Lower Bounds for Local Algorithms

Jukka Suomela
Aalto University, Finland

ADGA · Austin, Texas · 12 October 2014

LOCAL model

- **Input:** simple undirected **graph G**
  - communication network
  - nodes labelled with unique $O(\log n)$-bit identifiers

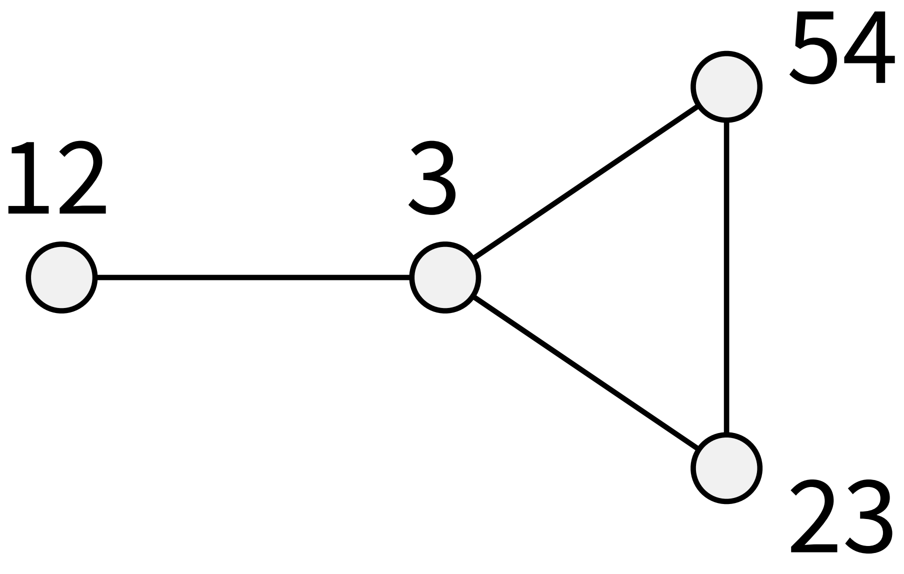

LOCAL model

- **Input:** simple undirected graph *G*
- **Output:** each node *v* produces a **local output**
  - graph colouring: colour of node *v*
  - vertex cover: 1 if *v* is in the cover
  - matching: with whom *v* is matched

LOCAL model

- Nodes exchange messages with each other, update local states
- Synchronous communication rounds
- Arbitrarily large messages

LOCAL model

- **Time = number of communication rounds**
  - until all nodes stop and produce their **local outputs**

LOCAL model

- **Time = number of communication rounds**
- **Time = distance:**
  - in *t* communication rounds,
    all nodes can learn everything
    in their **radius-*t* neighbourhoods**

time $t = 2$

LOCAL model

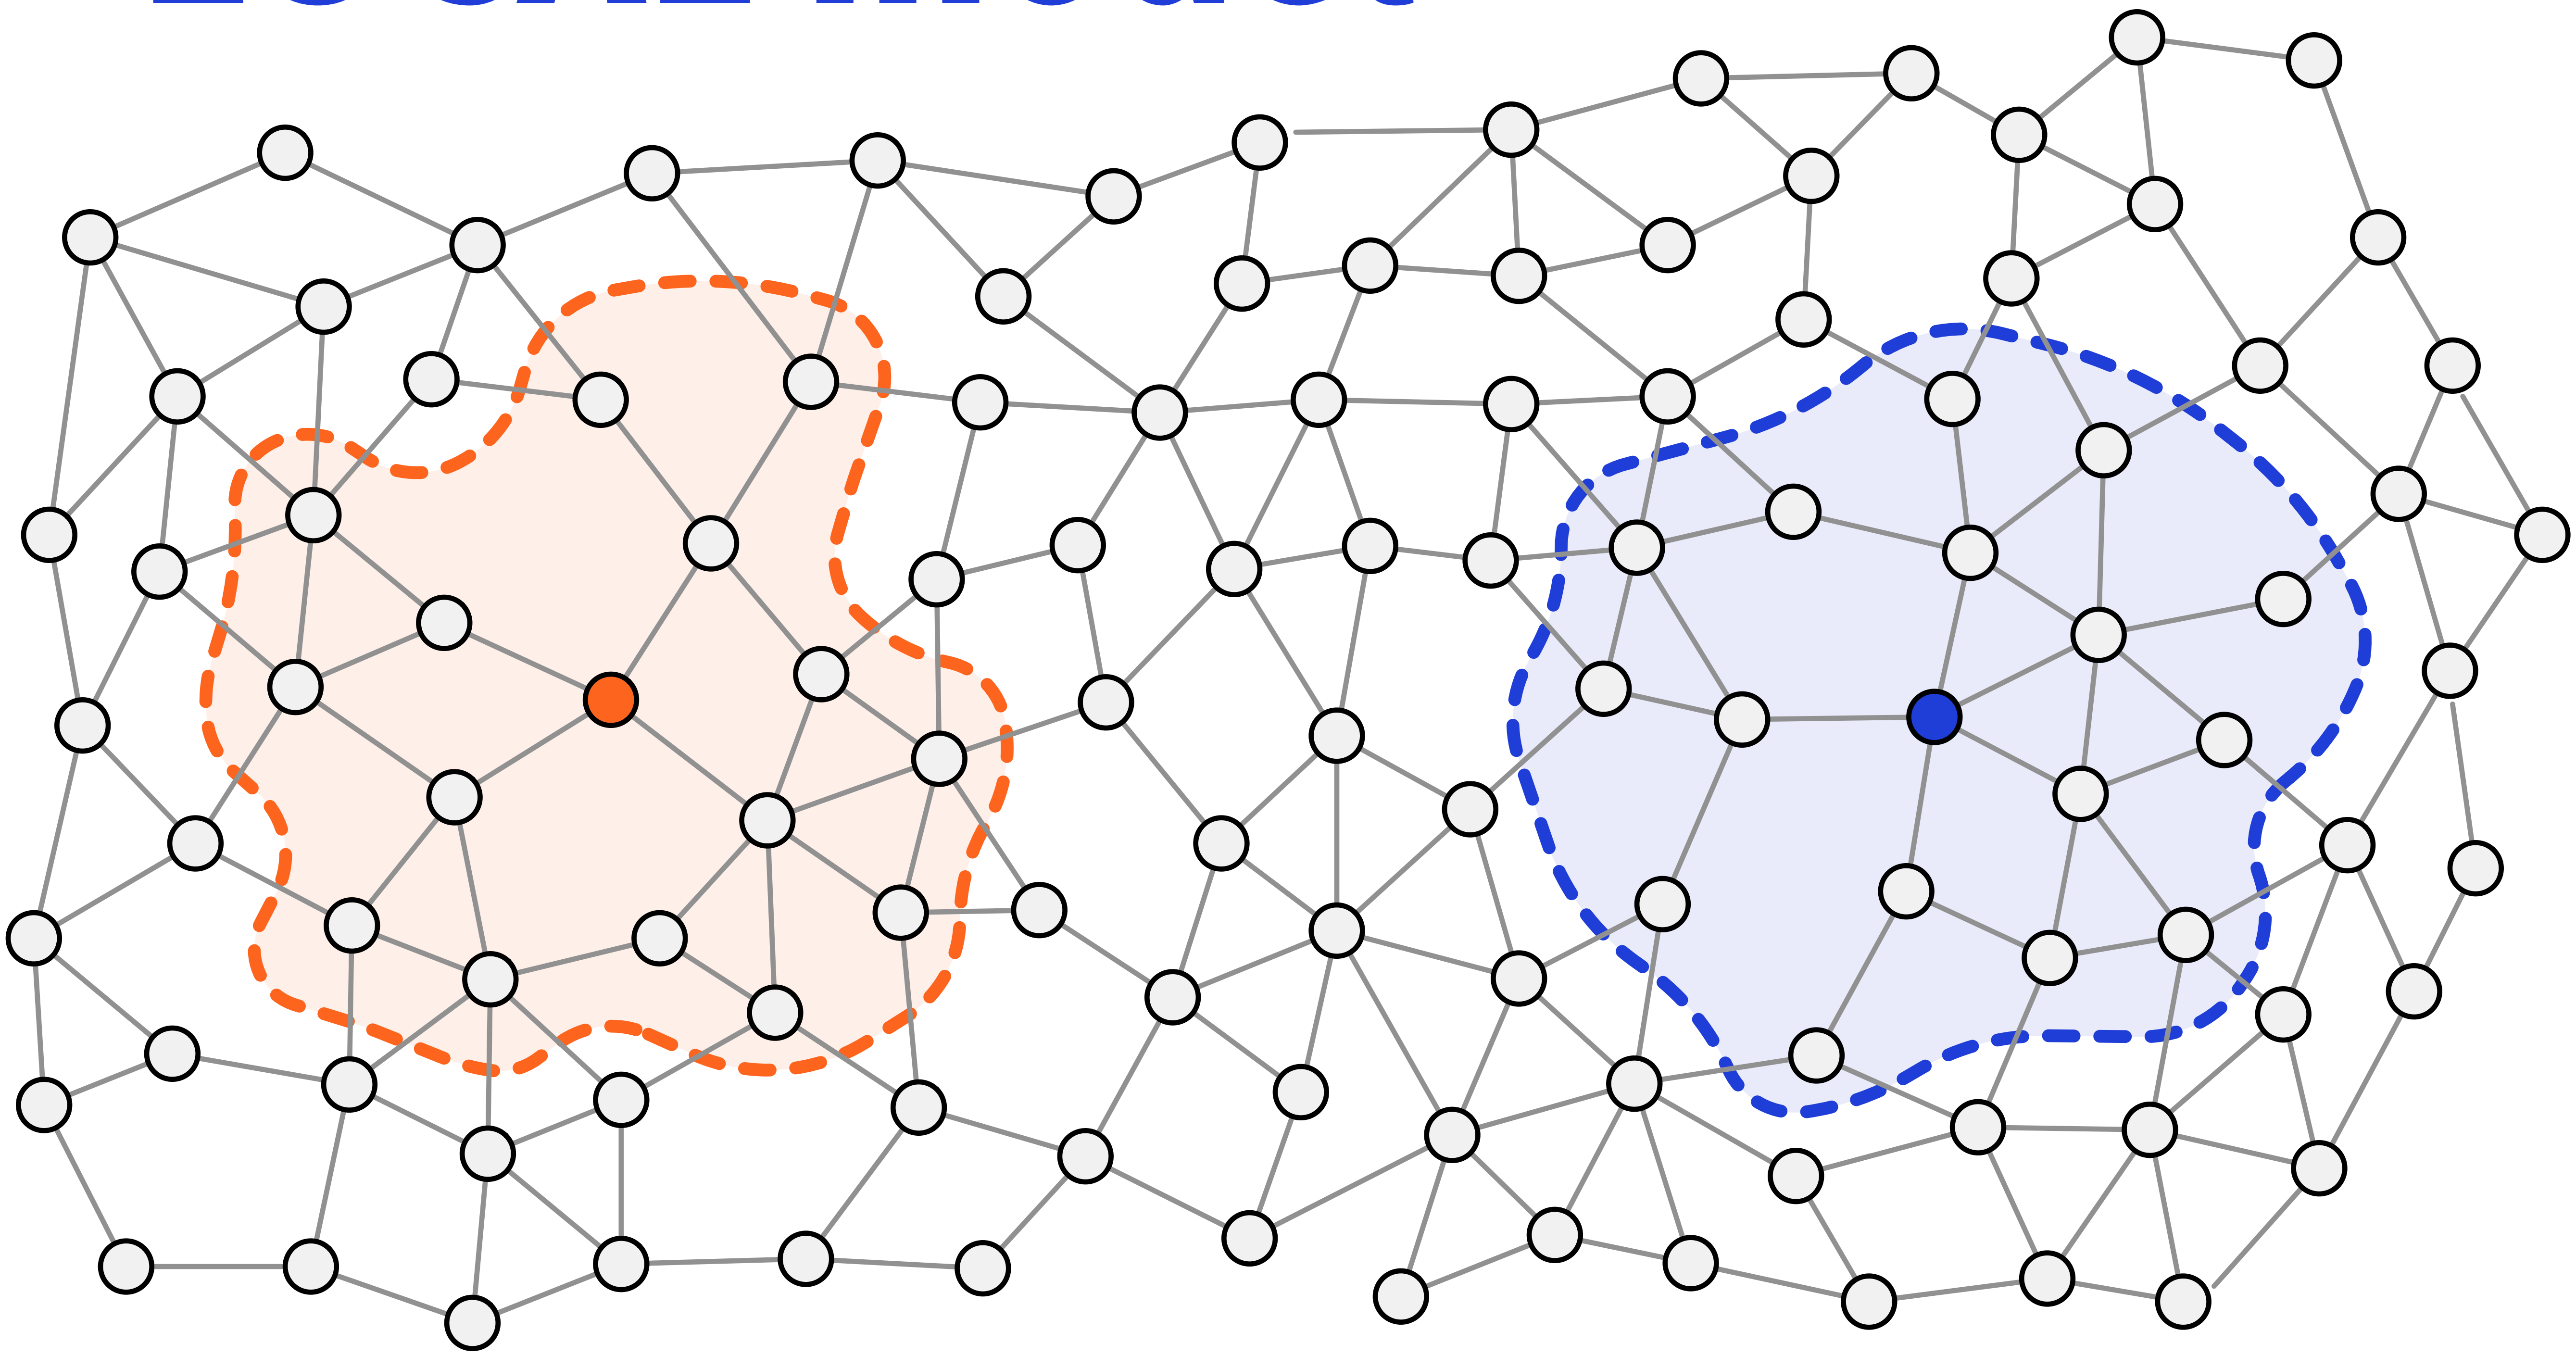

LOCAL model

A:

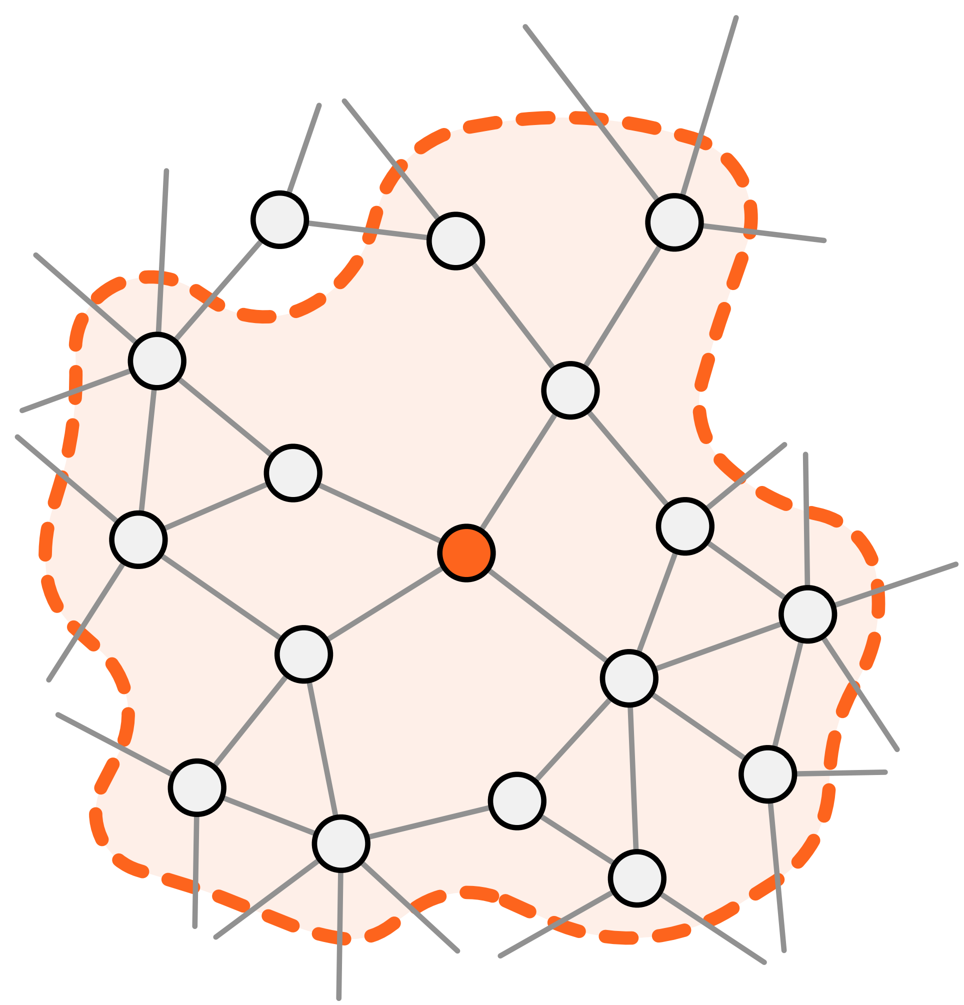

$\mapsto 1$

LOCAL model

- **Everything trivial in time diam(G)**
  - all nodes see whole G, can compute any function of G
- **What can be solved much faster?**

Distributed time complexity

- Smallest t such that the problem can be solved in time t

Distributed time complexity

- ***n* = number of nodes**
- **Δ = maximum degree**
  - $\Delta < n$
- **Time complexity $t = t(n, \Delta)$**

