Compact Distance Oracles with Large Sensitivity and Low Stretch

Davide Bilò, Keerti Choudhary, Sarel Cohen, Tobias Friedrich, **Simon Krogmann**, Martin Schirneck

UNIVERSITÀ DEGLI STUDI DELL'AQUILA

भारतीय प्रौद्योगिकी संस्थान दिल्ली
Indian Institute of Technology Delhi

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

HPI Hasso Plattner Institut

universität wien

The Problem

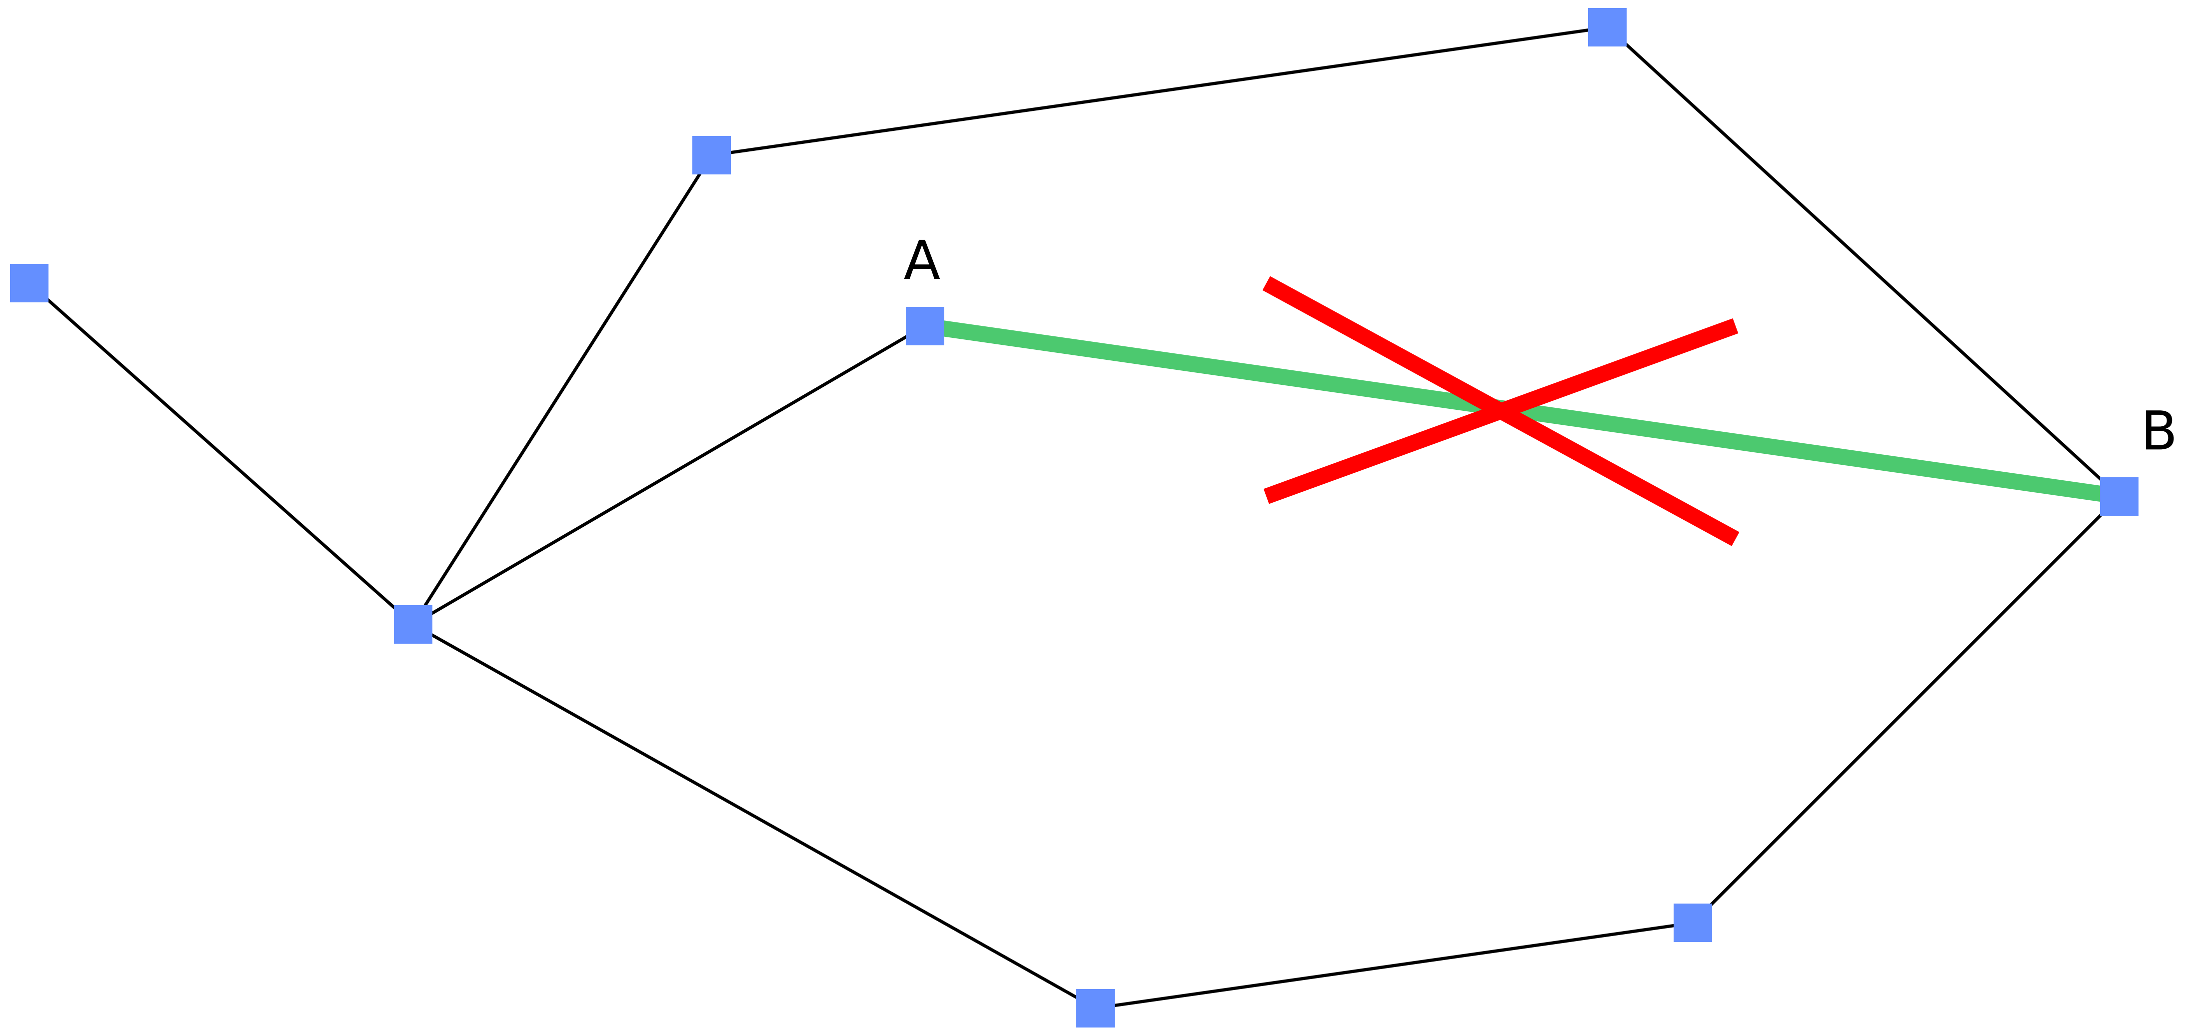

The Problem

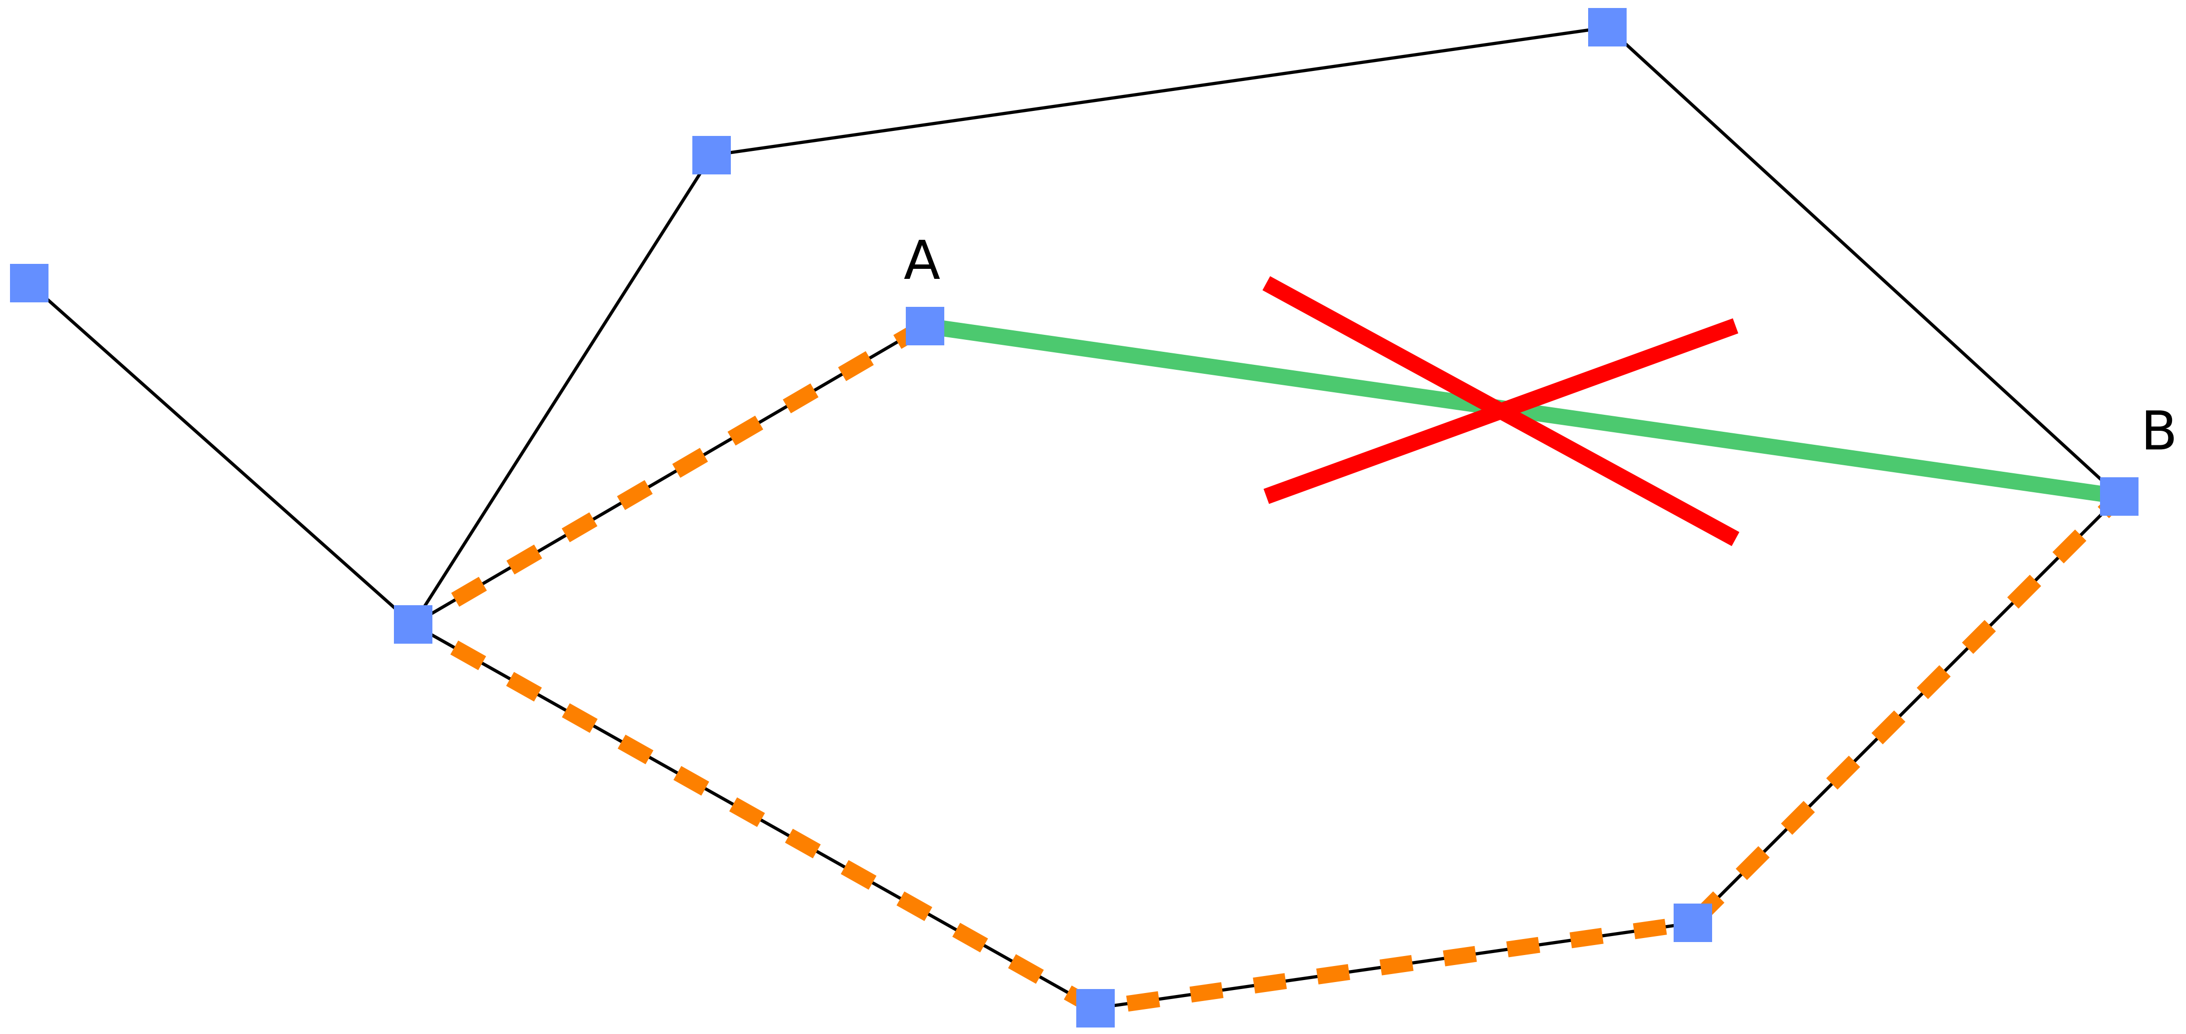

The Problem

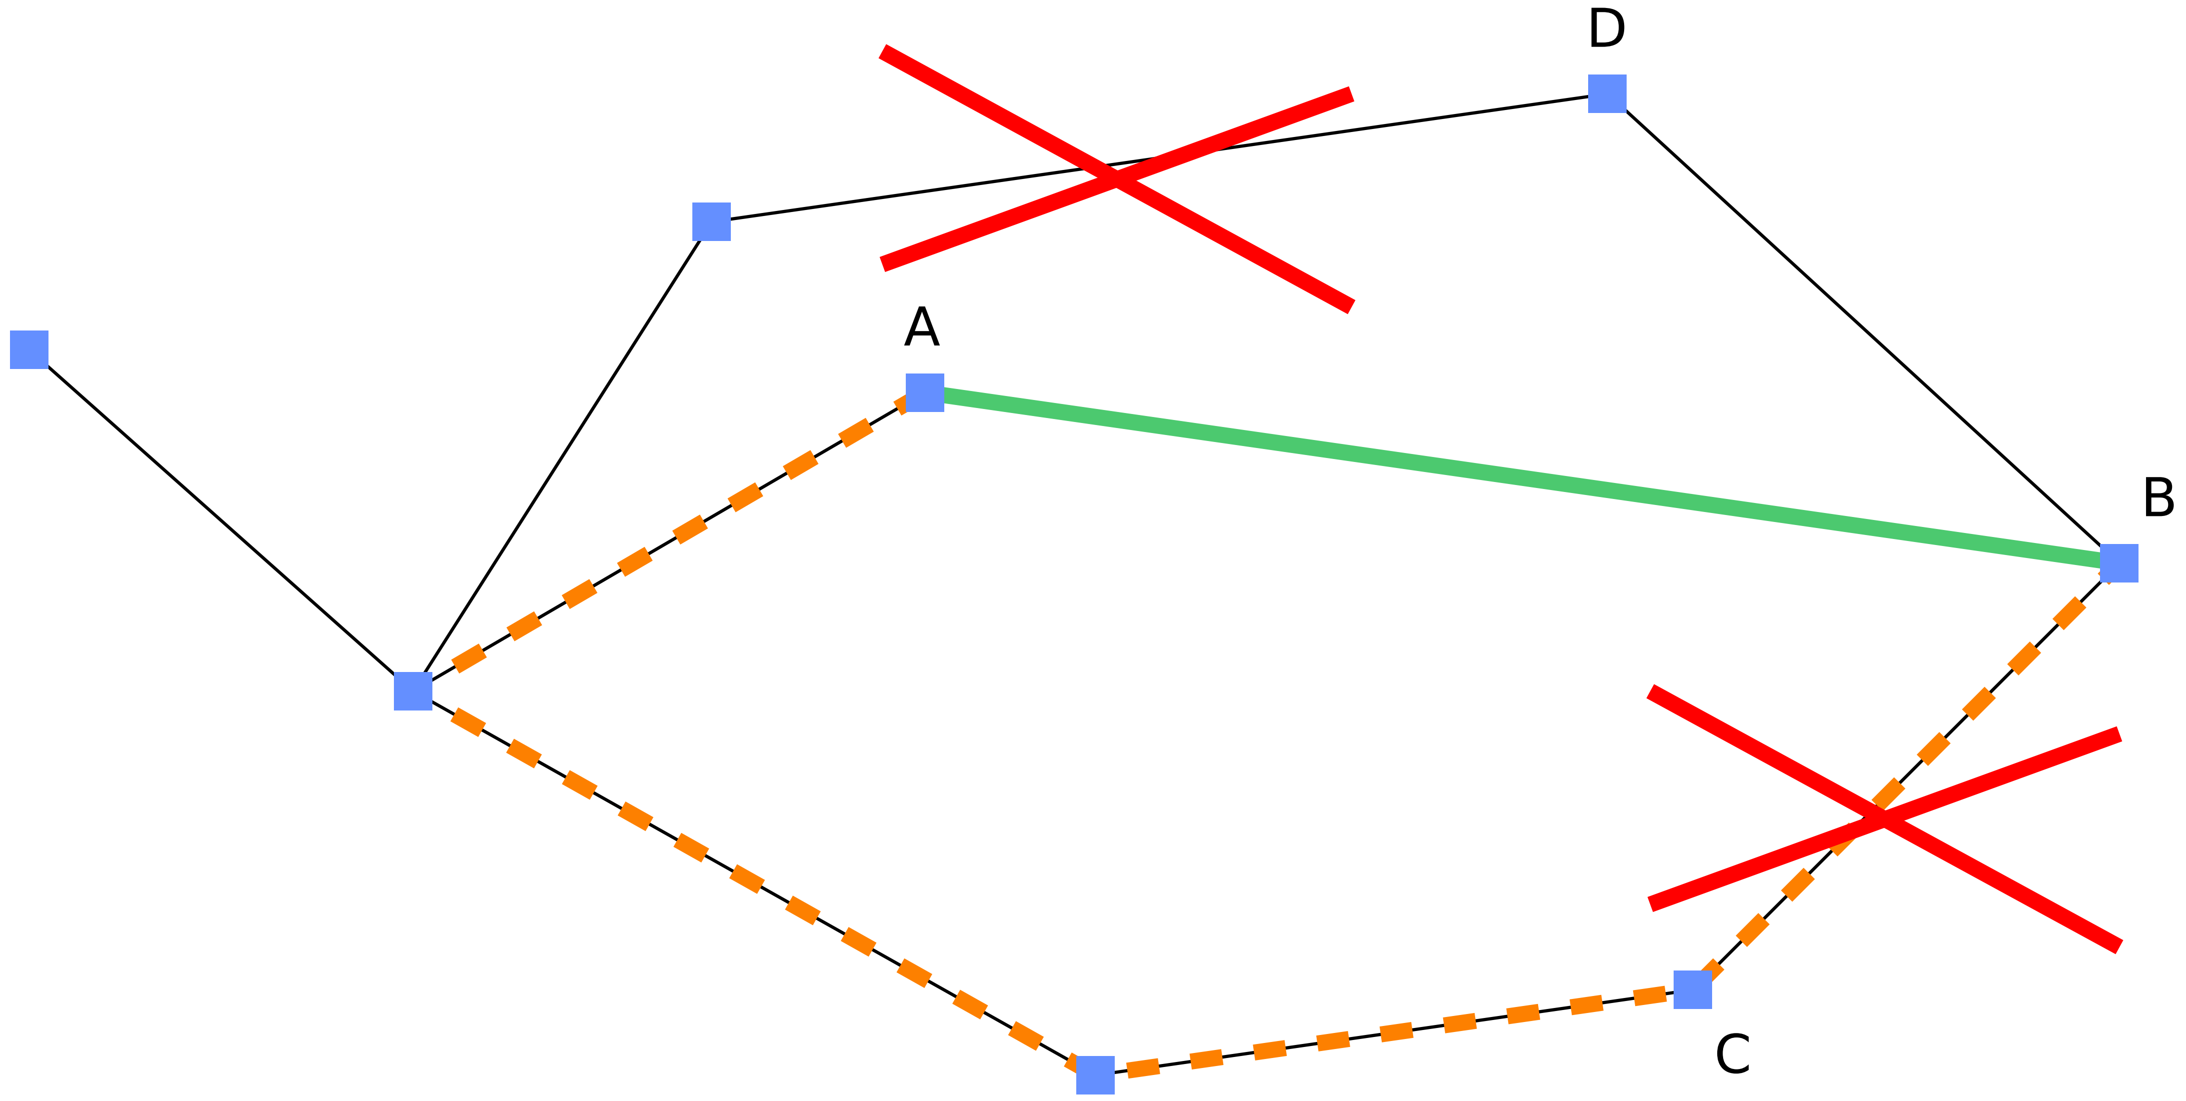

The Problem

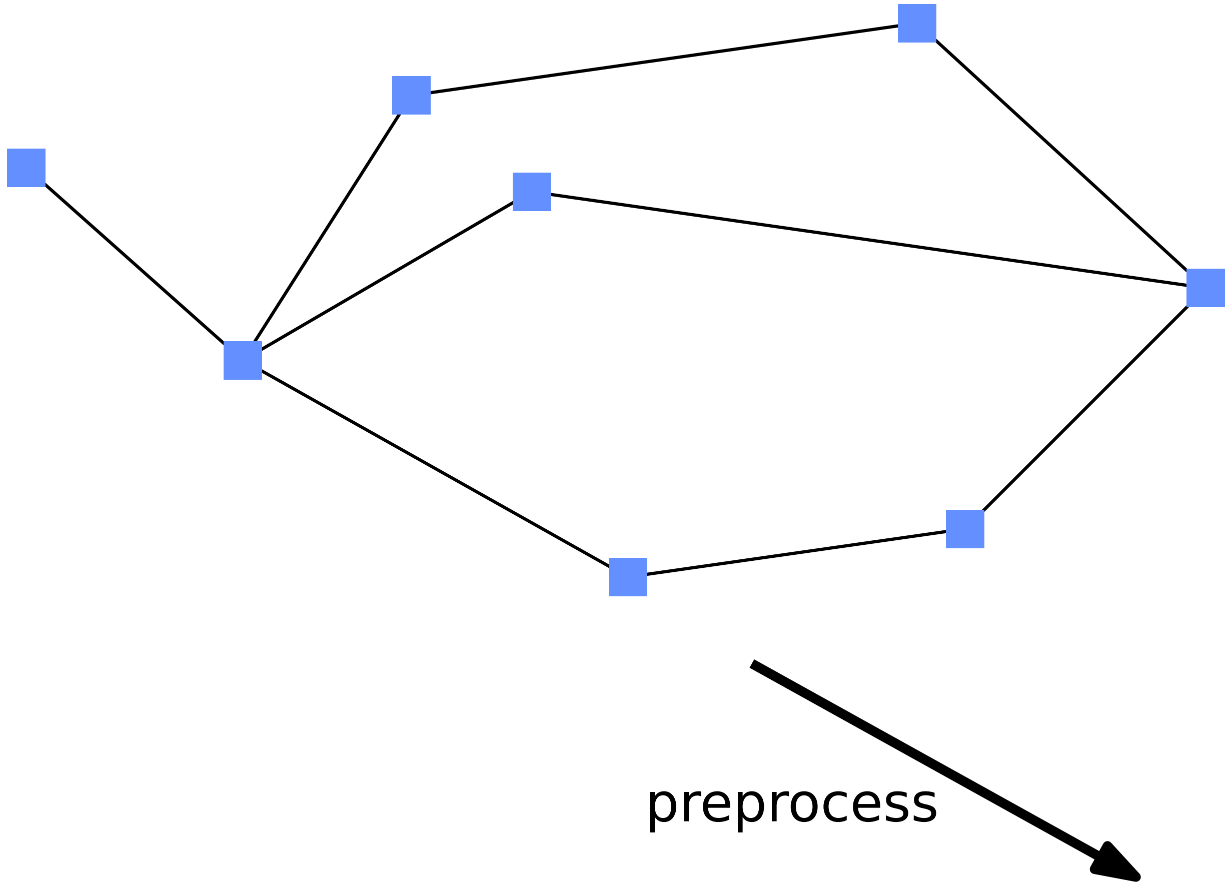

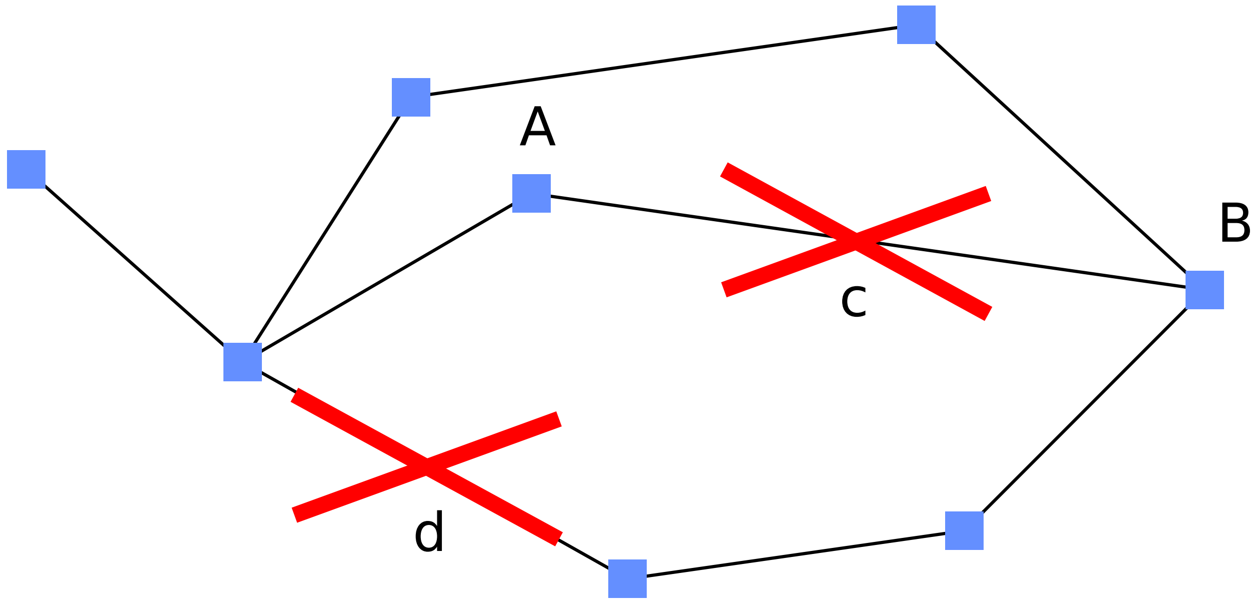

