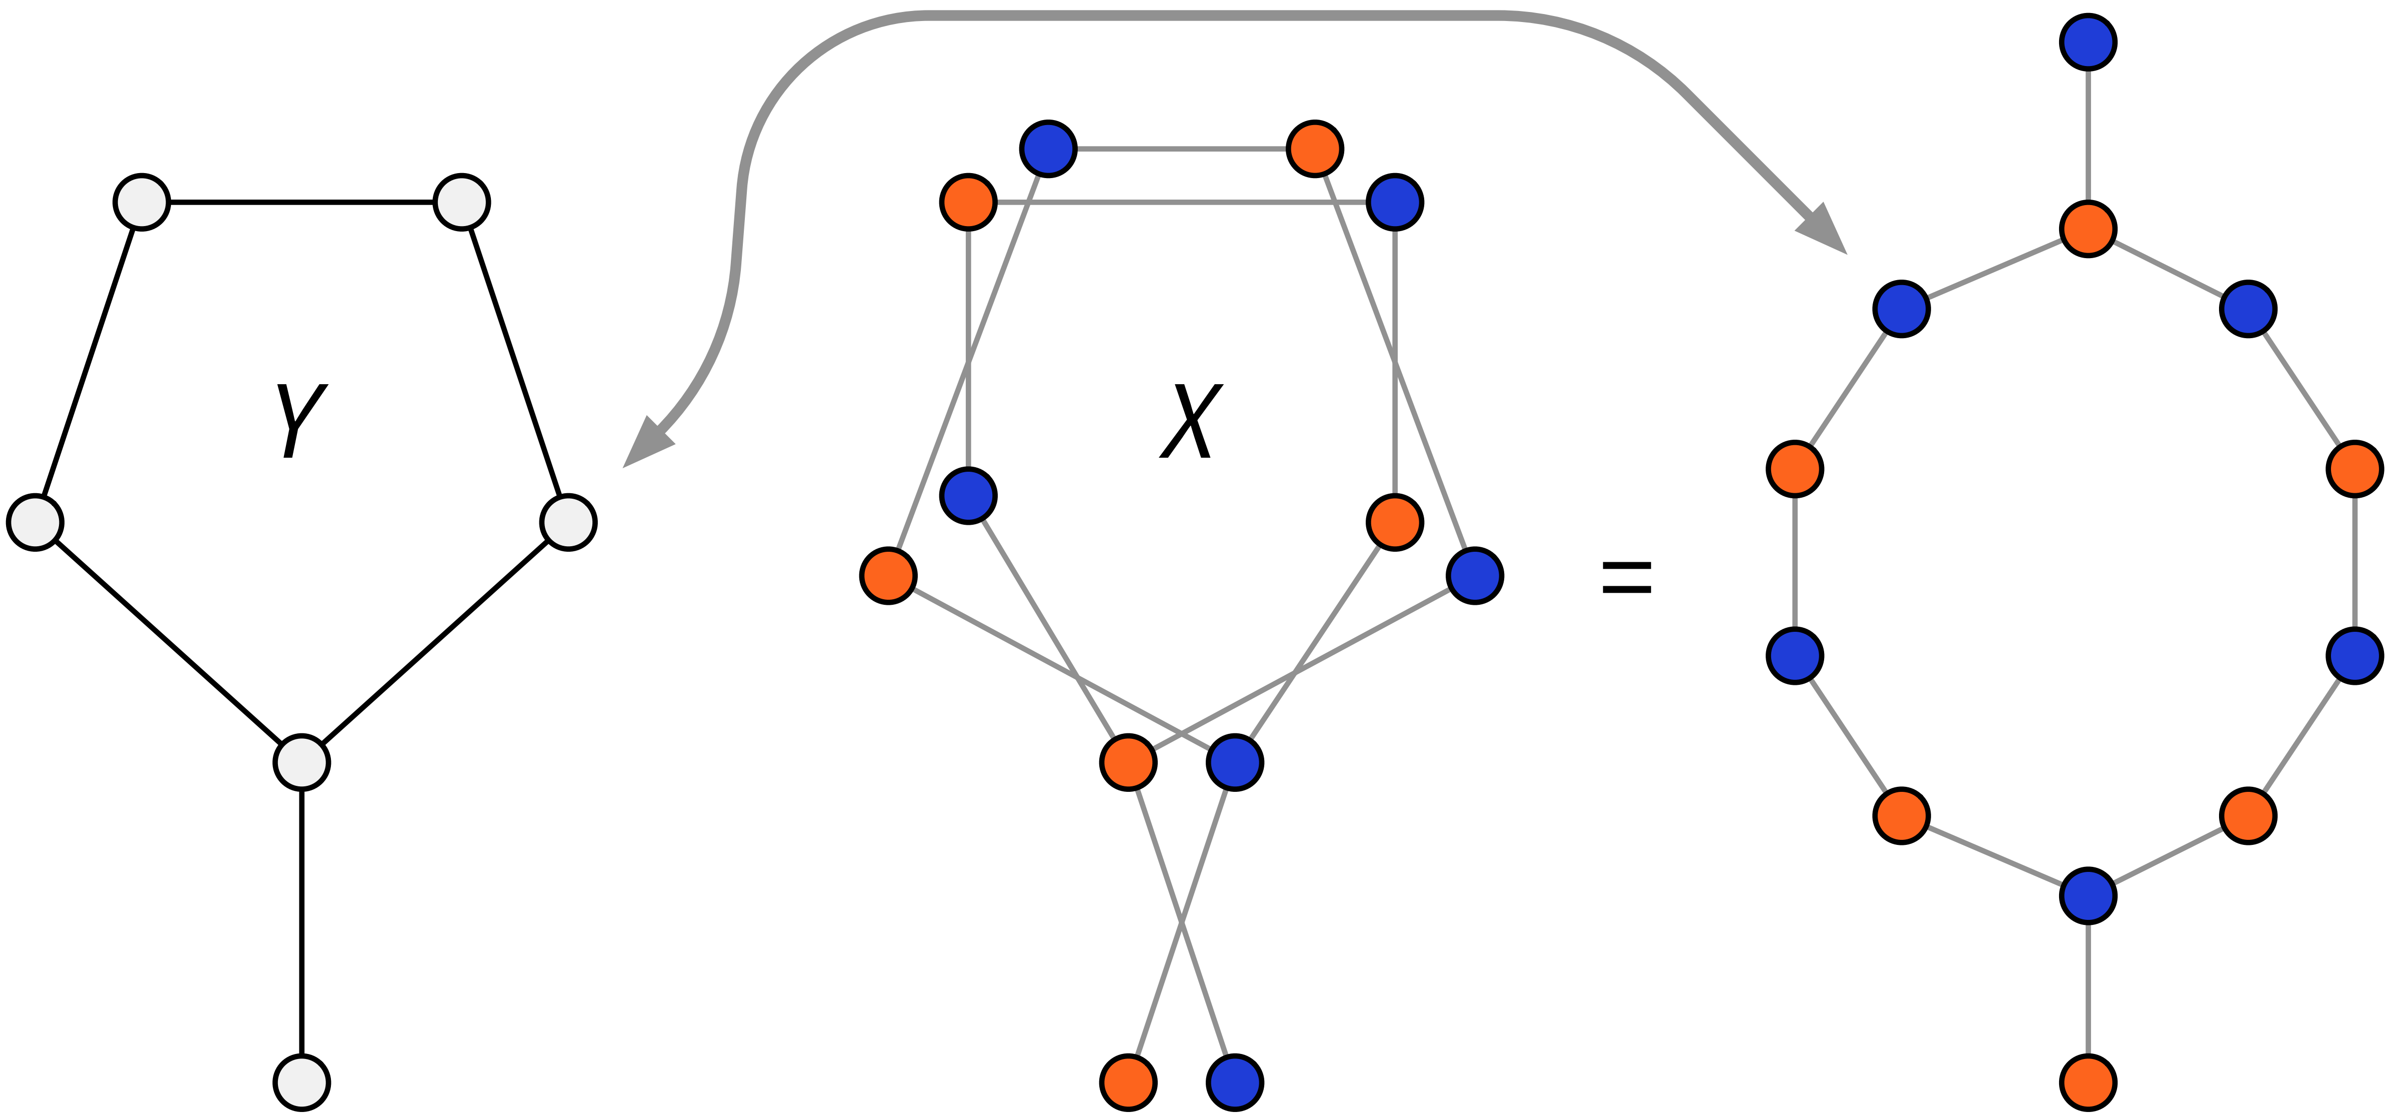

Identical Local Neighbourhoods

- **Edge *e* in *Y* — similar edges e_1 and e_2 in *X***

- **Pr[** edge e_1 in X is a cut edge **] =**
  Pr[edge e_2 in X is a cut edge **] =**
  Pr[edge e in Y is a cut edge **]**

- **E[** fraction of cut edges in X **] =**
  E[fraction of cut edges in Y **]**

Reasons for Non-Locality

- **Similar techniques work for many problems**
  - find a bad counterexample Y
  - construct an "easy" instance X
  - make sure X and Y look locally identical
  - local algorithm: similar behaviour in X and Y
  - poor approximation in X

Typical Counterexamples

- **Regular graph**
  - node degrees do not help
- **High girth**
  - locally looks like a regular tree
- **Expander graphs**

4. But What About More Realistic Networks?

– do locality lower bounds tell us anything about “typical” networks?

Locality in Real-World Networks

- **Local algorithms for "nice" graph families?**
- **Some progress:**
  - **bounded degrees**
  - bounded growth, bounded independence…
  - bounded arboricity, forbidden minors…
  - line graphs, planar graphs…

Locality in Real-World Networks

- **Distributed computing community focuses on graph families that look like "typical computer networks"**
  - bounded degrees ≈ wired networks
  - bounded growth ≈ wireless networks

Locality in Real-World Networks

- **Distributed computing community focuses on graph families that look like "typical computer networks"**
- **What about job markets, biological networks, social networks, ...?**
  - need to re-think the assumptions

5. Next Steps

– towards tight results in relevant graph families

Research Agenda: **Next Steps**

- **Radius of locality r vs. parameters of network family**
- **State of the art: r vs. maximum degree Δ**
  - $r = \Theta(1)$ — approximations of max-cut
  - $r = \Theta(\operatorname{polylog} \Delta)$ — approximations of LPs
  - $r = \Theta(\Delta)$ — maximal solutions to LPs

Research Agenda: **Next Steps**

- **Radius of locality *r* vs. parameters of network family**
- **Maximum degree:**
  - wrong parameter for social networks
  - tight bounds on *r* in networks with a small number of high-degree nodes?

Summary:
Locality in Networks

– how to go beyond
the traditional scope
of ***computer*** networks?