Landscape

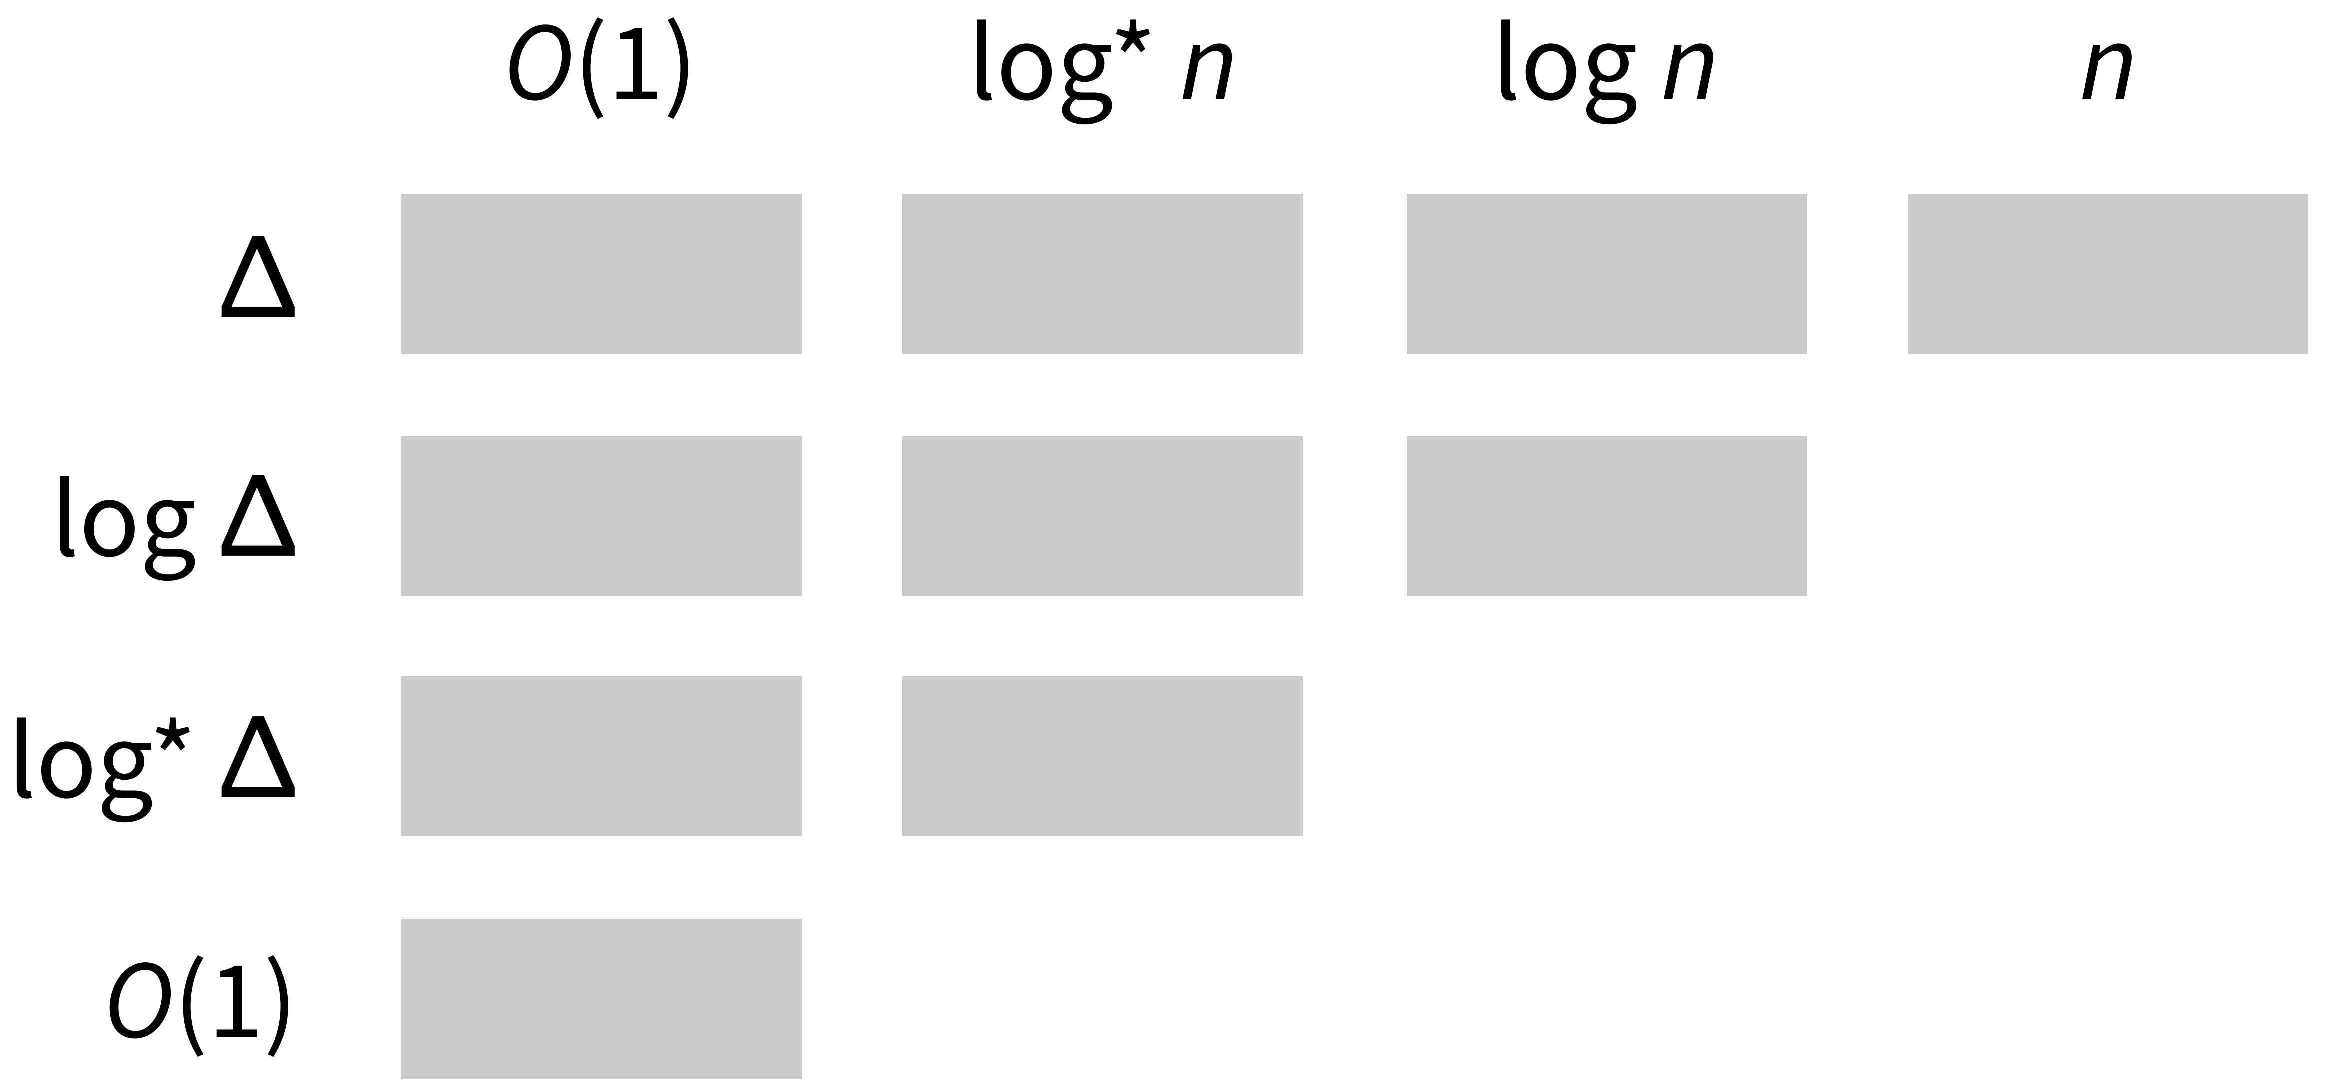

Landscape

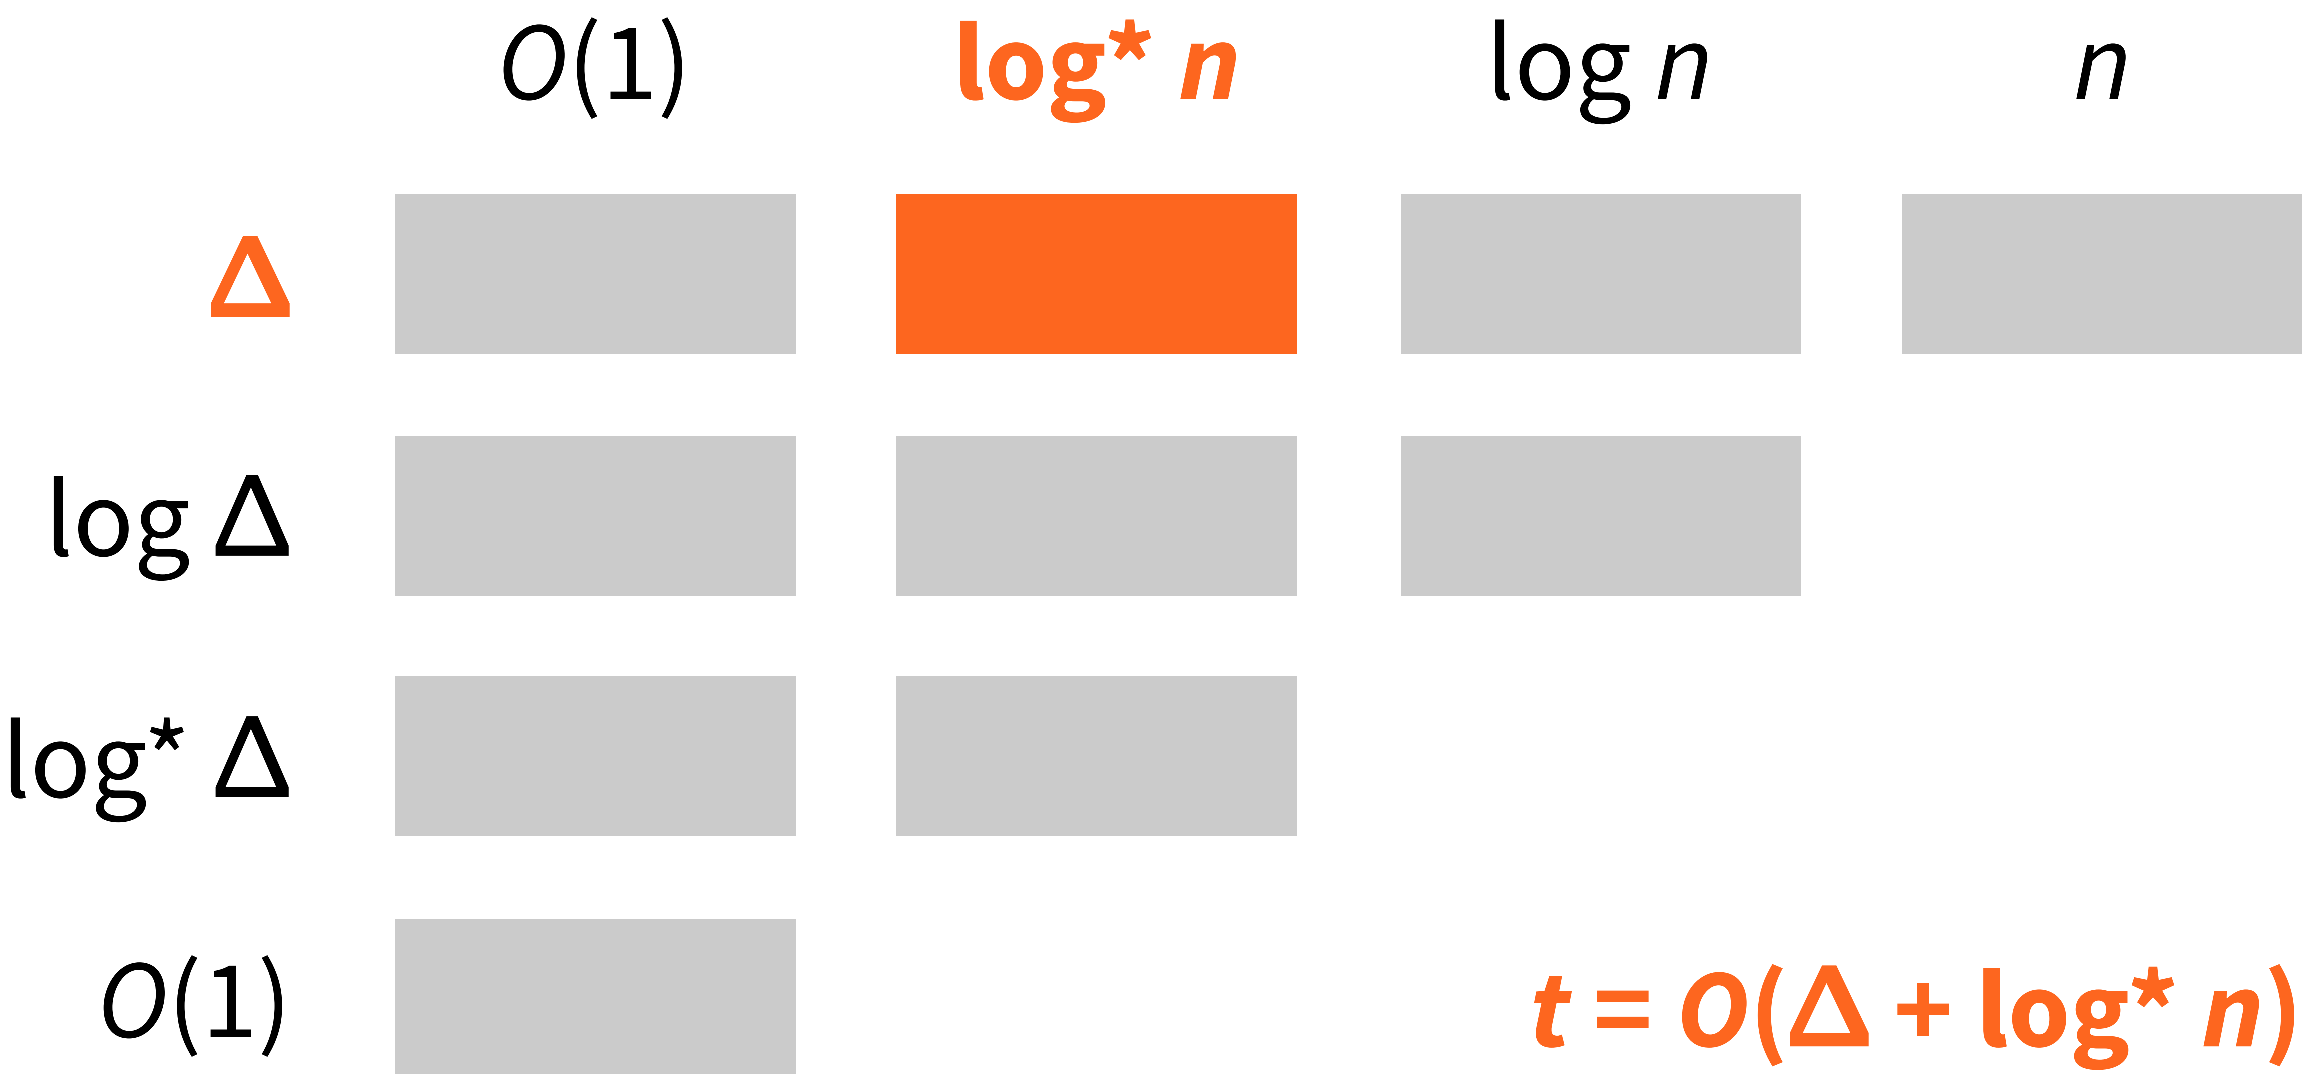

$t = O(\Delta + \log^* n)$

Landscape

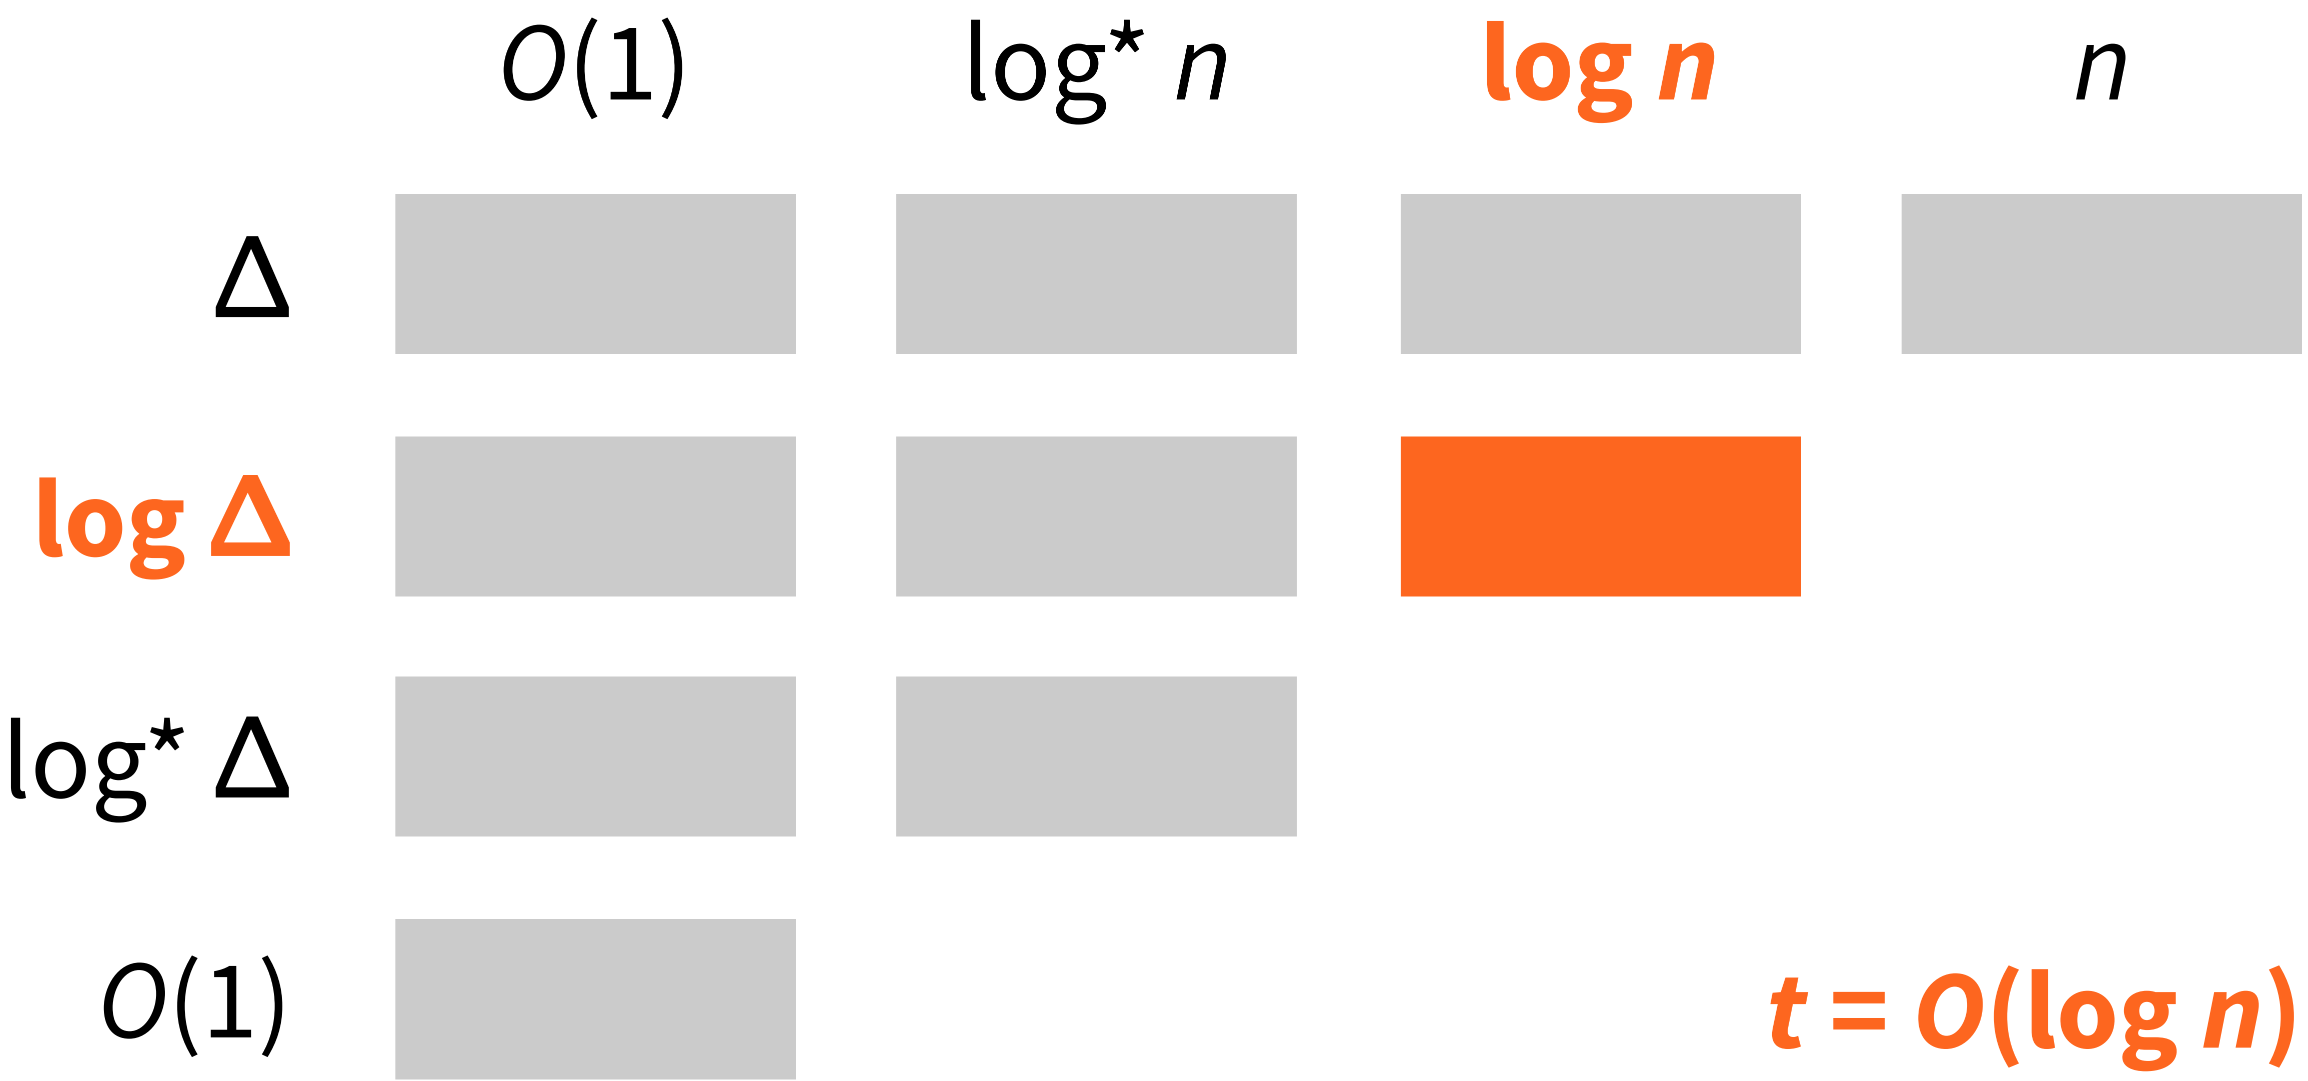

Landscape

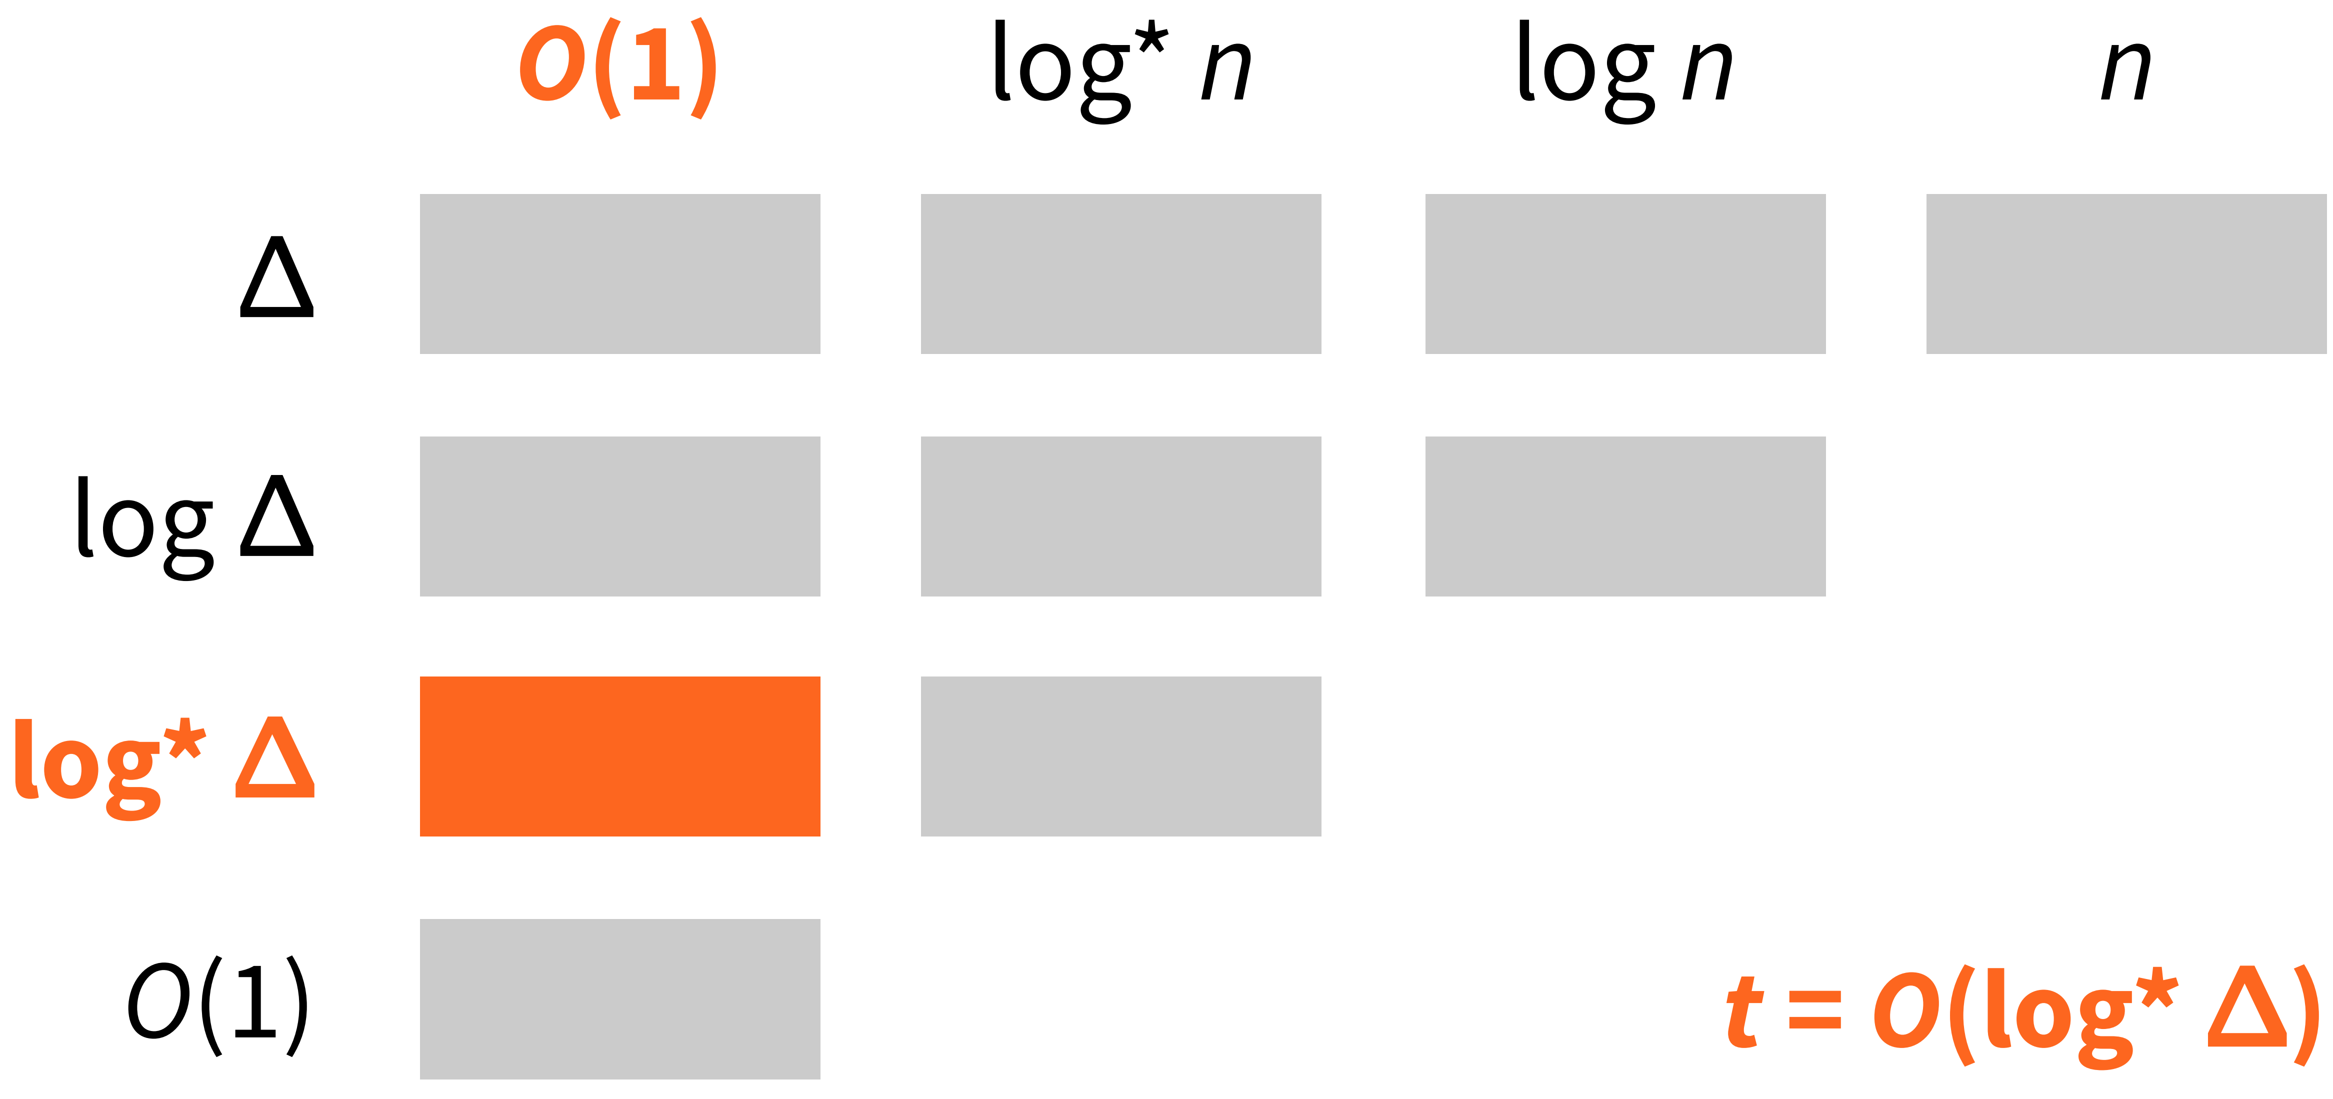