What is the shortest distance between **A** and **B** with failures on edges **c** and **d**?

It is 12.

And between **E** and **F** with a failure on edge **g**?

It is 7.

oracle

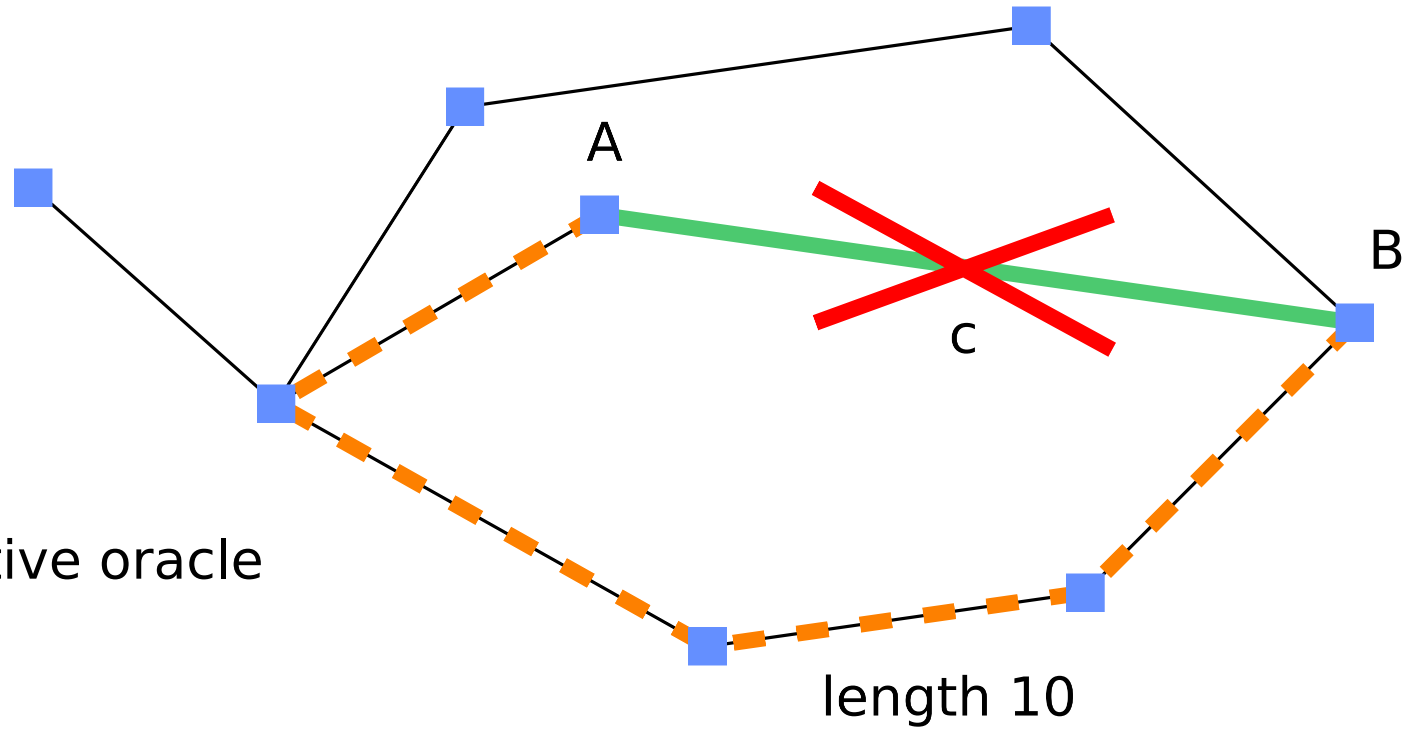

f-edge fault-tolerant distance sensitive oracle

with stretch σ

set F

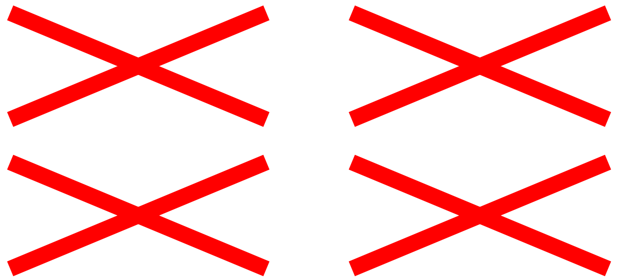

at most $f \in o(\log n / \log \log n)$ failing edges

query: $(A, B, \{c\})$

$\downarrow$

$10 \leq \text{answer} \leq \sigma \cdot 10$

Our Goal

goal: create an oracle with subquadratic space

however: undirected graph $\longrightarrow$ superquadratic space or stretch $\sigma \geq 3$ Thorup and Zwick [2005]

(worse for directed)

Our Result

stretch σ	**space**	**query time**	
$(8k-2)(f+1)$	$O(fkn^{1+1/k}\log(nW))$ integer $k \geq 1$	$\tilde{O}(f\log\log d_{G-F}(s,t))$	Chechik, Langberg, Peleg, Roditty [Algorithmica 2012]
$2k-1$	$O(f^{1-1/k}n^{1+1/k})$	$\Omega(n^{1+1/k})$	using a spanner
$3+\epsilon$	$O(n^{2-\frac{\alpha}{f+1}})$ $0 < \alpha < \frac{1}{2}$	$O(n^{\alpha})$	Bilò, Chechik, Choudhary, Cohen, Friedrich, K, Schirneck [STOC 2023]
$2k-1$	$O(n^{1+\frac{1}{k}+\alpha+o(1)})$ $0 < \alpha < 1$ integer $k \geq 2$	$\tilde{O}(n^{1+\frac{1}{k}-\frac{\alpha}{k(f+1)}})$	our result

Overview

small hop diameter $< L$

many graphs with random edges missing

one of them is probably good

derandomized with error-correcting codes

long hop diameter $> L$

subgraph of pivots

replacement path goes through at least one pivot

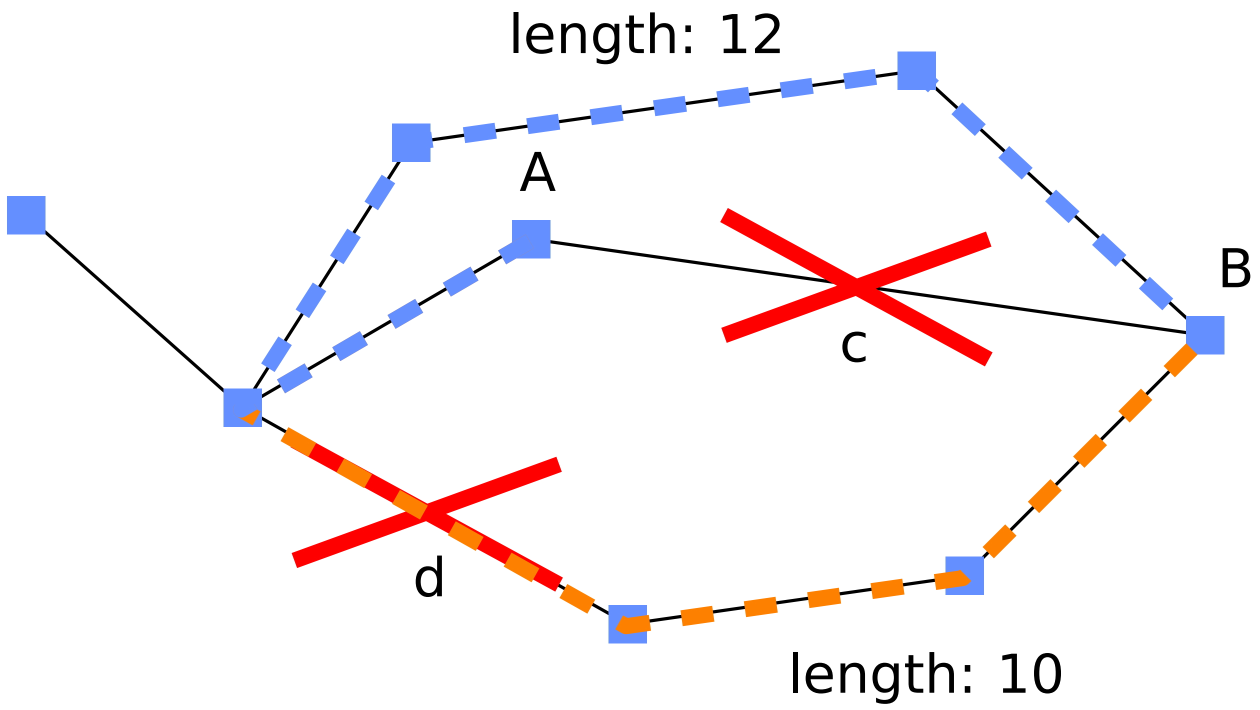

What is the shortest distance between **A** and **B** with failures on edges **c**?

It is 12.

Small Hop Diameter

idea from Weimann and Yuster [2013]