Landscape

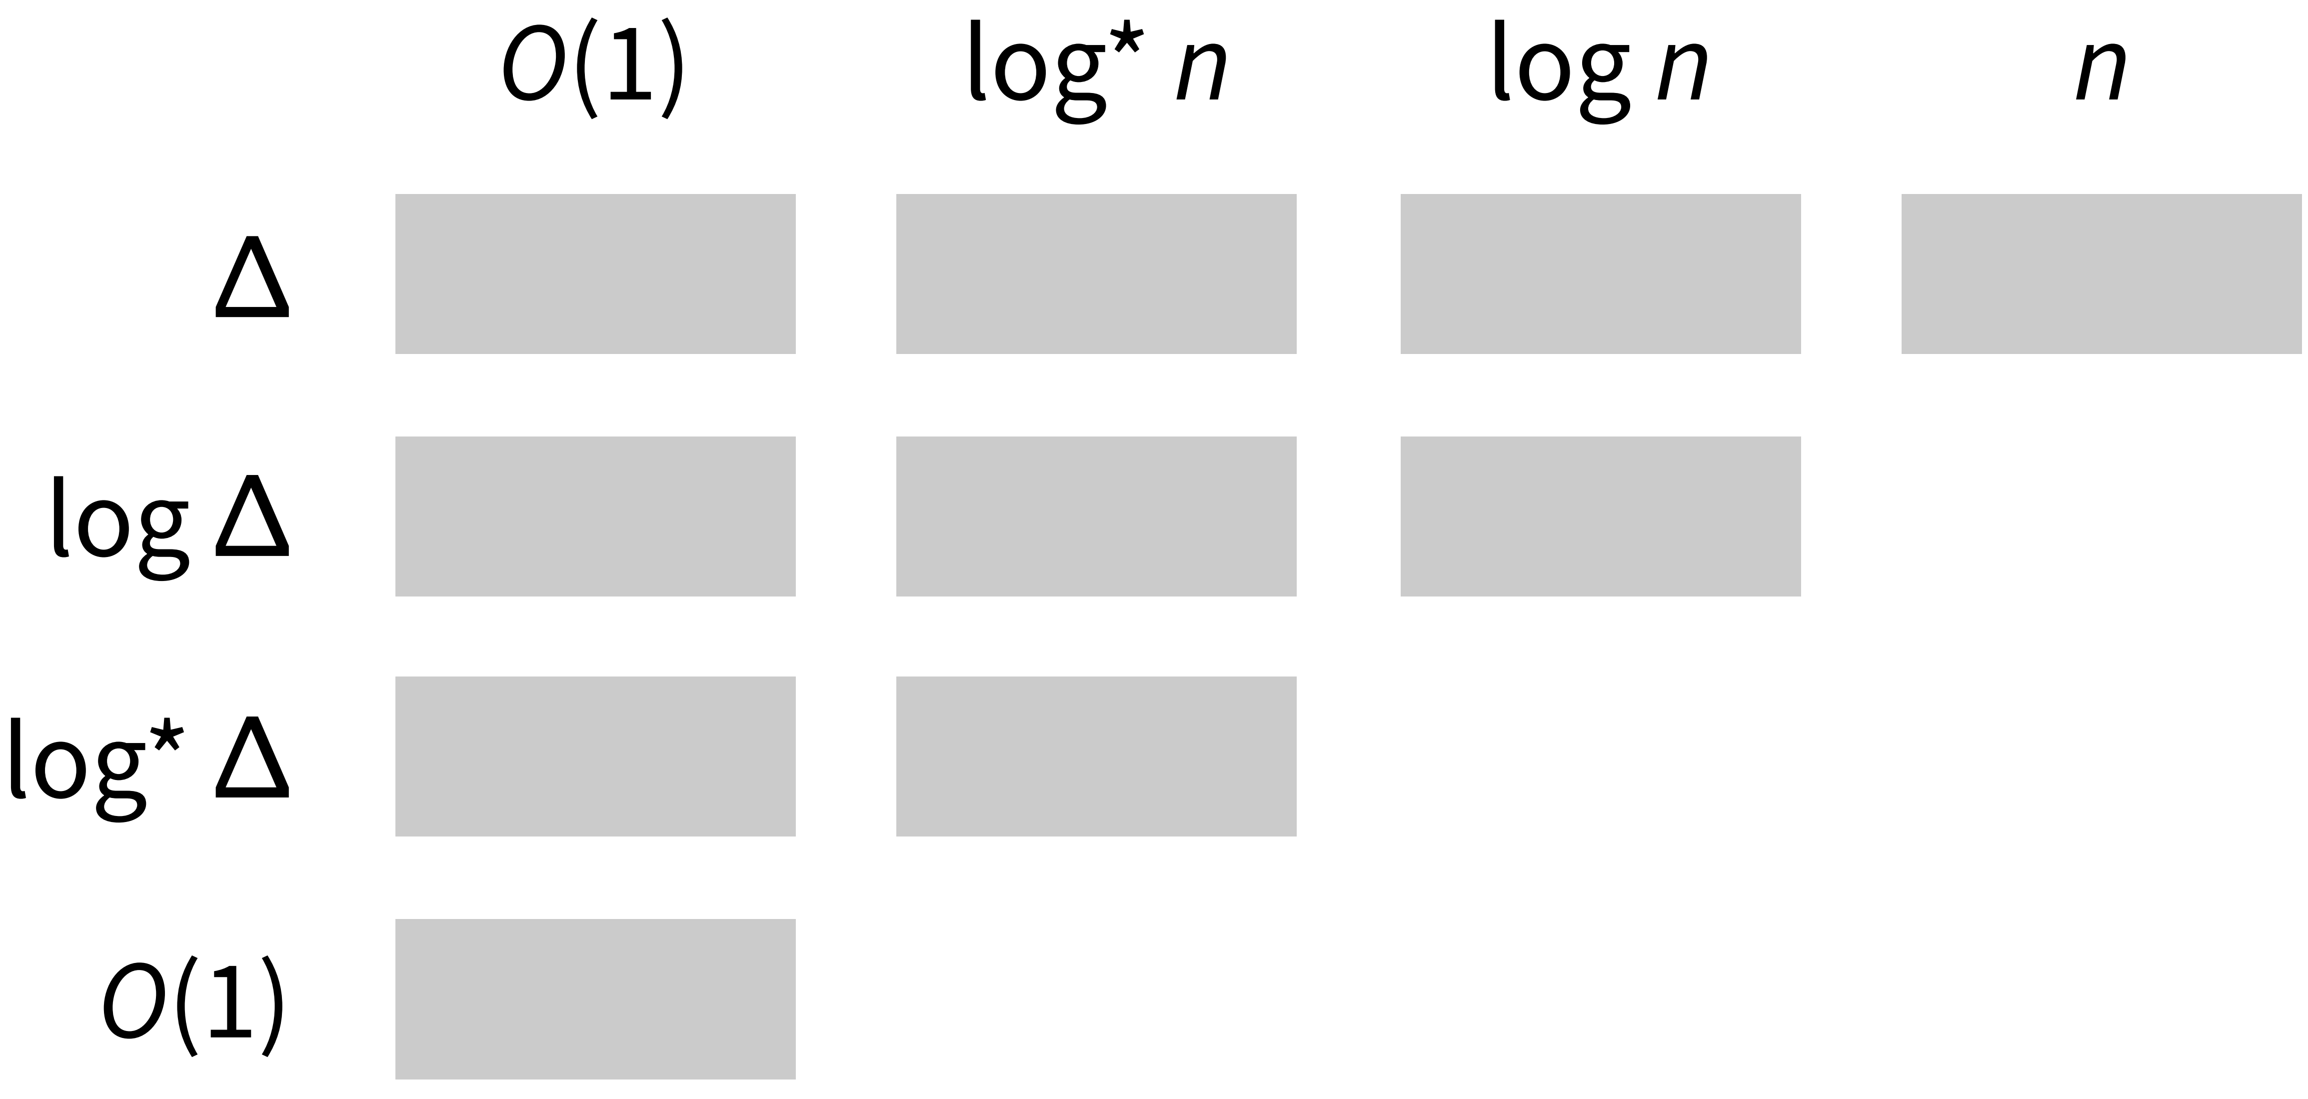

Landscape

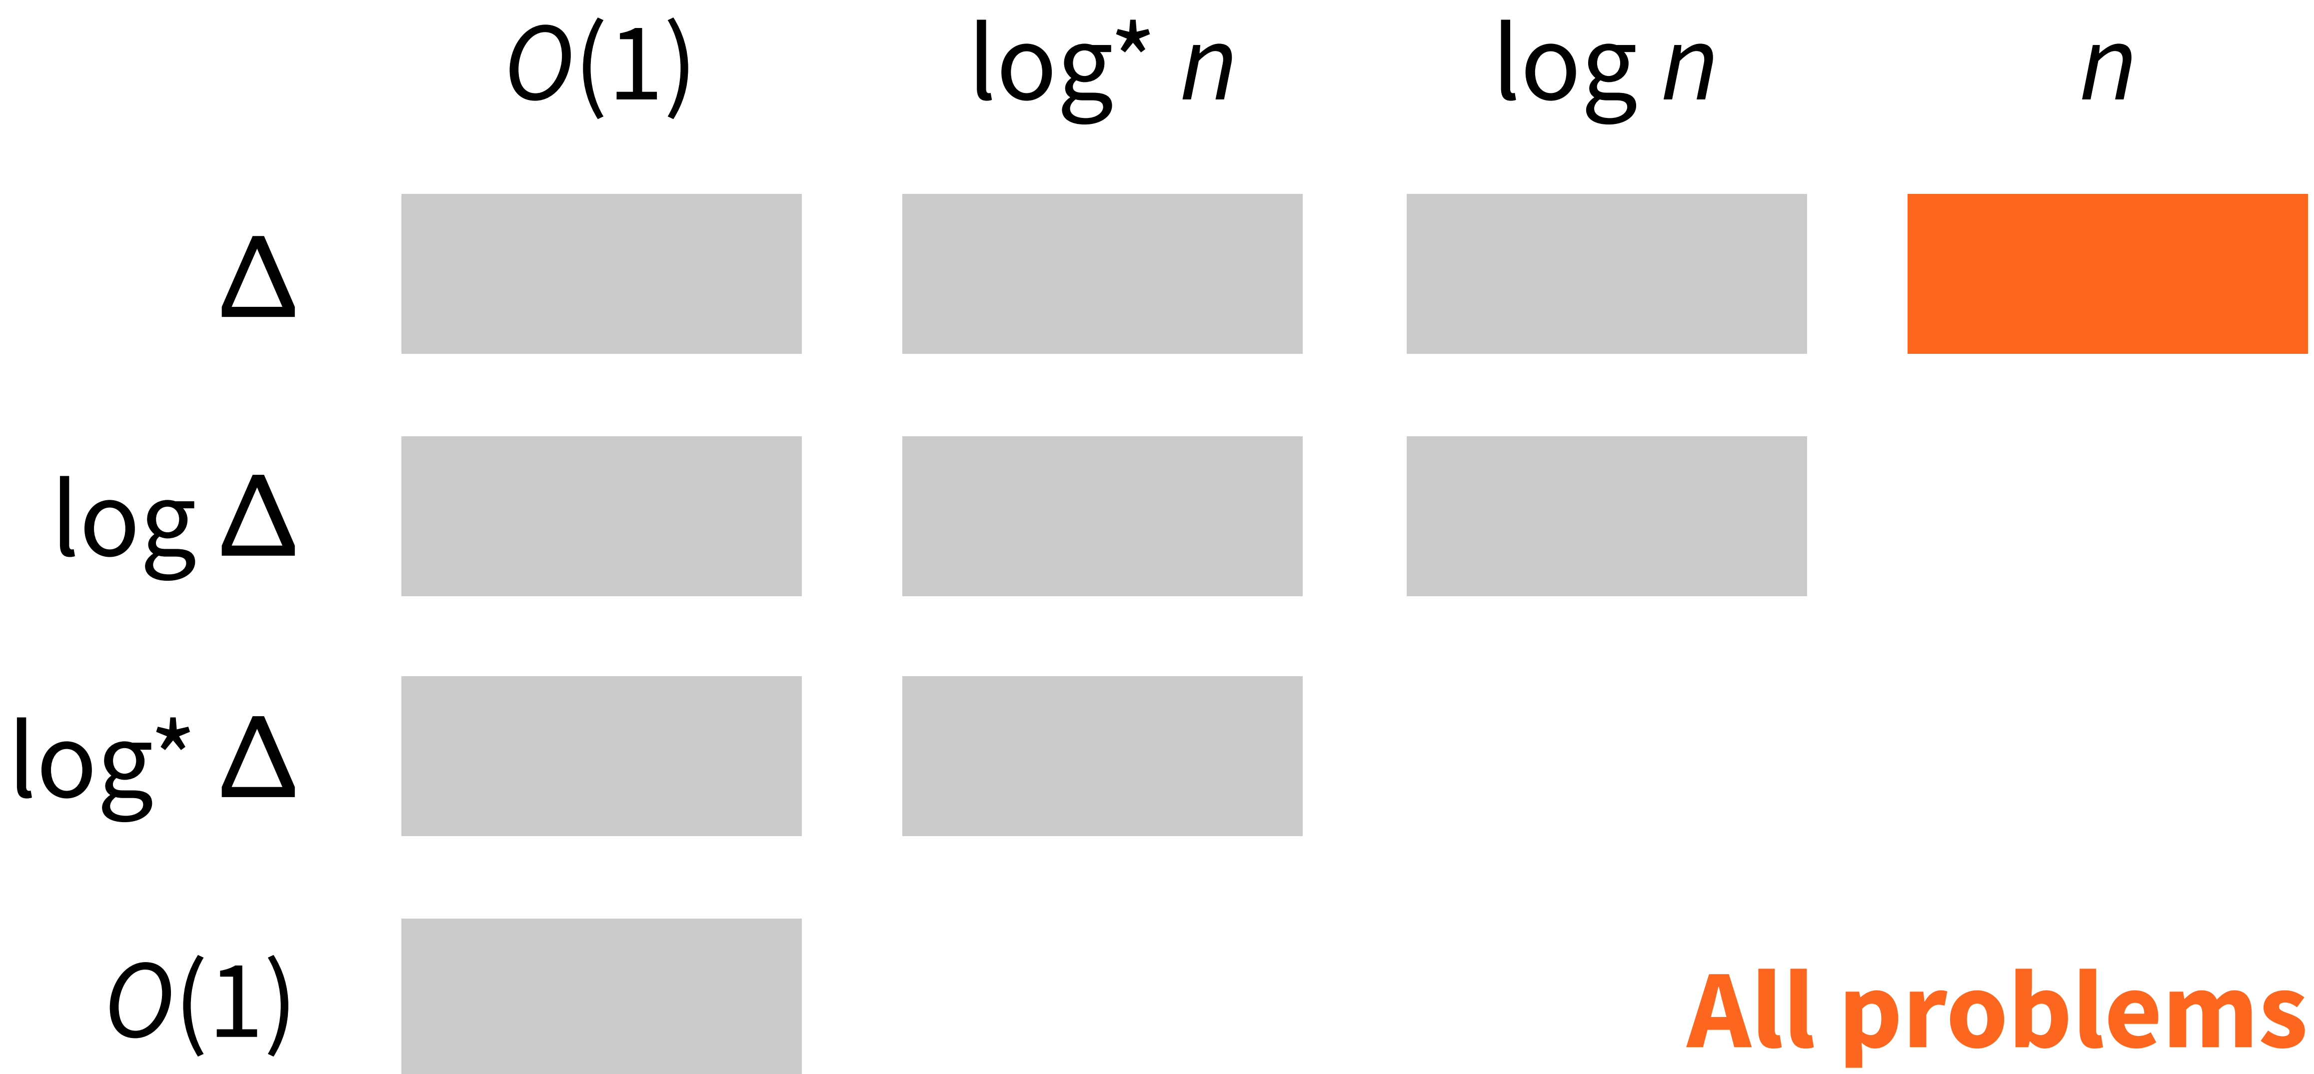

Landscape

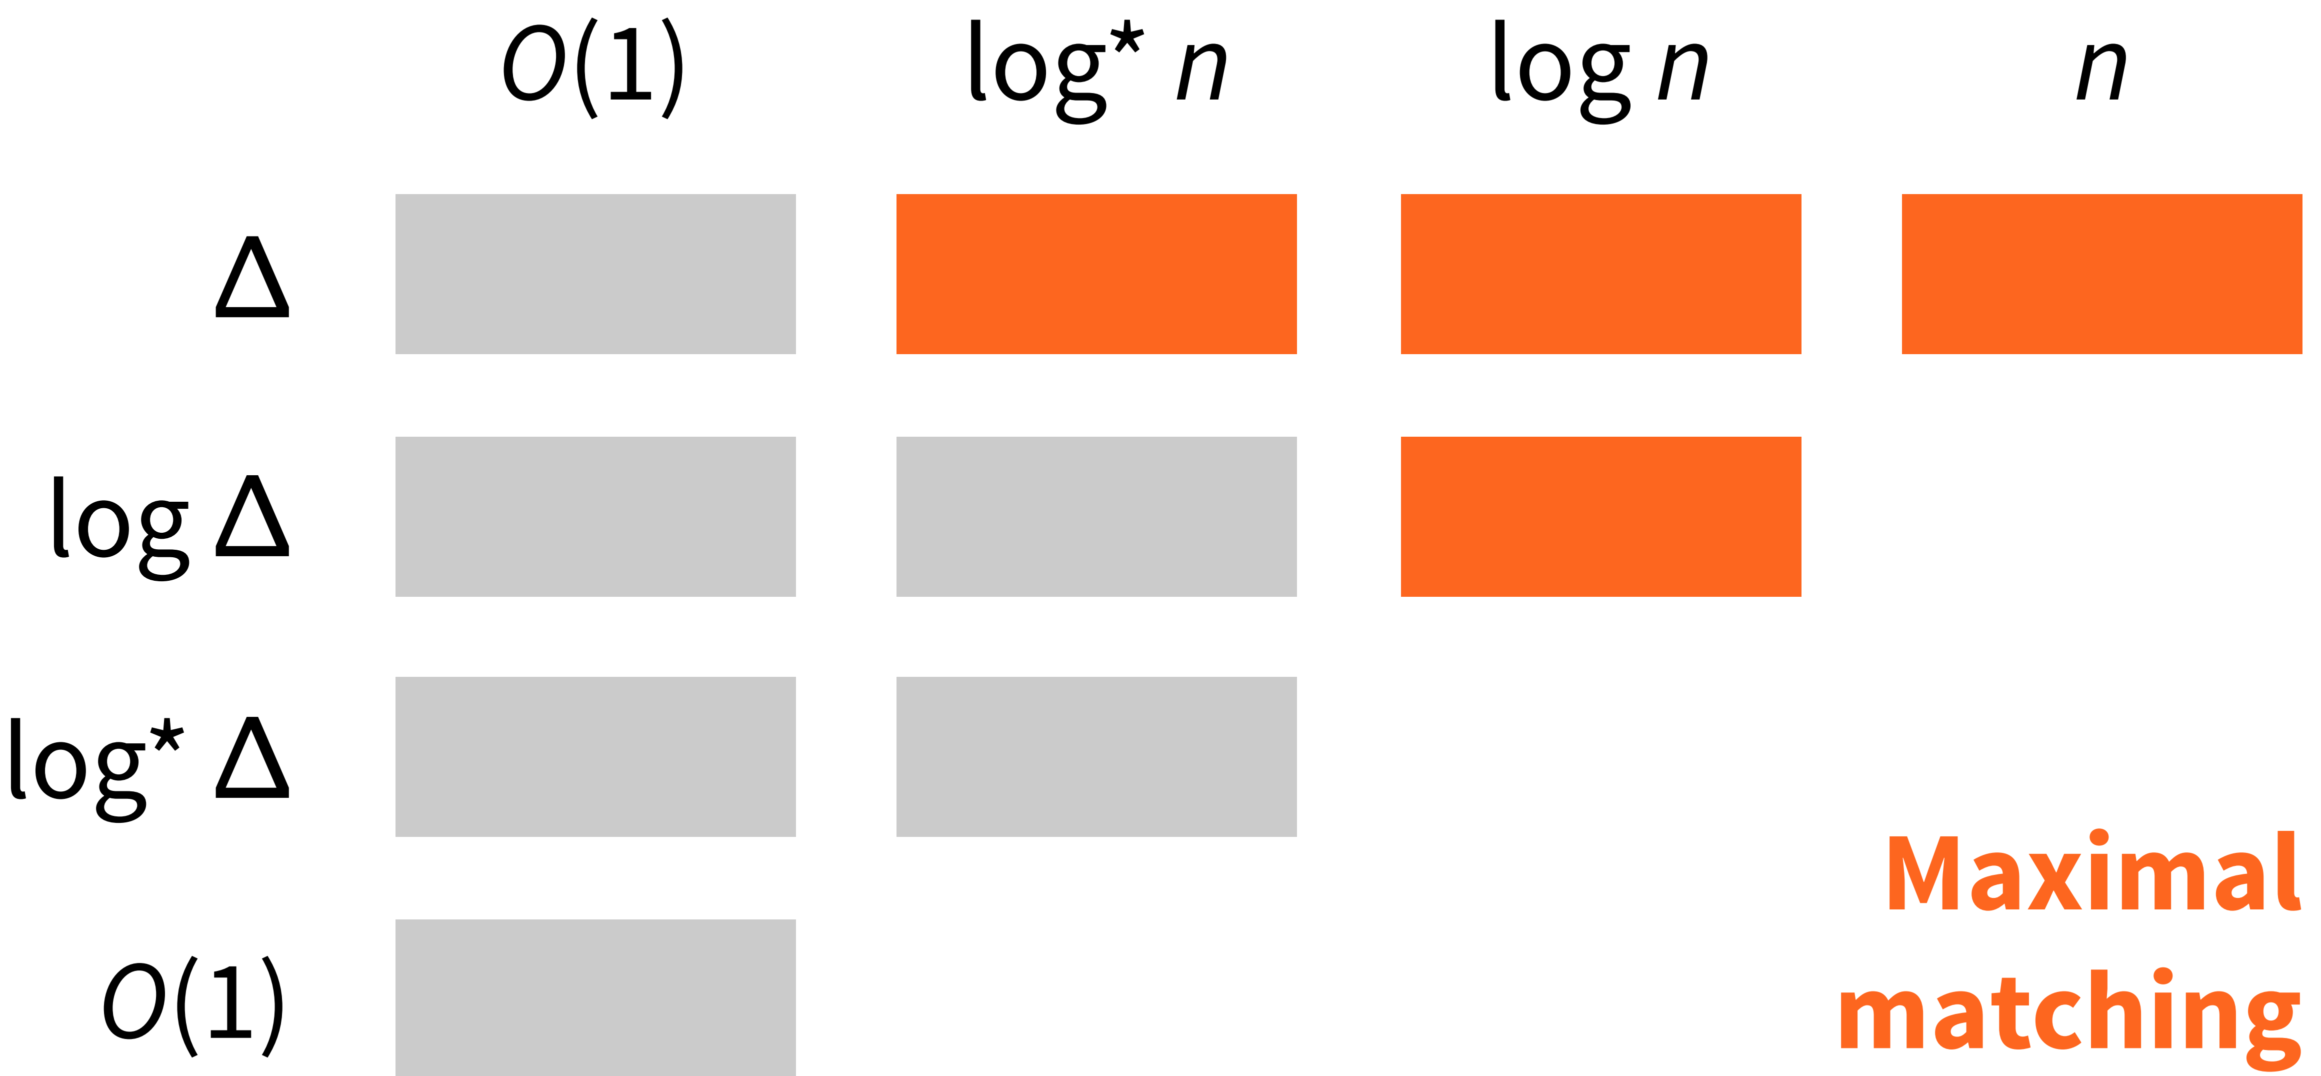

Landscape

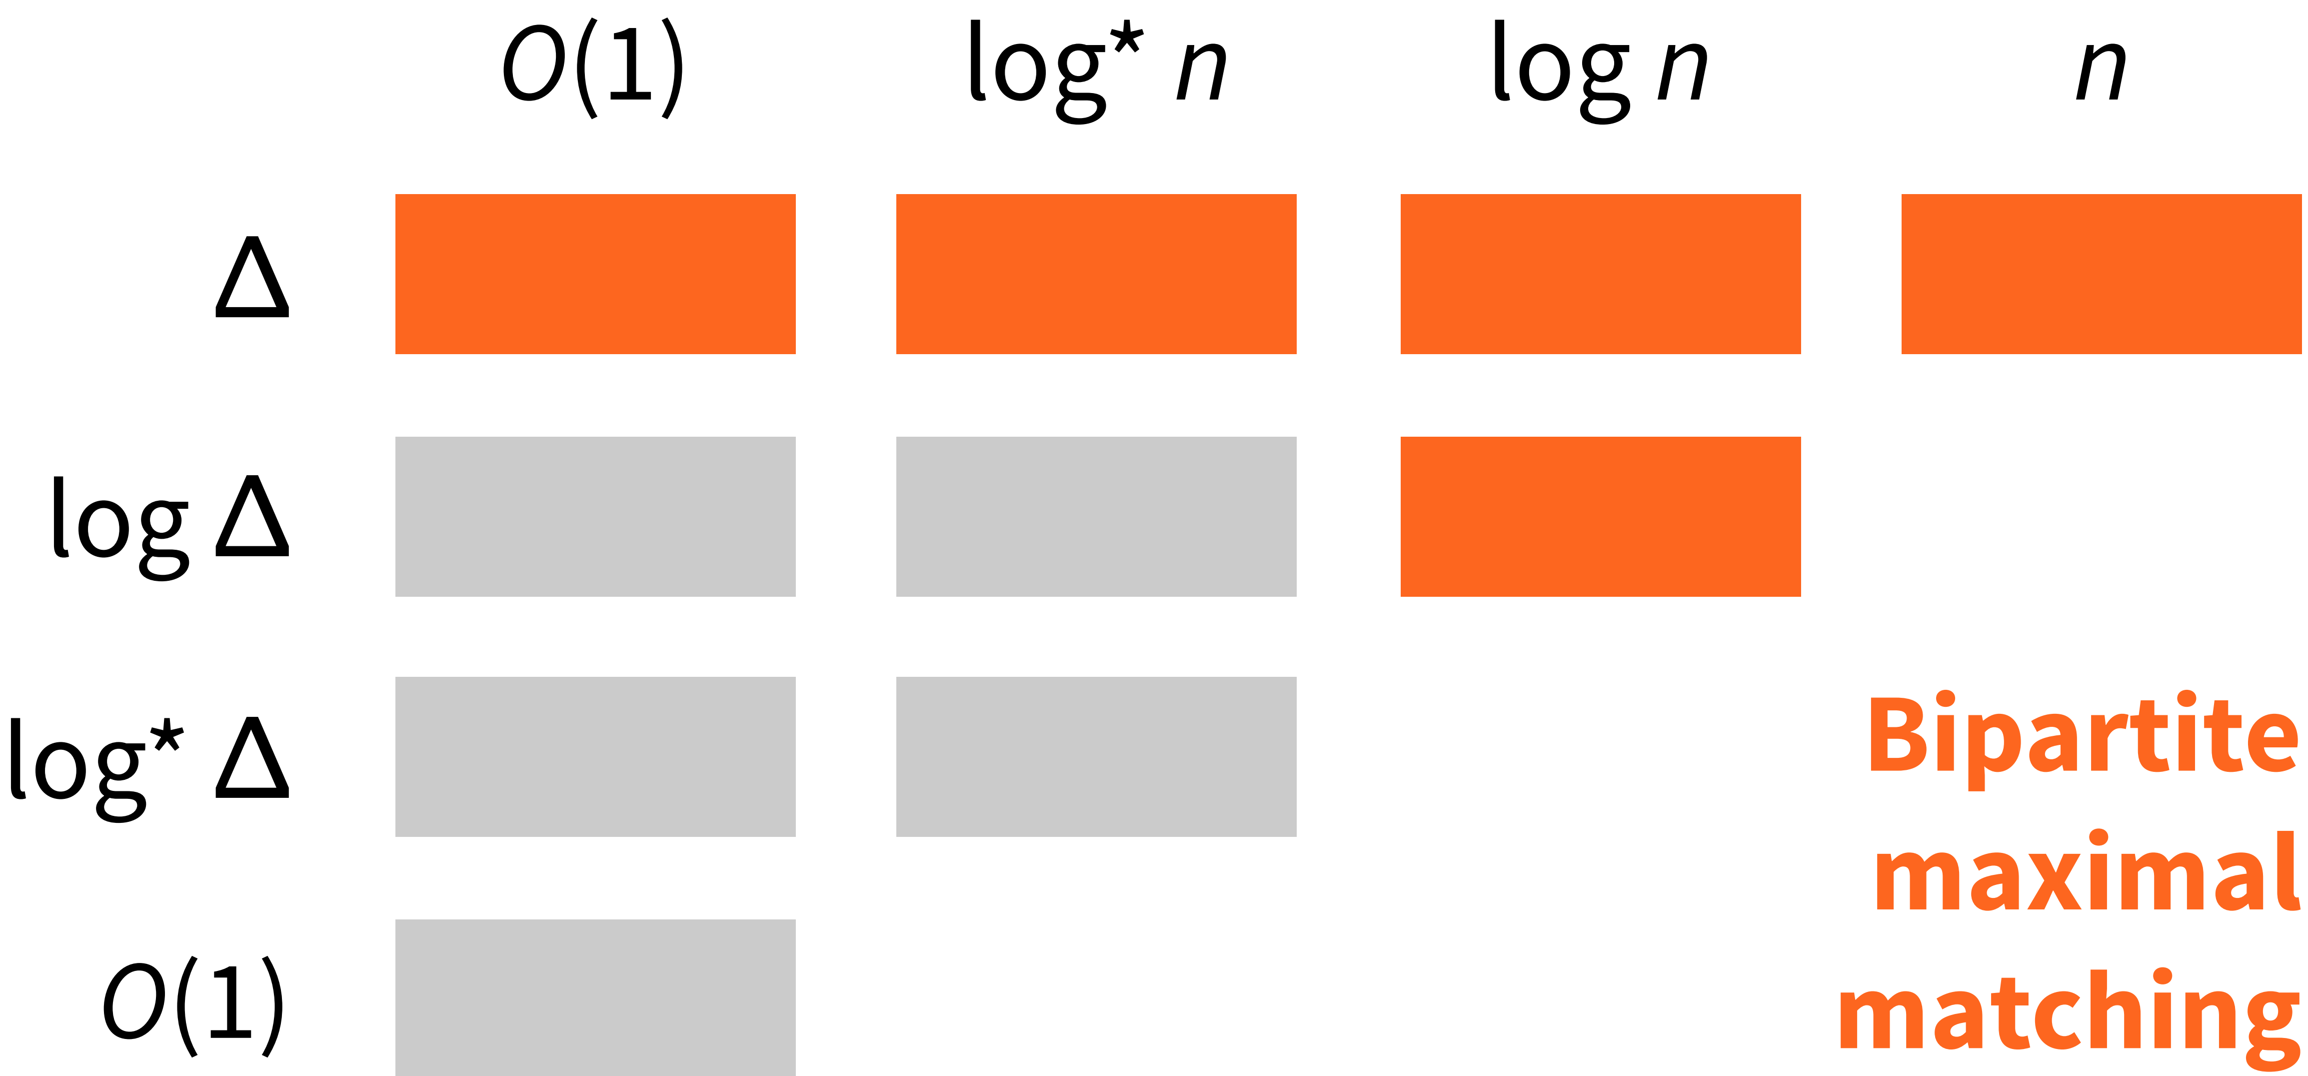
$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
Bipartite maximal matching

Landscape

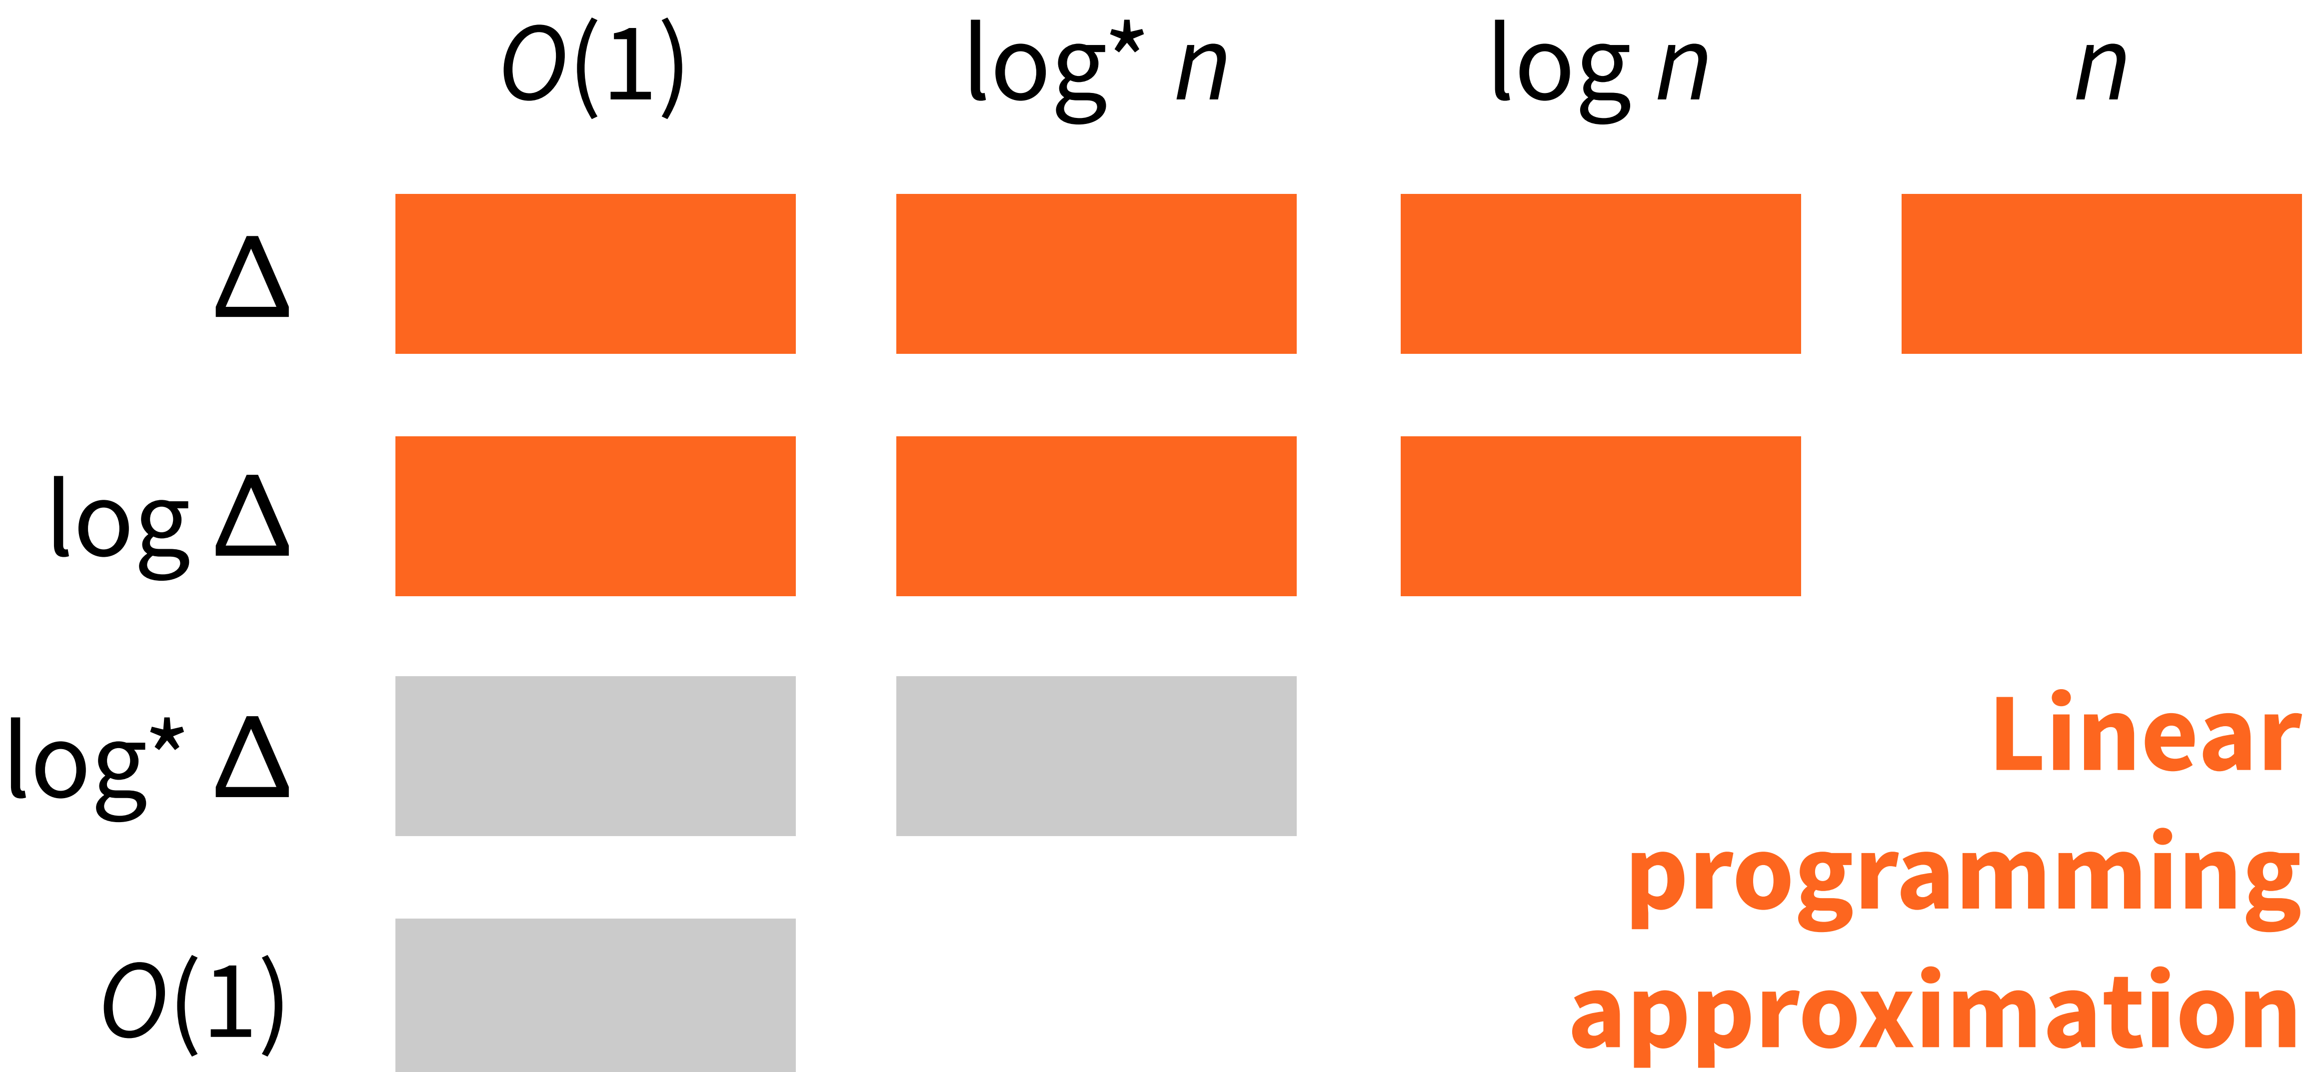
O(1)
log* n
log n
n
Δ
log Δ
log* Δ
O(1)
Linear
programming
approximation