deleted with probability $\frac{1}{L}$

$\tilde{O}(fL^f)$ subgraphs

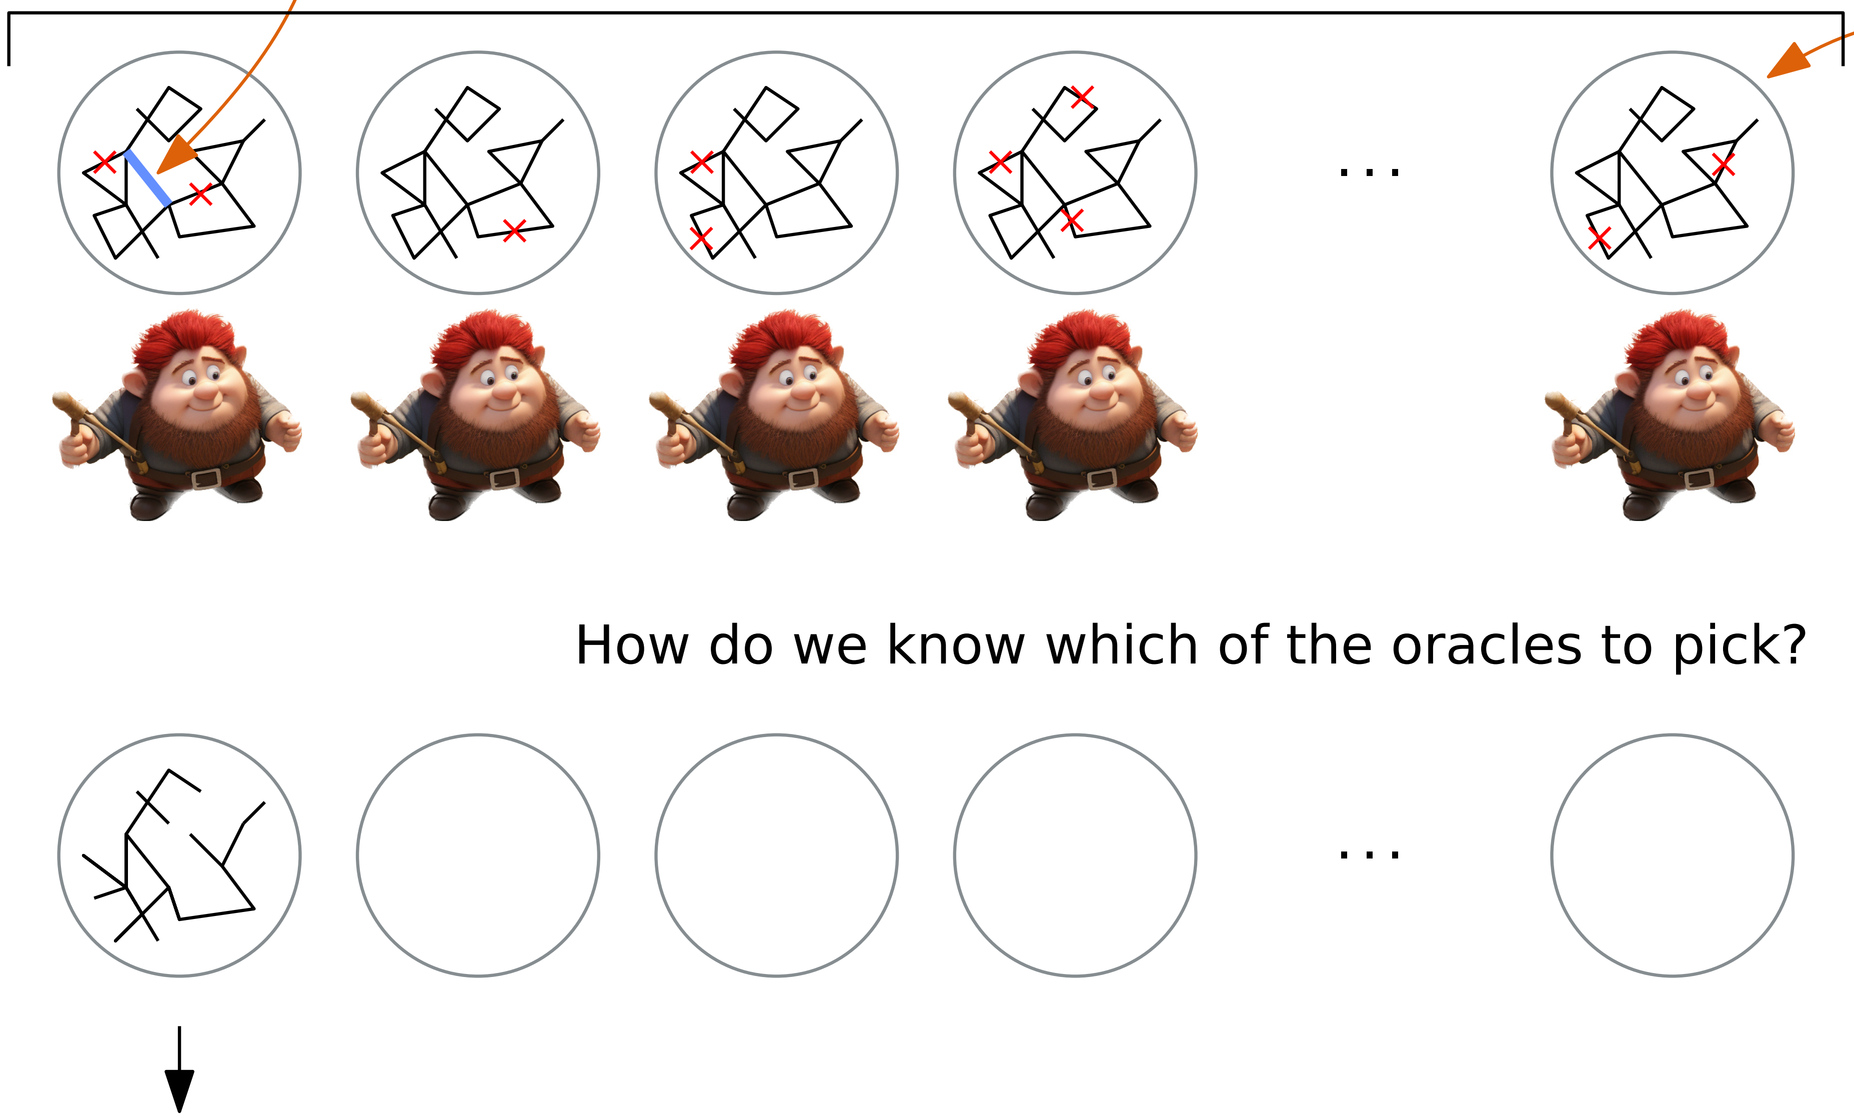

too large

replace with DOs by
Thorup and Zwick [2005]

How do we know which of the oracles to pick?

spanners from
Thorup and Zwick [2005]

static dictionary
with $O(1)$ lookup

lemma: spanner has no failing edges → oracle is applicable

Small Hop Diameter

query algorithm:

1. select spanners with *F* (or more) missing

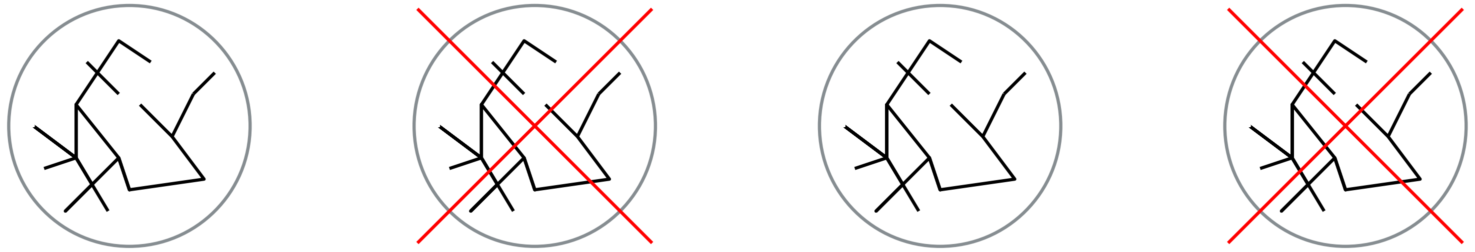

2. query corresponding oracles

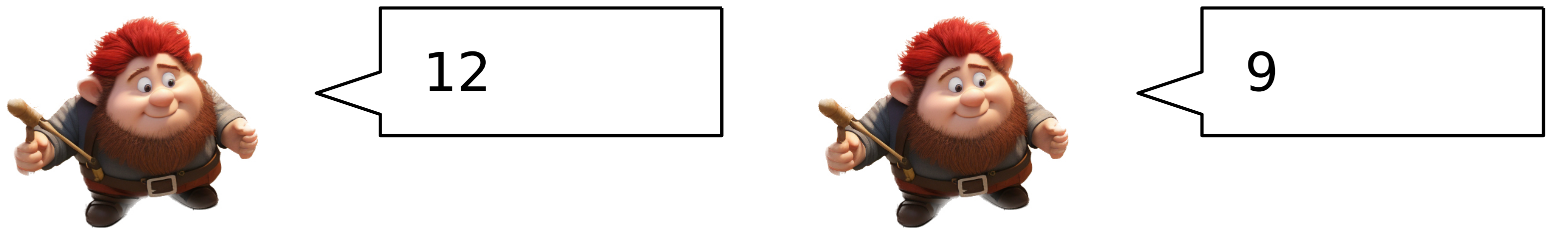

3. return lowest distance

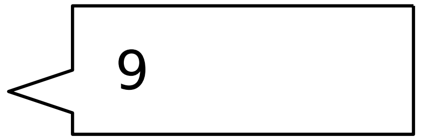

Small Hop Diameter - Derandomization

how to do this deterministically?

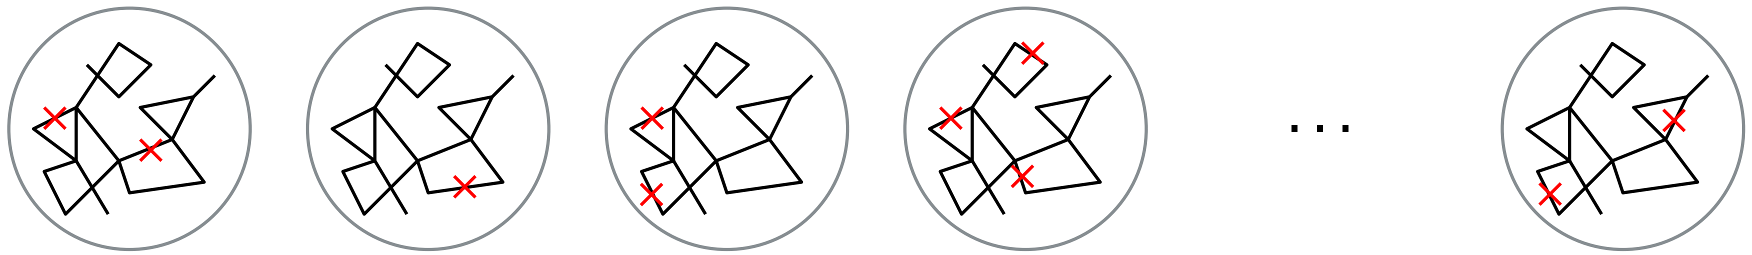

edge *a*	0	0	1	0	1	1
edge *b*	0	1	0	1	0	1
edge *c*	1	0	0	1	1	0
edge *d*	1	1	1	0	0	0

ideally: one oracle for each possible F

↓

obviously too many

instead: Reed-Solomon codes (adapted by Karthik and Parter [2021])

$\rightarrow$ identify correct oracles in $O(1)$ no spanners needed

Small Hop Diameter - Derandomization

derandomization tradeoffs:

more space

more preprocessing time

less query time

Large Hop Diameter

select pivots $B \subset V$

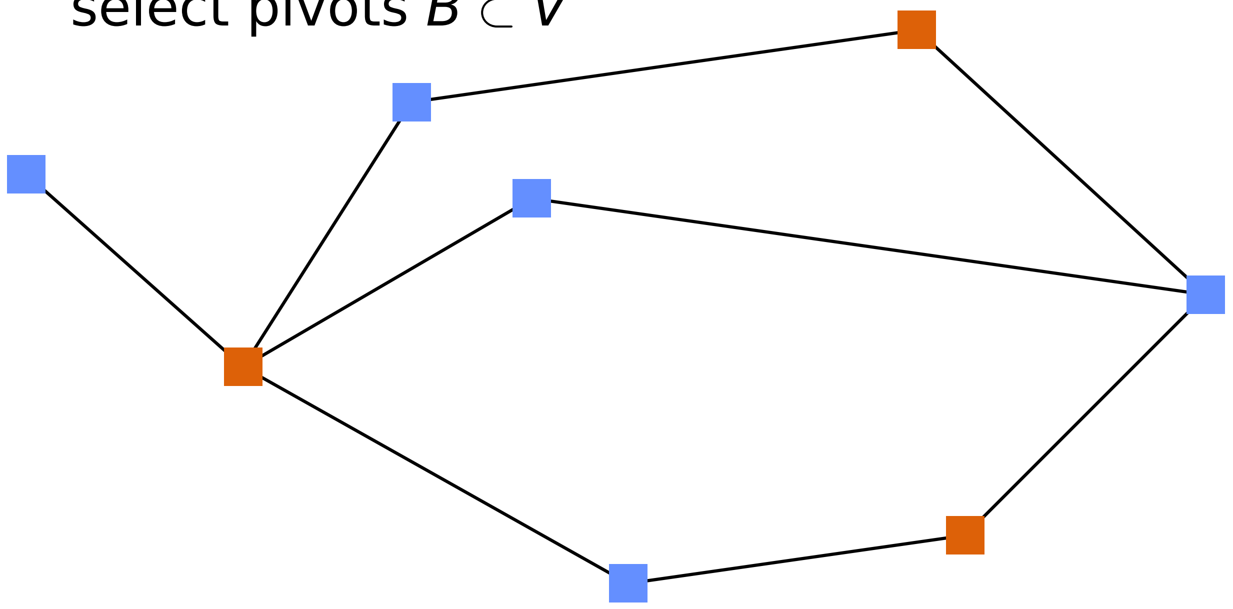

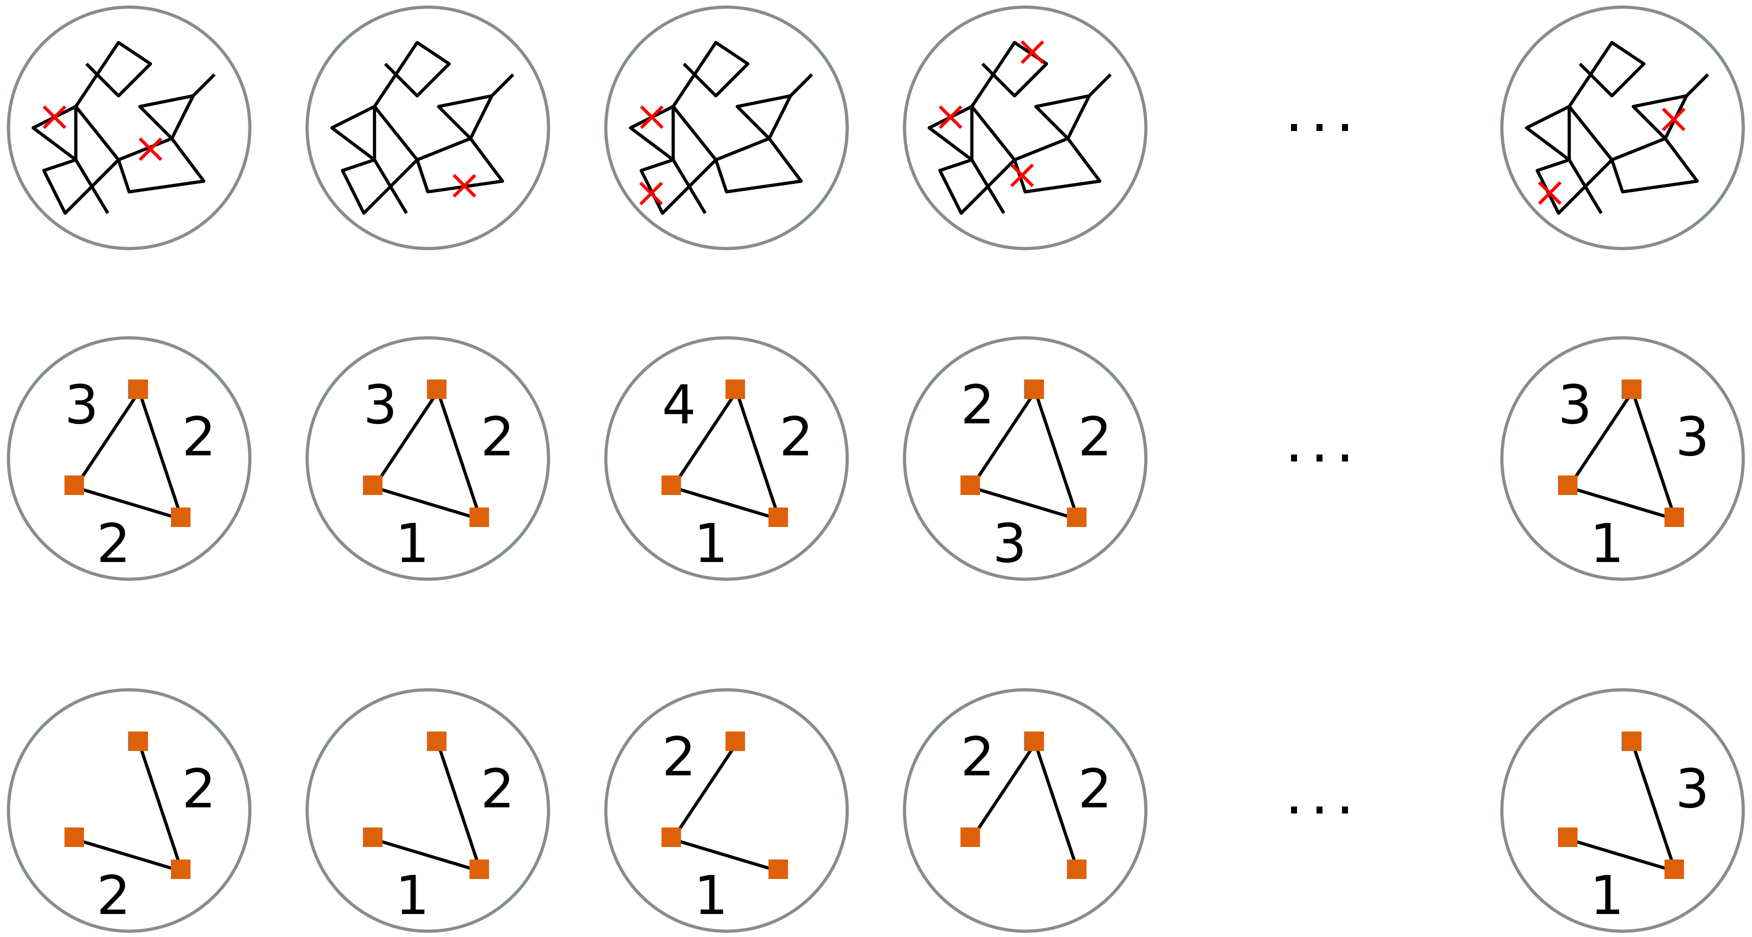

all pivot distances

pivot spanners
(Roddity, Thorup, Zwick [2005])

Large Hop Diameter

query algorithm:

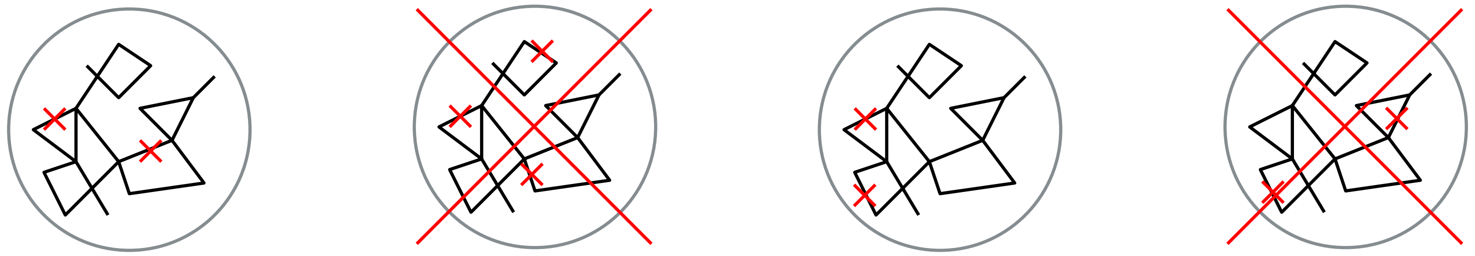

1. select graphs with F (or more) missing
2. ask short hop diameter oracle first

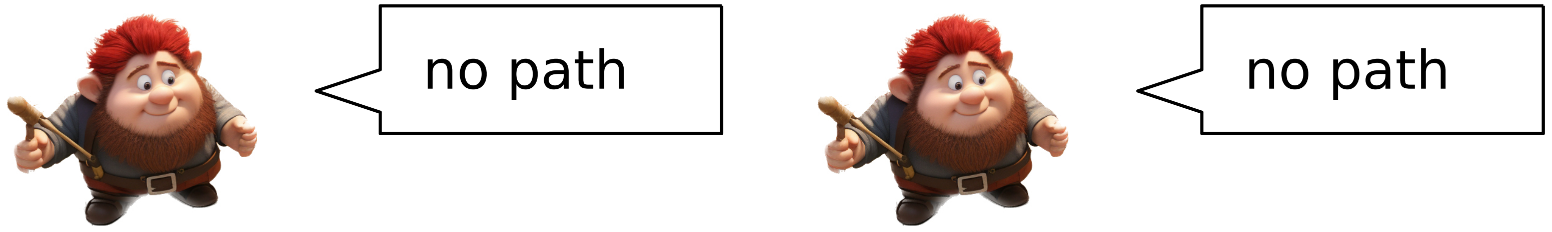

3. merge pivot spanners

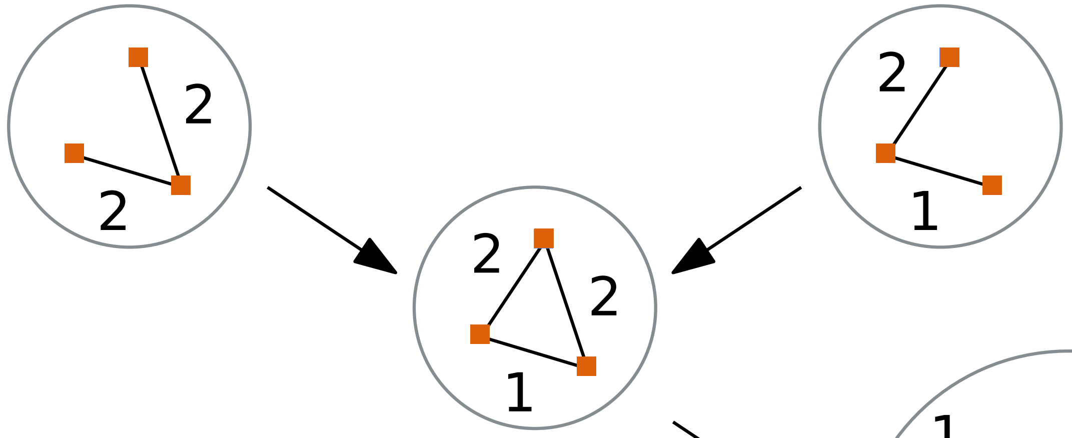

4. add connections to s and t for each pivot and each graph

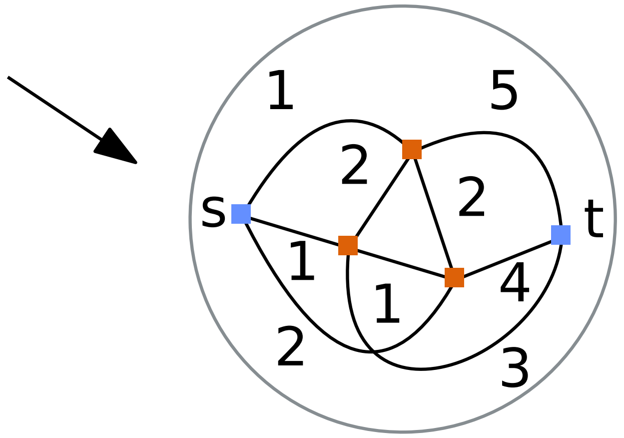

5. return minimum of shortest path from step 2 and 4

Summary

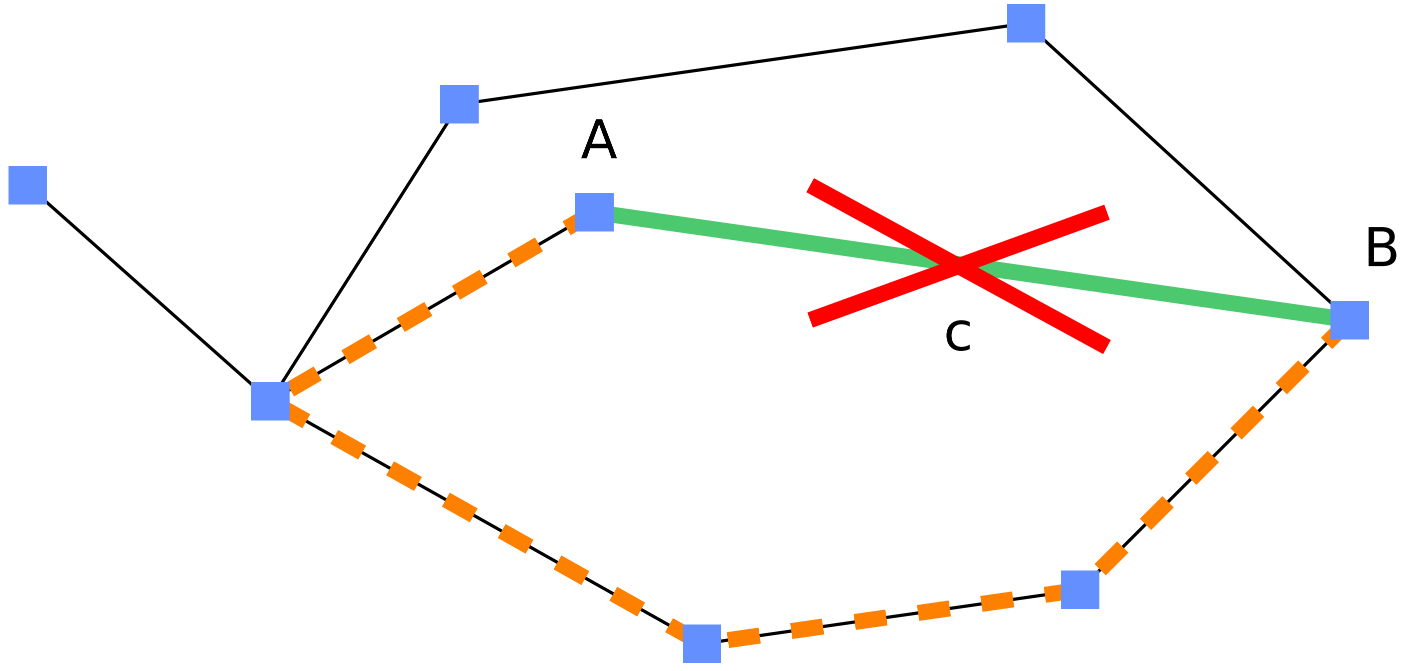

oracle

stretch σ	**space**	**query time**	**preprocessing time**
$2k-1$	$O(n^{1+\frac{1}{k}+\alpha+o(1)})$ $0<\alpha<1$ integer $k \geq 2$	$\tilde{O}(n^{1+\frac{1}{k}-\frac{\alpha}{k(f+1)}})$	$kmn^{1+\alpha+\frac{1}{k}+o(1)}$

Thanks for listening!