Landscape

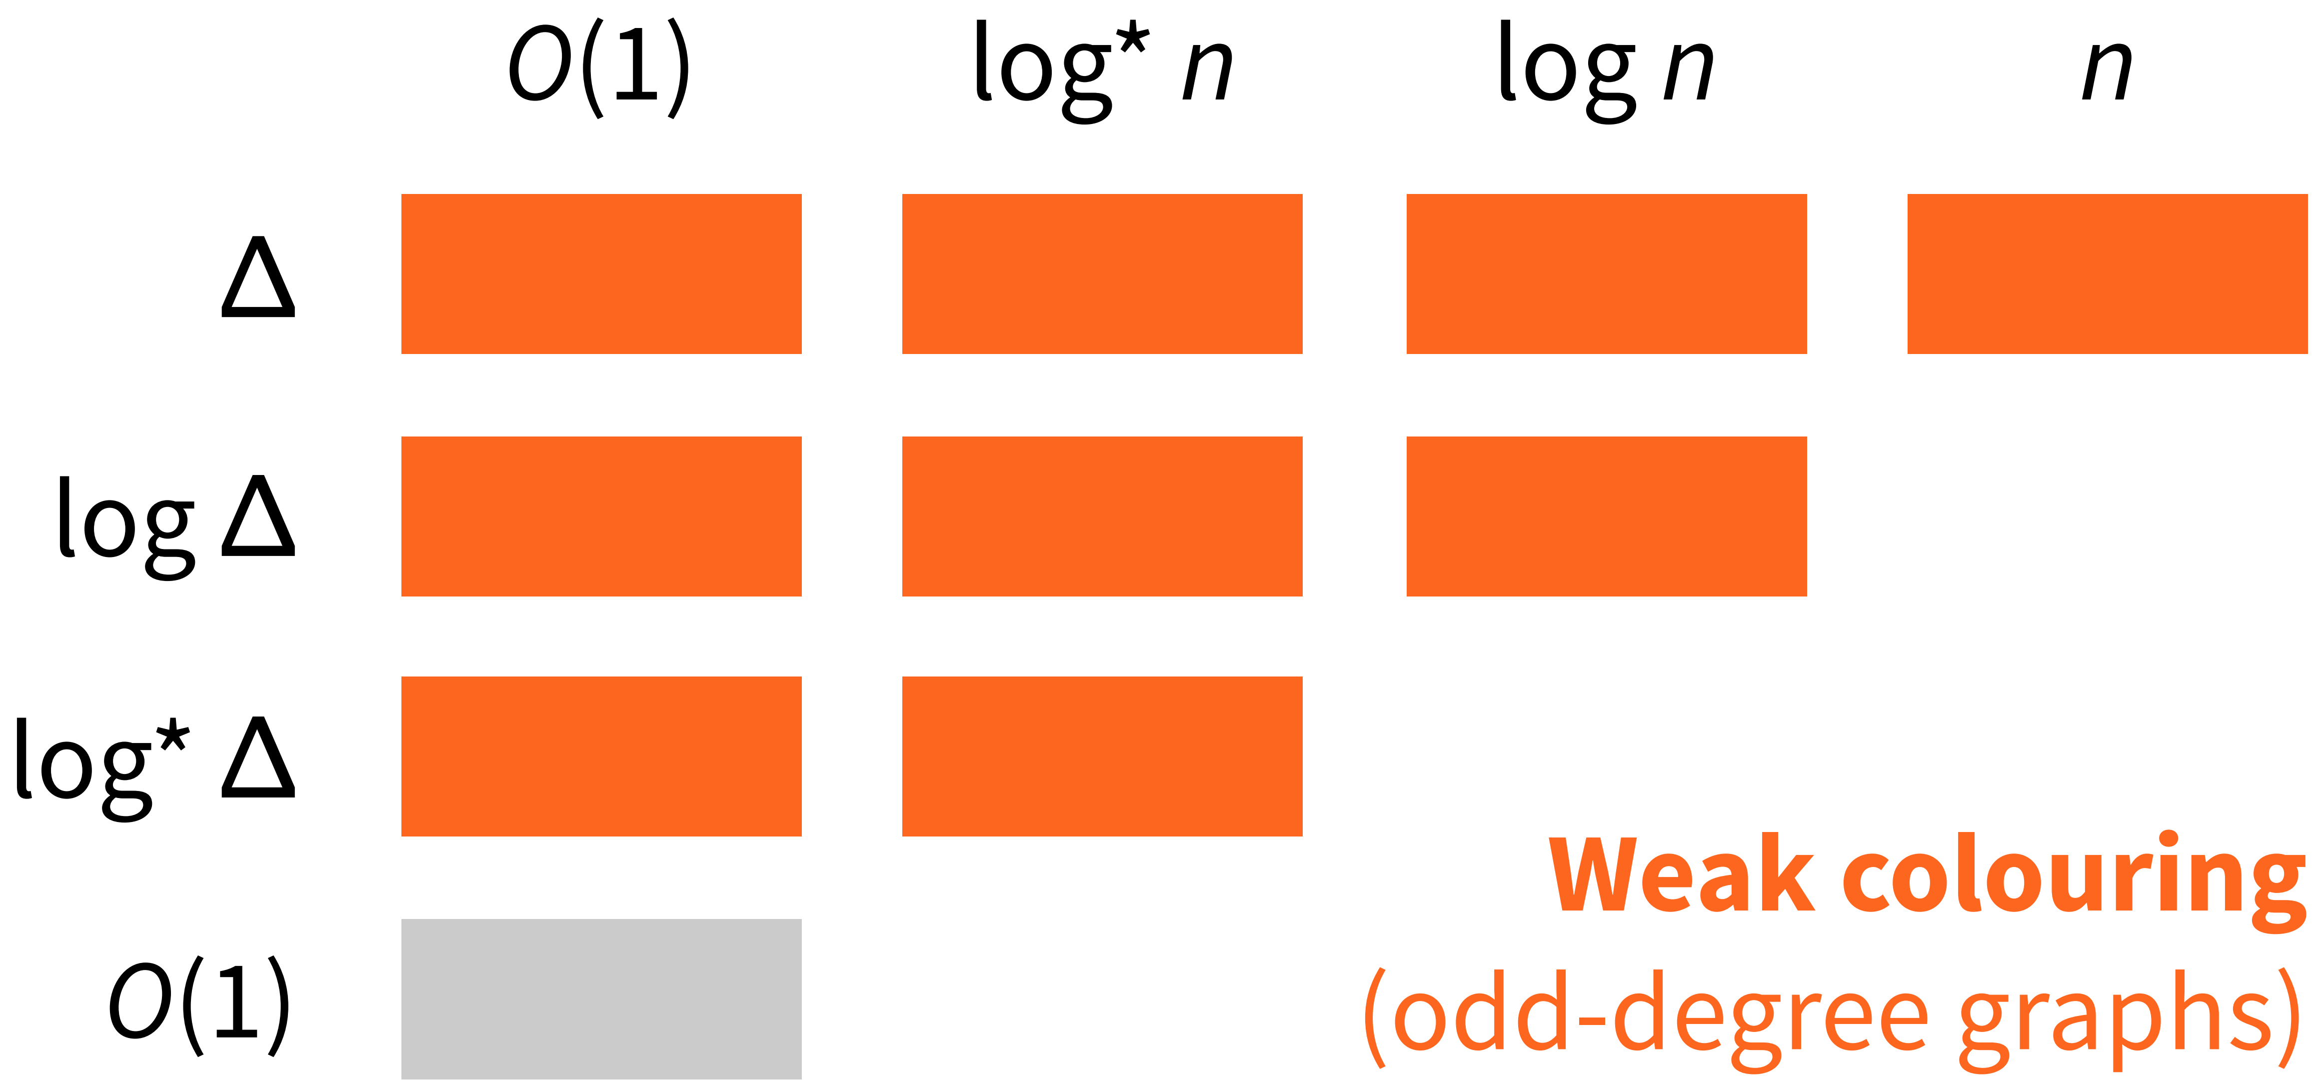

$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
Weak colouring
(odd-degree graphs)

Landscape

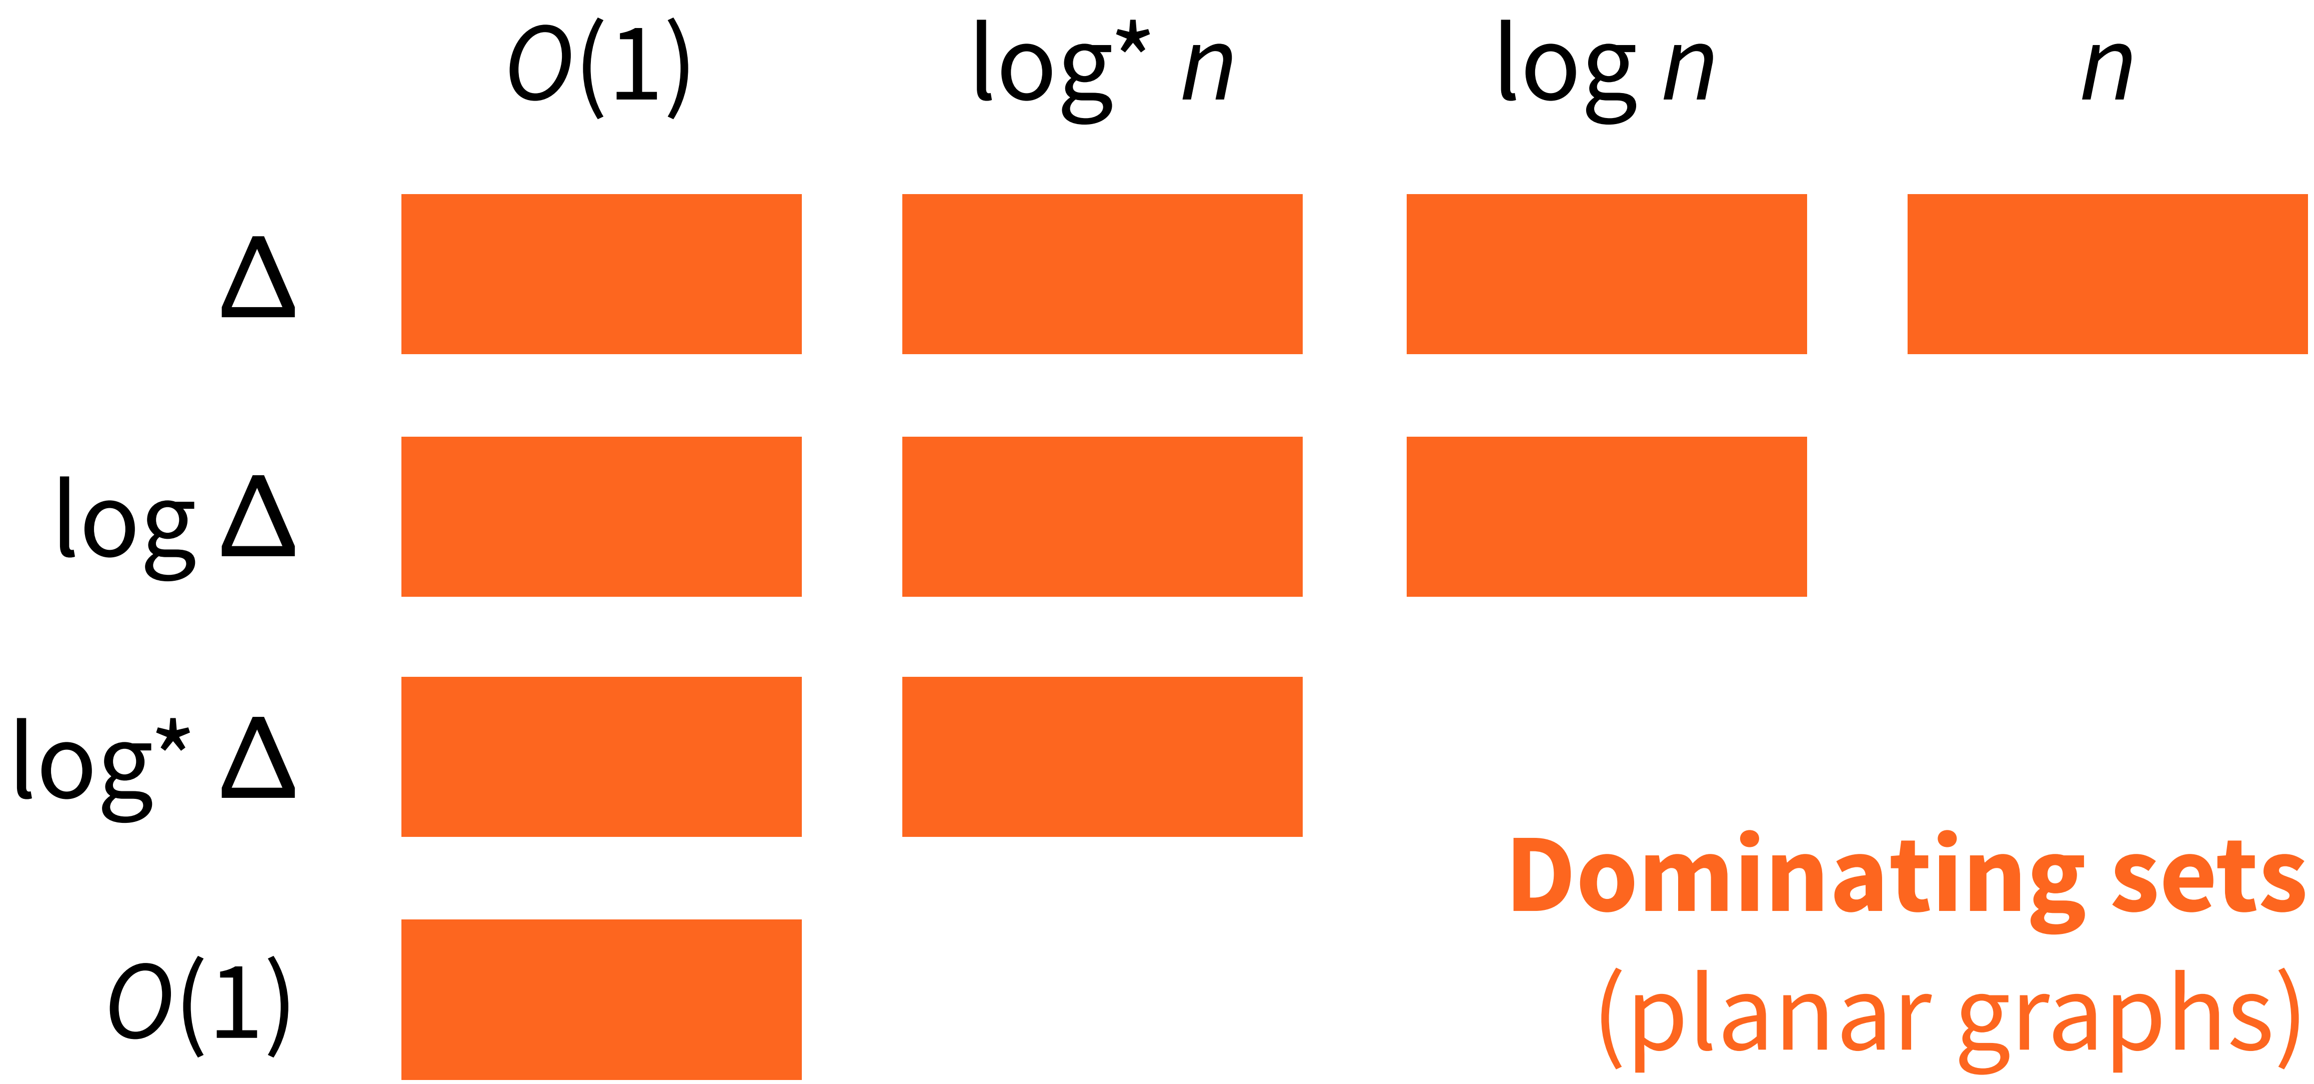

Dominating sets
(planar graphs)

Landscape

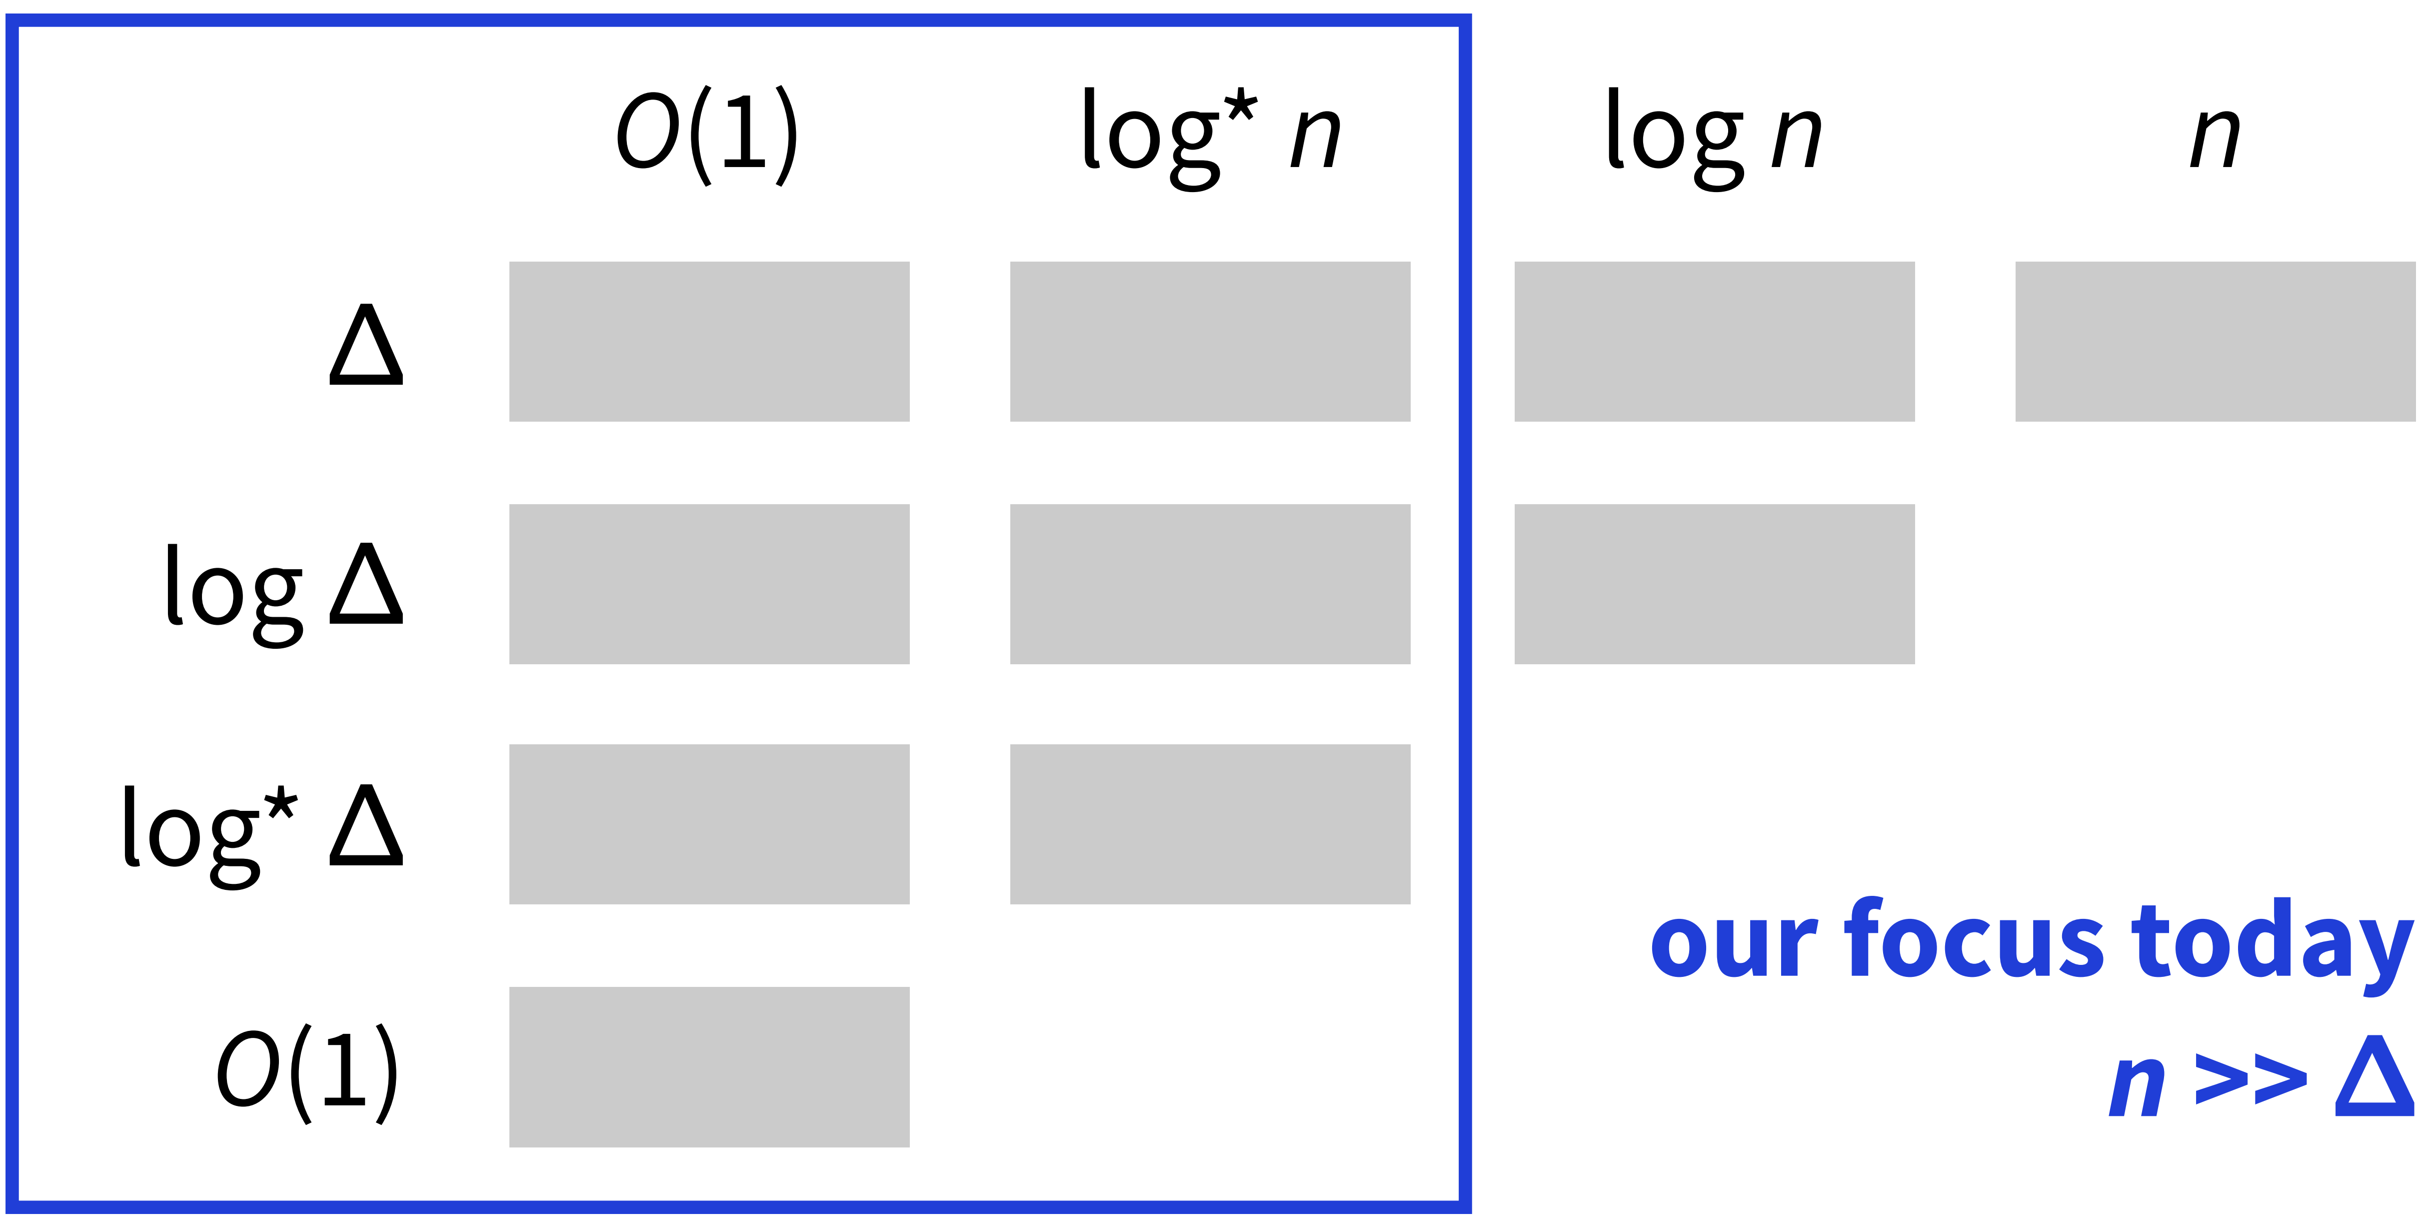

our focus today
$n >> \Delta$

Typical state of the art

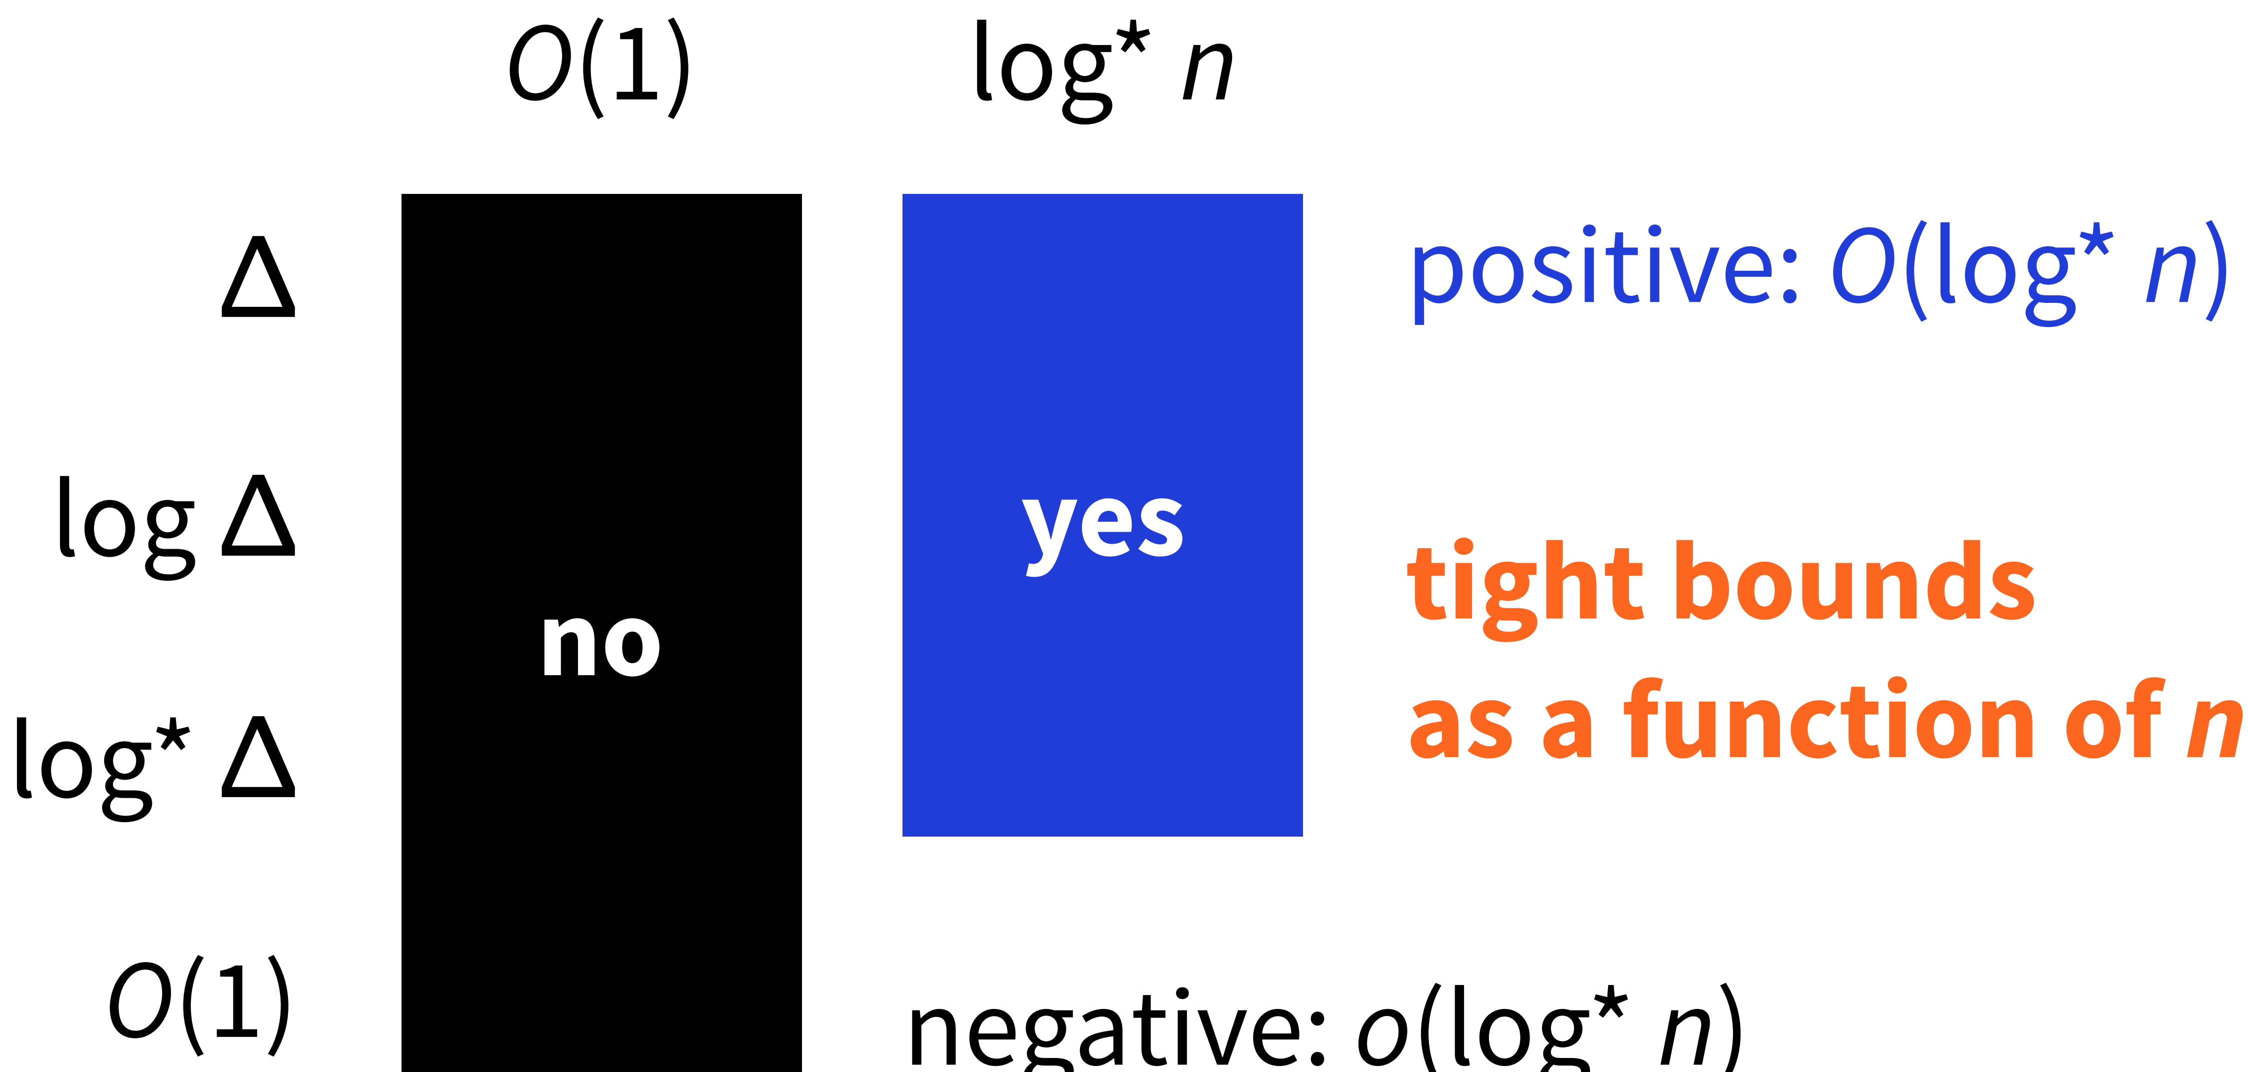

Typical state of the art

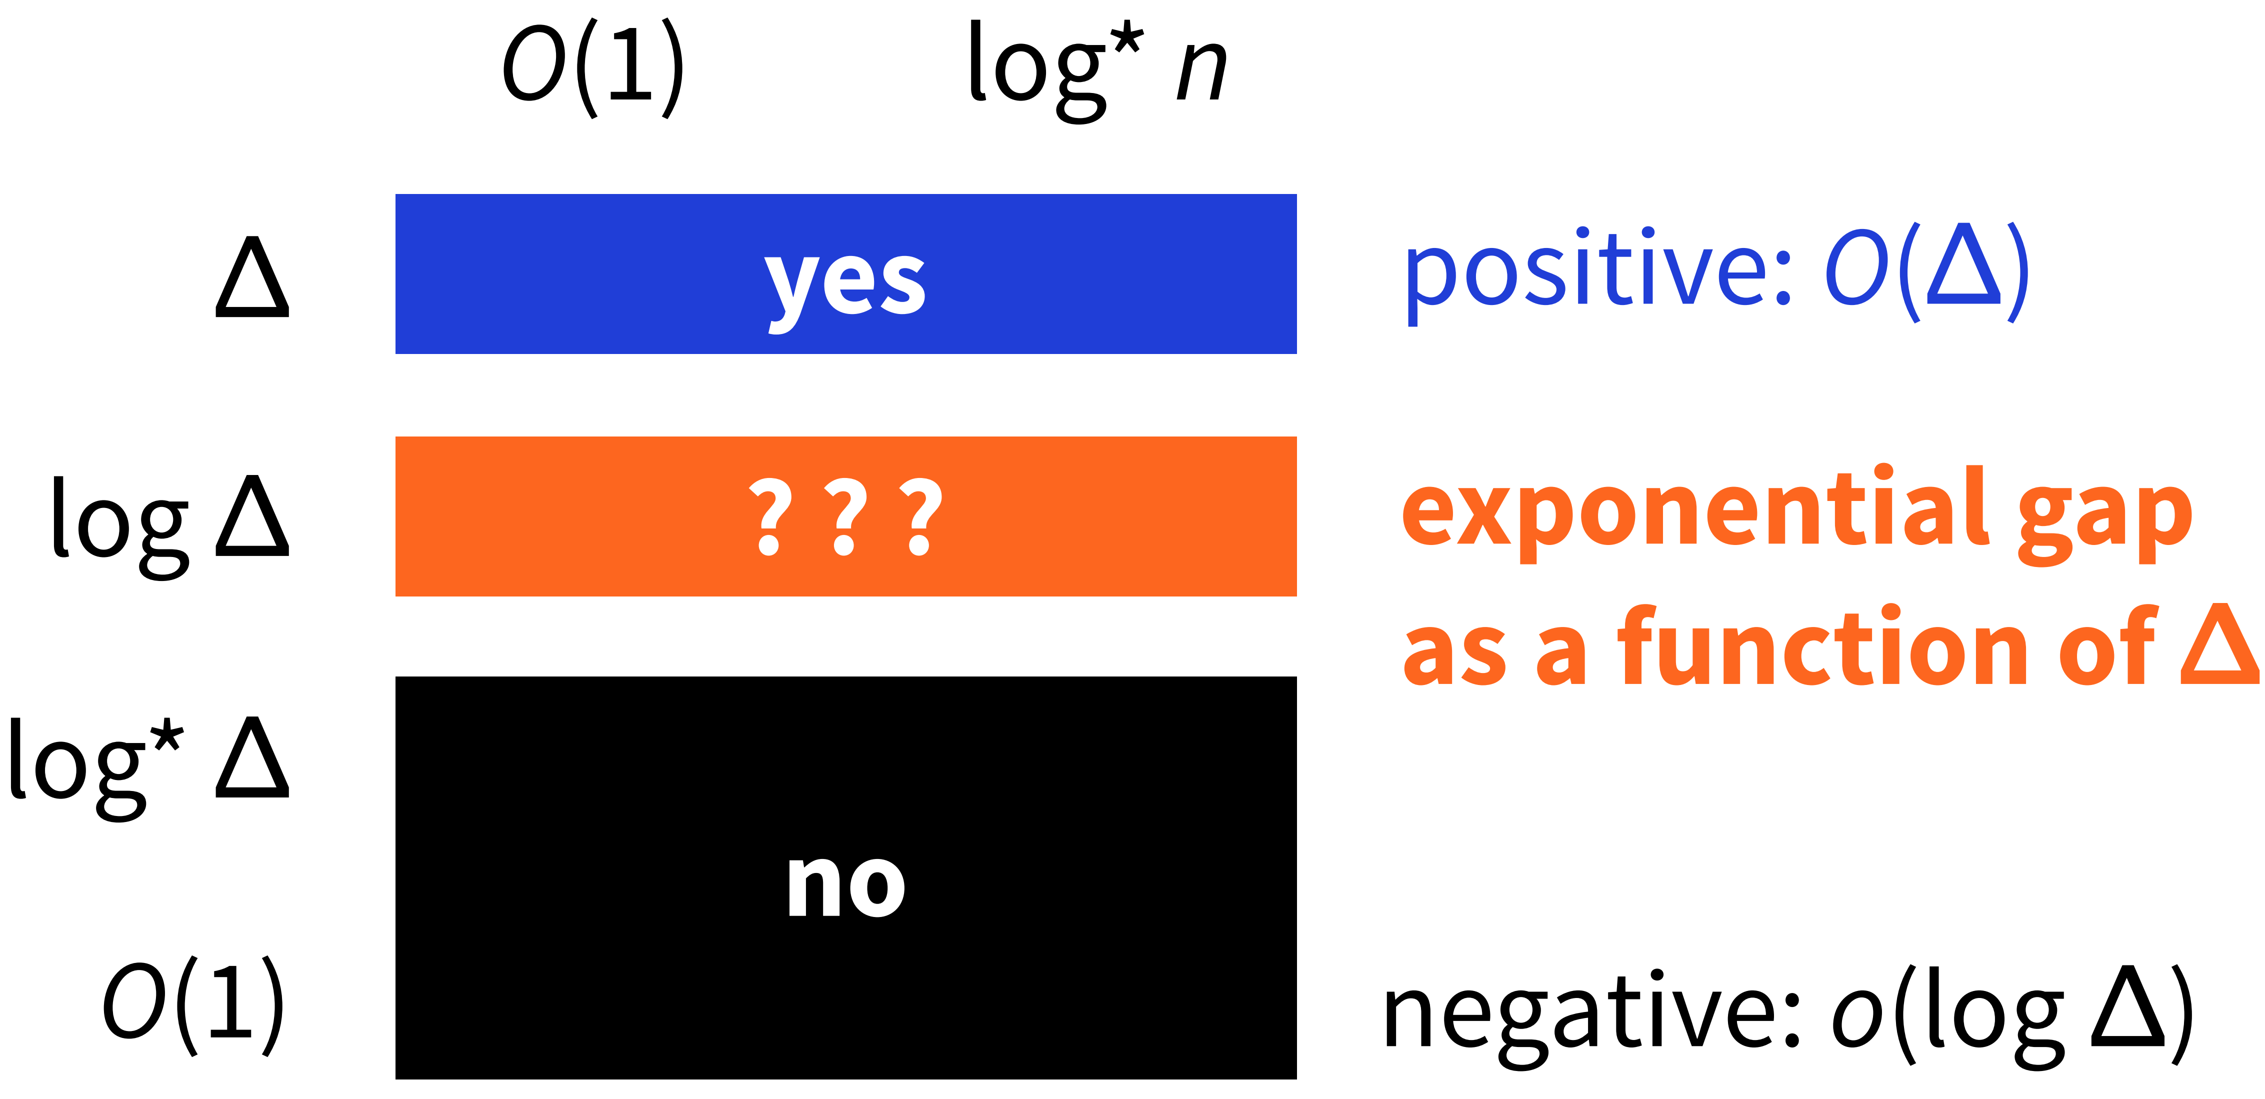

Typical state of the art

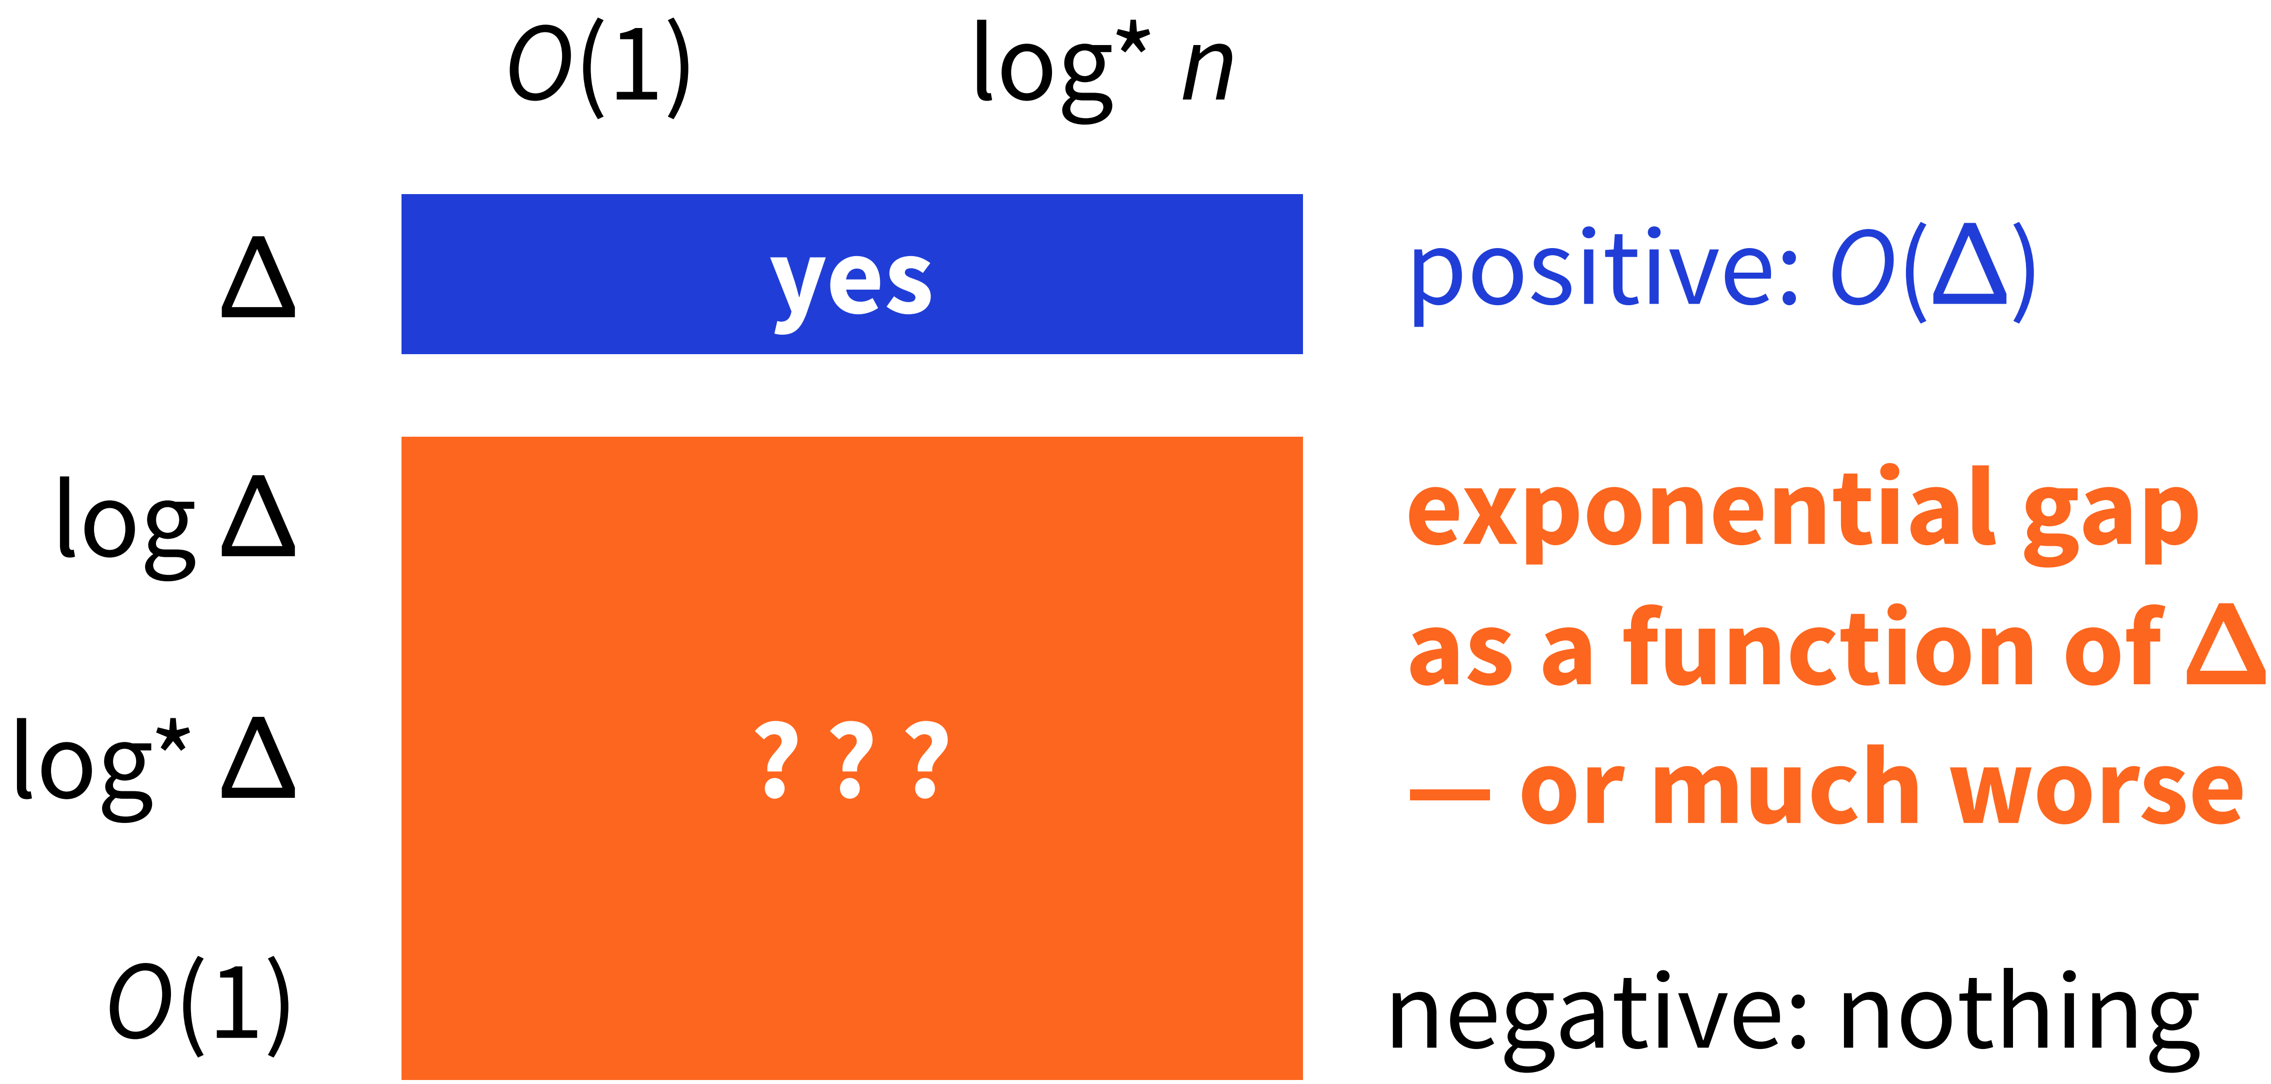

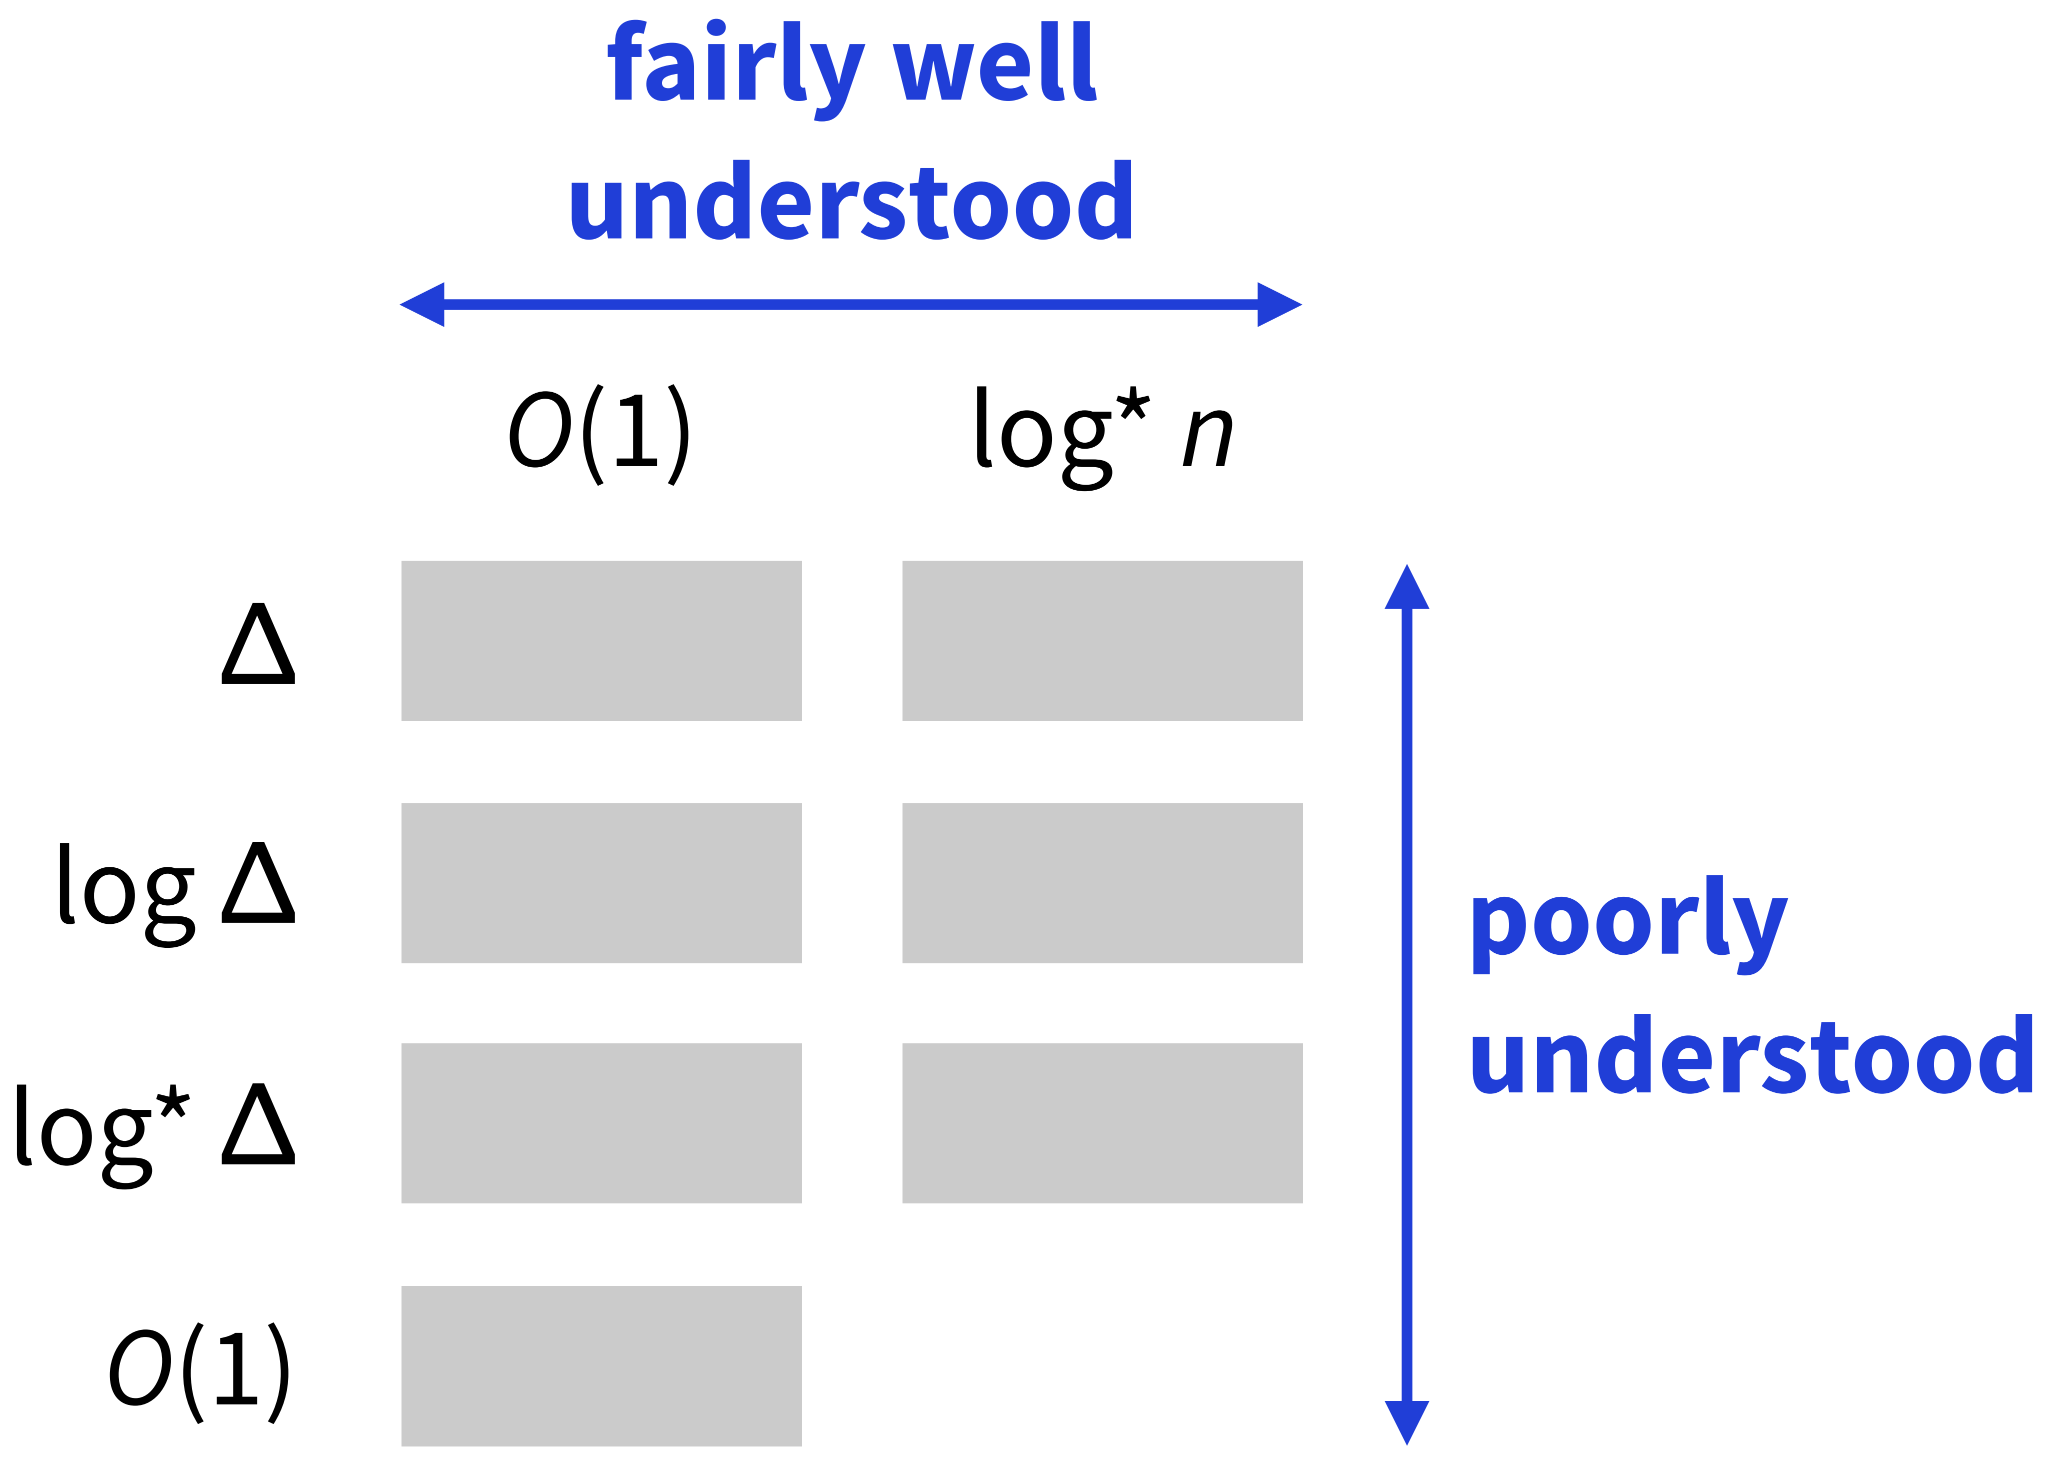

fairly well understood
$O(1)$
$\log^* n$
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
poorly understood

Example:
LP approximation

- **$O(\log \Delta)$: possible**
  - Kuhn et al. (2004, 2006)
- **$o(\log \Delta)$: not possible**
  - Kuhn et al. (2004, 2006)

Example:
Maximal matching

- **$O(\Delta + \log^* n)$: possible**
  - Panconesi & Rizzi (2001)
- **$O(\Delta) + o(\log^* n)$: not possible**
  - Linial (1992)
- **$o(\Delta) + O(\log^* n)$: unknown**

Example:
(Δ+1)-colouring

- **$O(\Delta + \log^* n)$: possible**
  - Barenboim & Elkin (2008), Kuhn (2008)
- **$O(\Delta) + o(\log^* n)$: not possible**
  - Linial (1992)
- **$o(\Delta) + O(\log^* n)$: unknown**

Example: **Bipartite maximal matching**

- ***O*(Δ): trivial**
  - Hańćkowiak et al. (1998)
- ***o*(Δ): unknown**

Example:
Semi-matching

- ***O*(Δ⁵): possible**
  - Czygrinow et al. (2012)
- ***o*(Δ⁵): unknown**

Example:
Semi-matching

- ***O*(Δ^5^): possible**
  - Czygrinow et al. (2012)
- ***o*(Δ^5^): unknown**
- ***o*(Δ): unknown**

Example:

Weak colouring

- ***O*(log* Δ): possible** (in odd-degree graphs)
  - Naor & Stockmeyer (1995)
- ***o*(log* Δ): unknown**

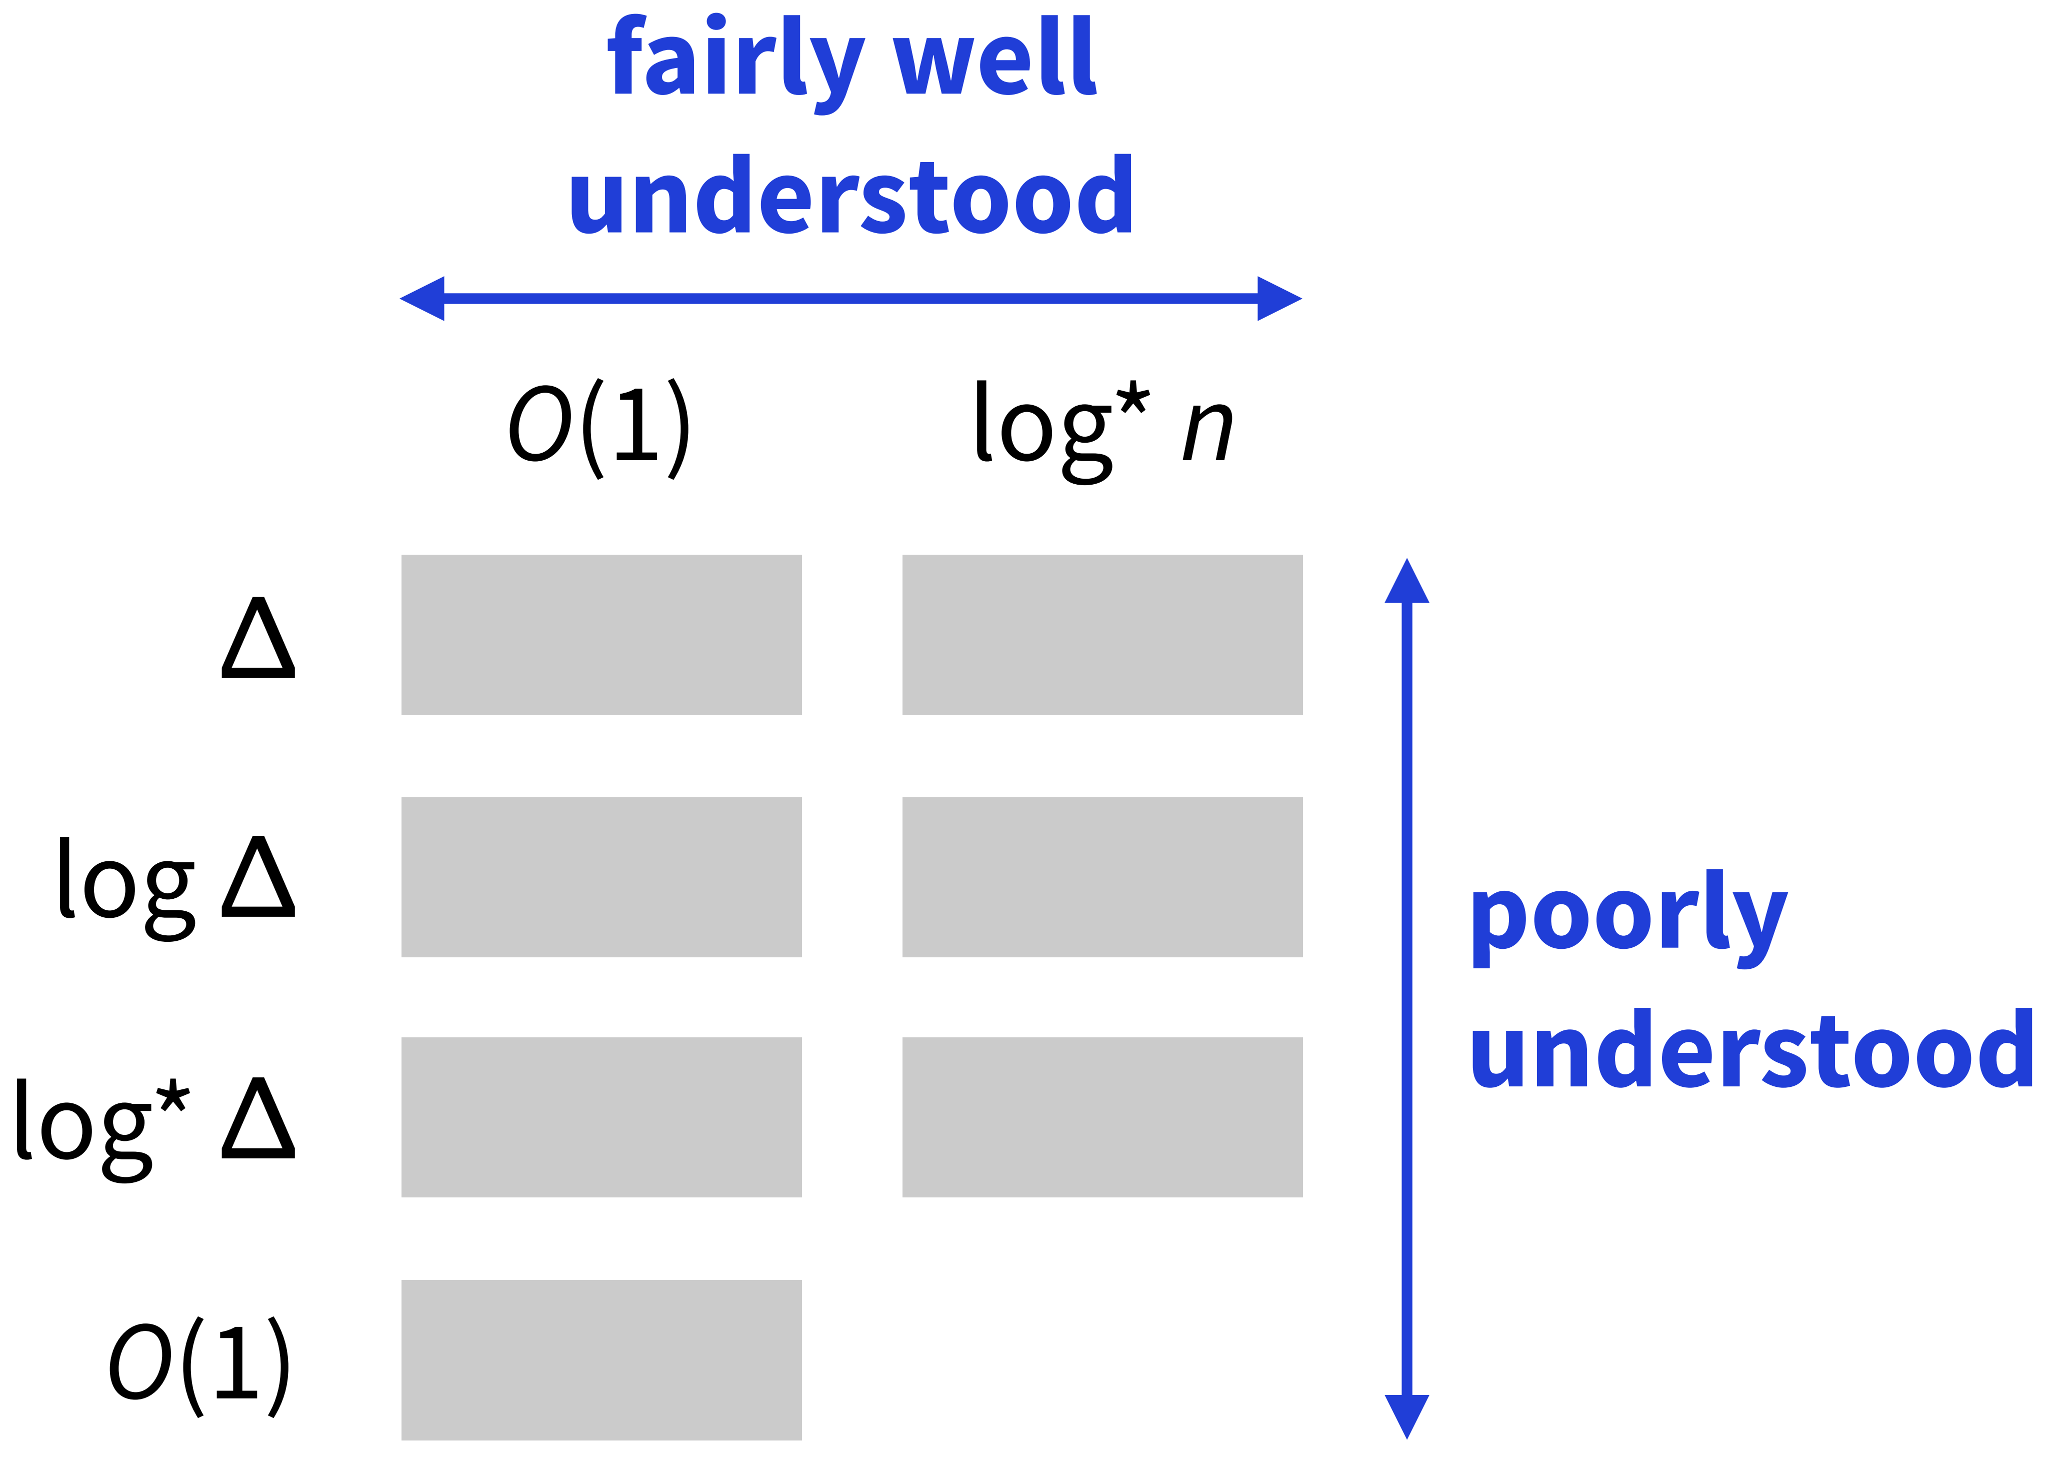
fairly well understood
$O(1)$
$\log^* n$
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
poorly understood

Orthogonal challenges?

- **n: “symmetry breaking”**
  - fairly well understood
  - Cole & Vishkin (1986), Linial (1992), Ramsey theory …
- **Δ: “local coordination”**
  - poorly understood

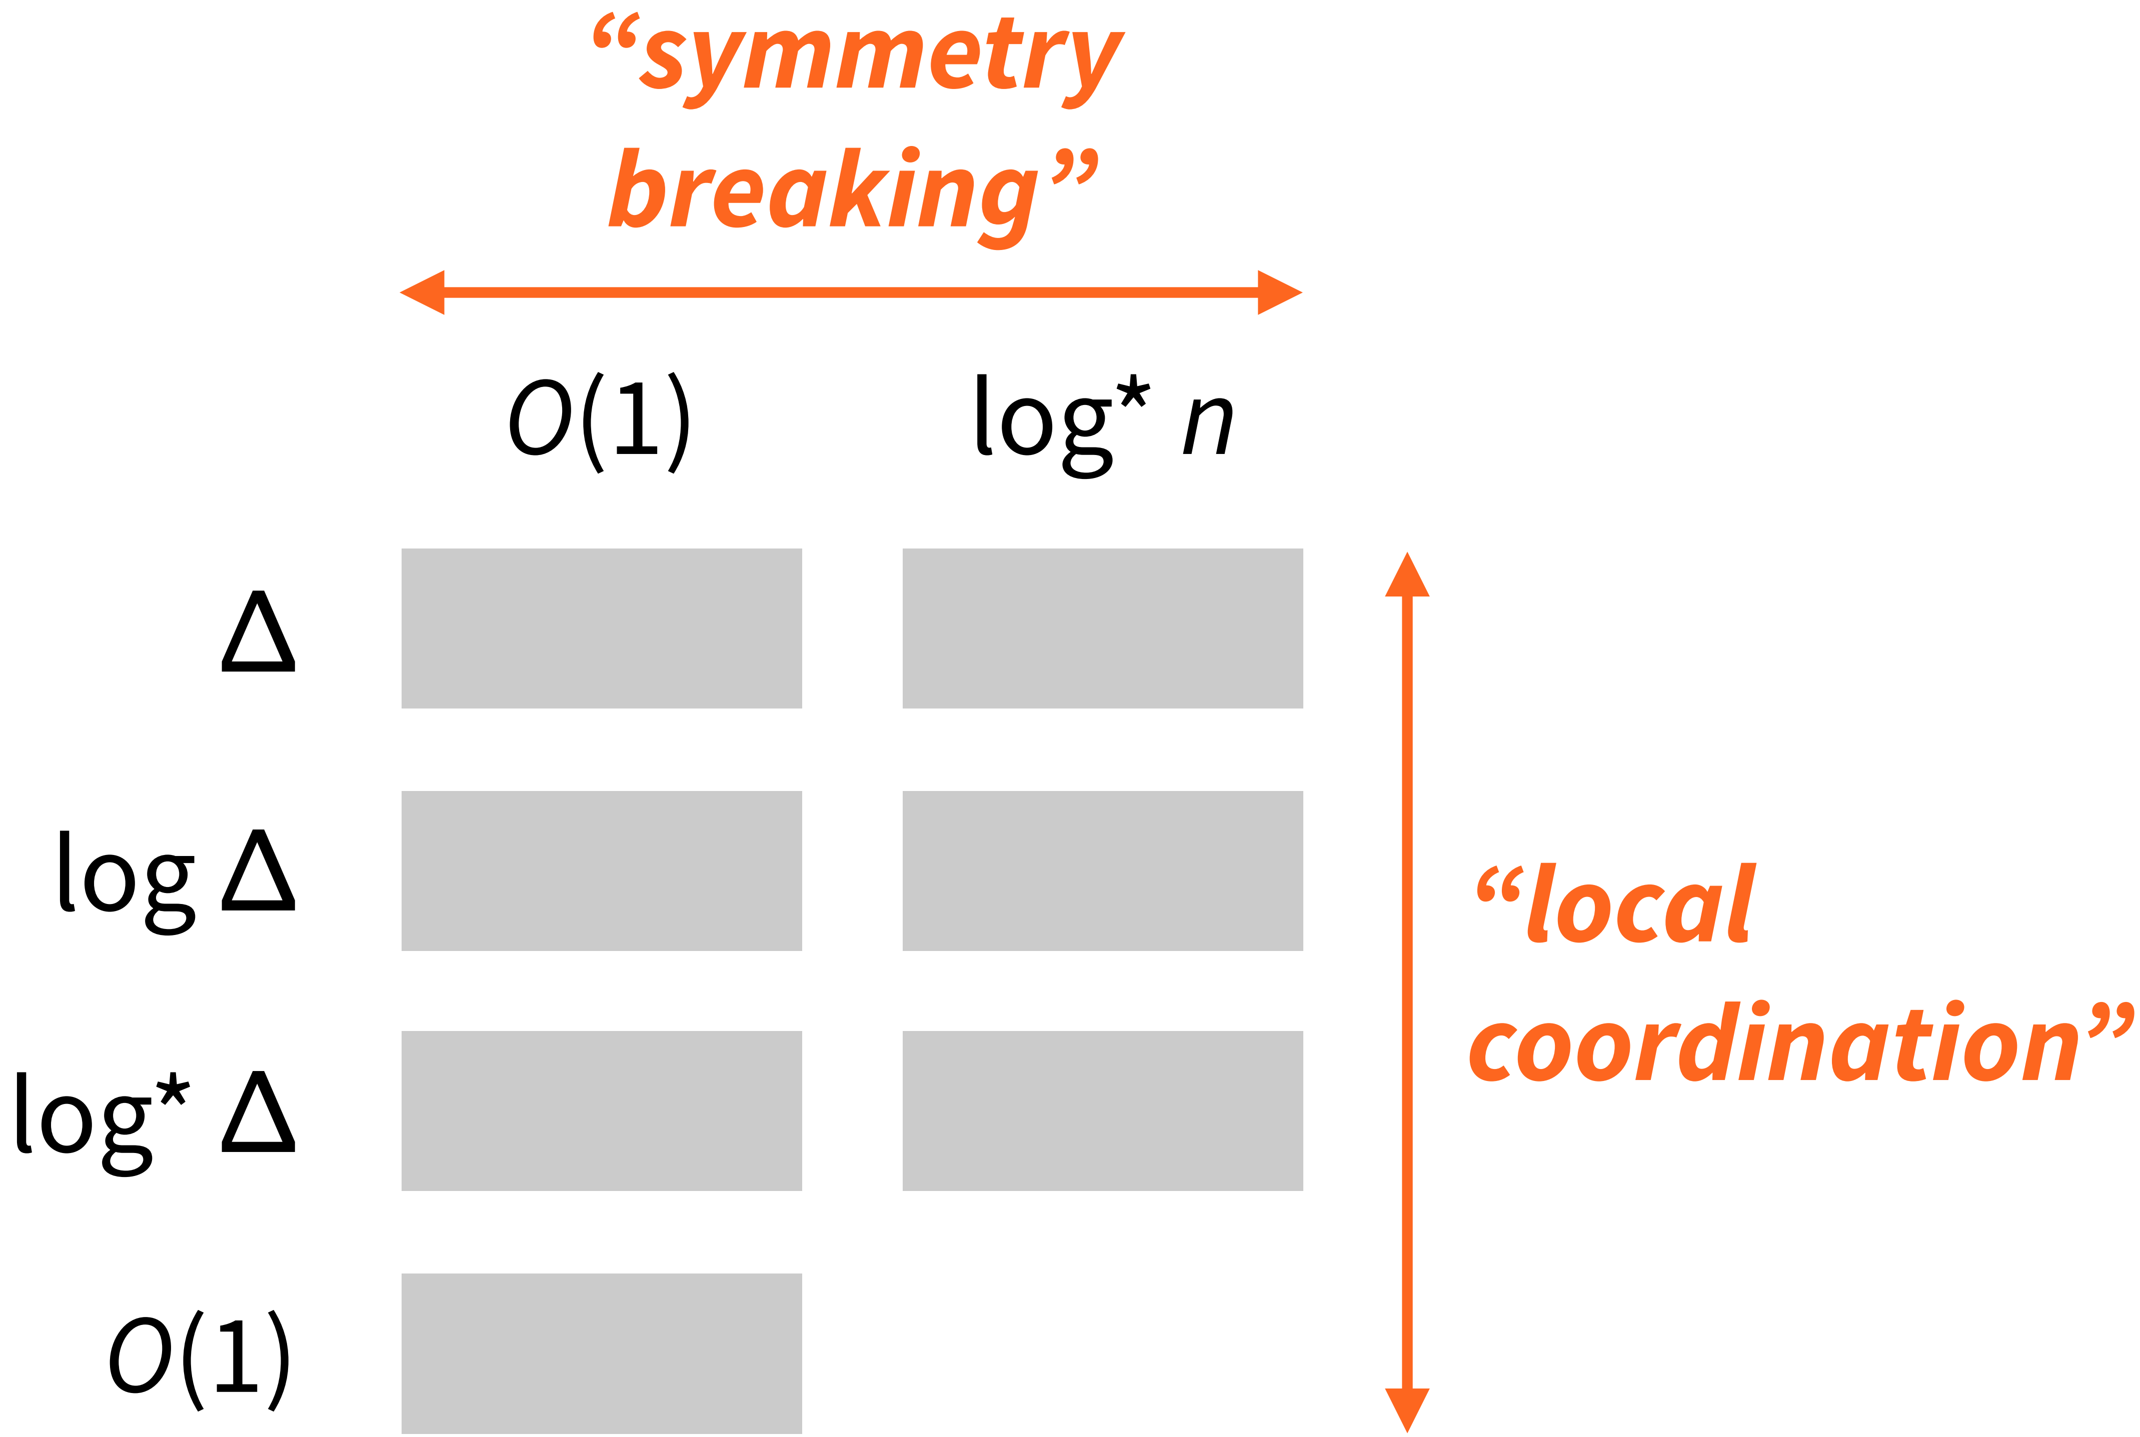
"symmetry breaking"
O(1)
log* n
Δ
log Δ
log* Δ
O(1)
"local coordination"

Orthogonal challenges

- **Example: maximal matching, $O(\Delta + \log^* n)$**
- **Restricted versions:**
  - pure symmetry breaking, $O(\log^* n)$
  - pure local coordination, $O(\Delta)$

Orthogonal challenges

- **Example: maximal matching, $O(\Delta + \log^* n)$**
- **Pure symmetry breaking:**
  - input = cycle
  - no need for local coordination
  - $O(\log^* n)$ is possible and tight

Orthogonal challenges

- **Example: maximal matching, $O(\Delta + \log^* n)$**
- **Pure local coordination:**
  - input = 2-coloured graph
  - no need for symmetry breaking
  - $O(\Delta)$ is possible — is it tight?

Maximal matching in 2-coloured graphs

- **Trivial algorithm:**
  - black nodes send proposals to their neighbours, one by one
  - white nodes accept the first proposal that they get

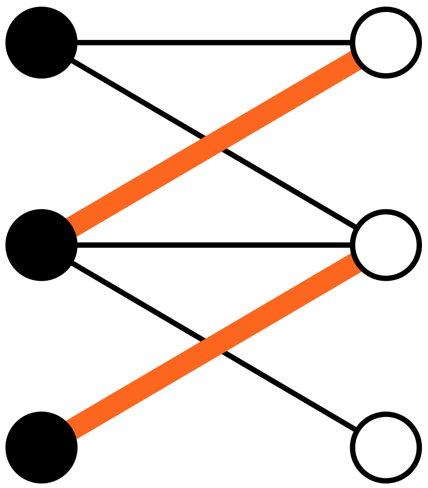

- **“Coordination” ≈ one by one traversal**

Maximal matching in 2-coloured graphs

- **Trivial algorithm:**
  - black nodes send proposals to their neighbours, one by one
  - white nodes accept the first proposal that they get

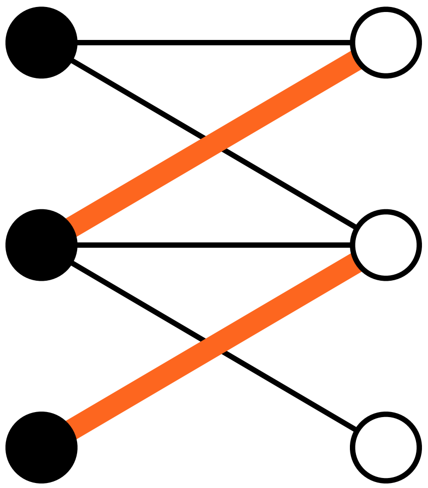

- **Clearly $O(\Delta)$, but is this tight?**

Maximal matching in 2-coloured graphs

- **General case:**
  - upper bound: $O(\Delta)$
  - lower bound: $\Omega(\log \Delta)$ — Kuhn et al.
- **Regular graphs:**
  - upper bound: $O(\Delta)$
  - lower bound: nothing!

Linear-in-Δ bounds

- **Many combinatorial problems seem to require "local coordination", takes $O(\Delta)$ time?**
- **Lacking: linear-in-Δ lower bounds**
  - known for restricted algorithm classes (Kuhn & Wattenhofer 2006)
  - not previously known for the LOCAL model

Recent progress

- **Maximal *fractional* matching**
- **$O(\Delta)$: possible**
  - SPAA 2010
- **$o(\Delta)$: not possible**
  - PODC 2014

Matching

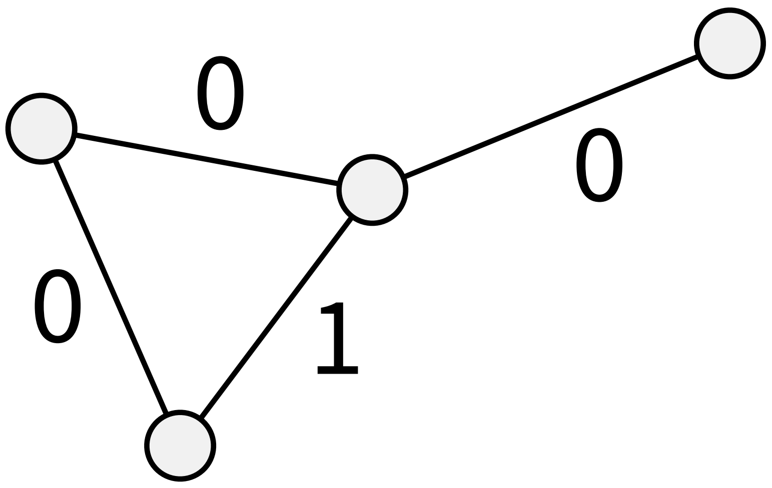

- **Edges labelled with integers {0, 1}**
- **Sum of incident edges at most 1**
- **Maximal matching:**
  cannot increase the value of any label

Fractional matching

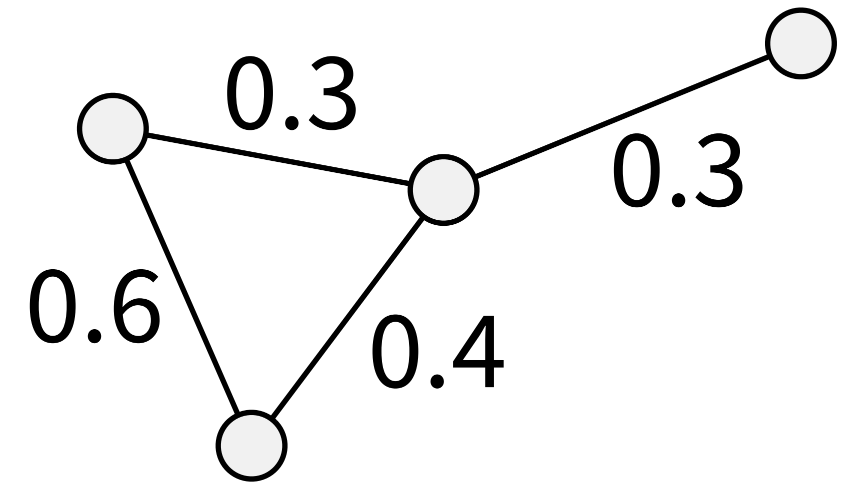

- **Edges labelled with real numbers [0, 1]**
- **Sum of incident edges at most 1**
- **Maximal fractional matching:**
  cannot increase the value of any label

Maximal fractional matching

- **Possible in time $O(\Delta)$**
  - does **not** require symmetry breaking
  - d-regular graph: label all edges with $1/d$
- **Nontrivial part:** graphs that are not regular…

Maximal fractional matching

- **Not possible in time $o(\Delta)$, independently of n**
  - note: we do not say anything about e.g. possibility of solving in time $o(\Delta) + O(\log^* n)$
- **Key ingredient of the proof:**
  analyse many **different models** of distributed computing

ID: unique identifiers

Nodes have unique identifiers,
output may depend on them

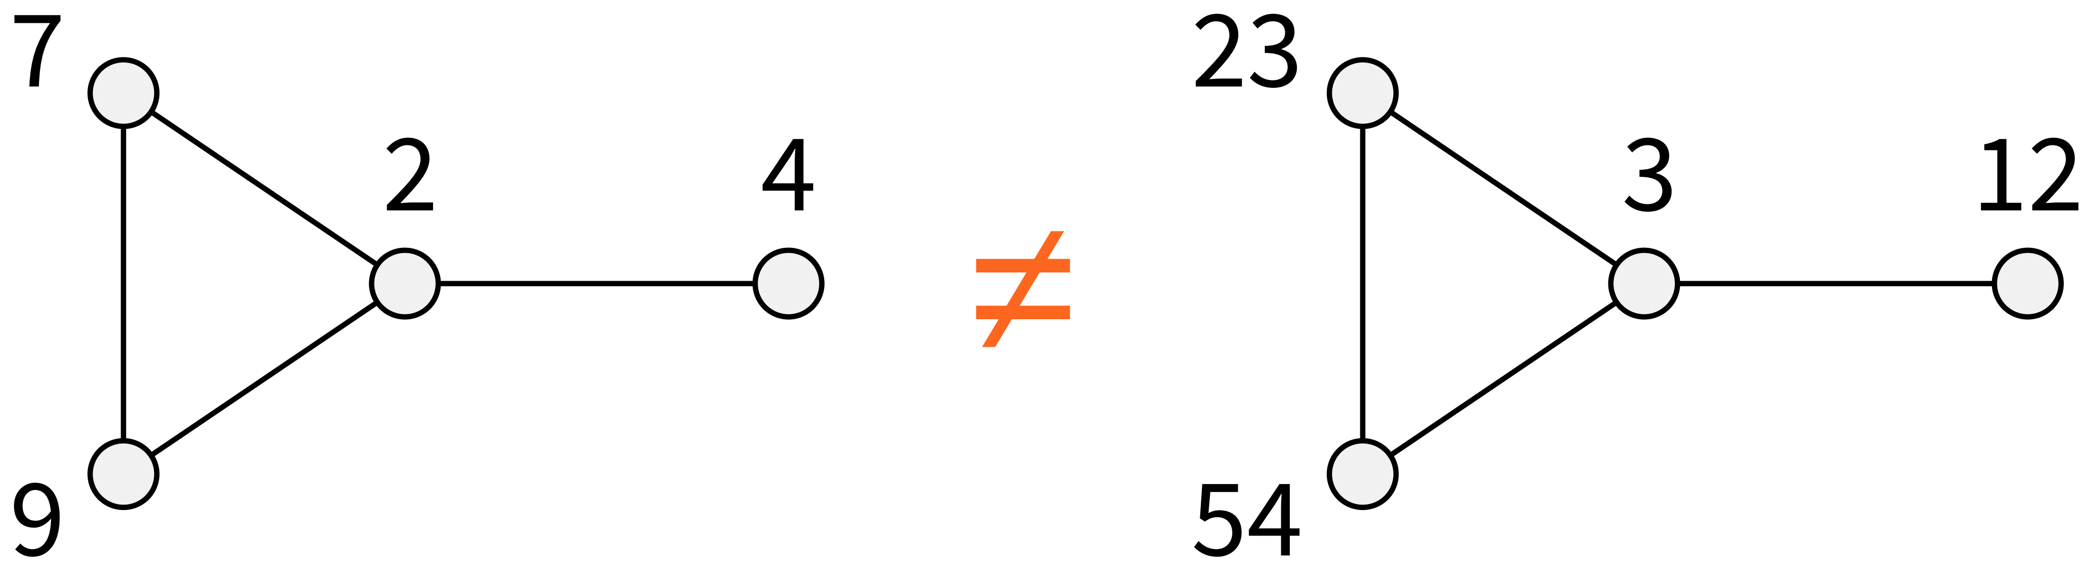

OI: order invariant

Output does not change if we change identifiers but keep their relative order

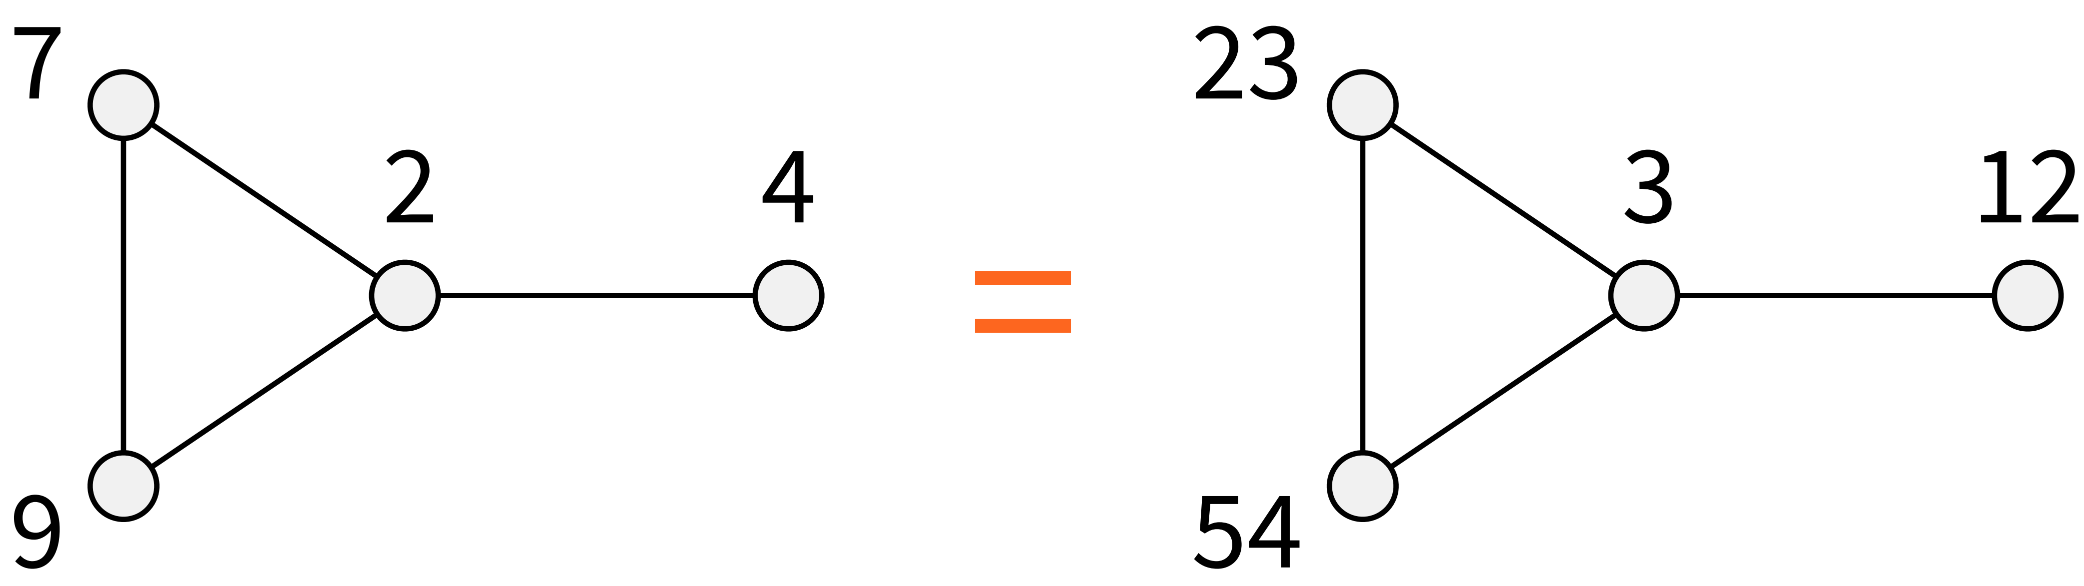

PO: ports & orientation

No identifiers

Node *v* labels incident edges with 1, ..., deg(*v*)

Edges oriented

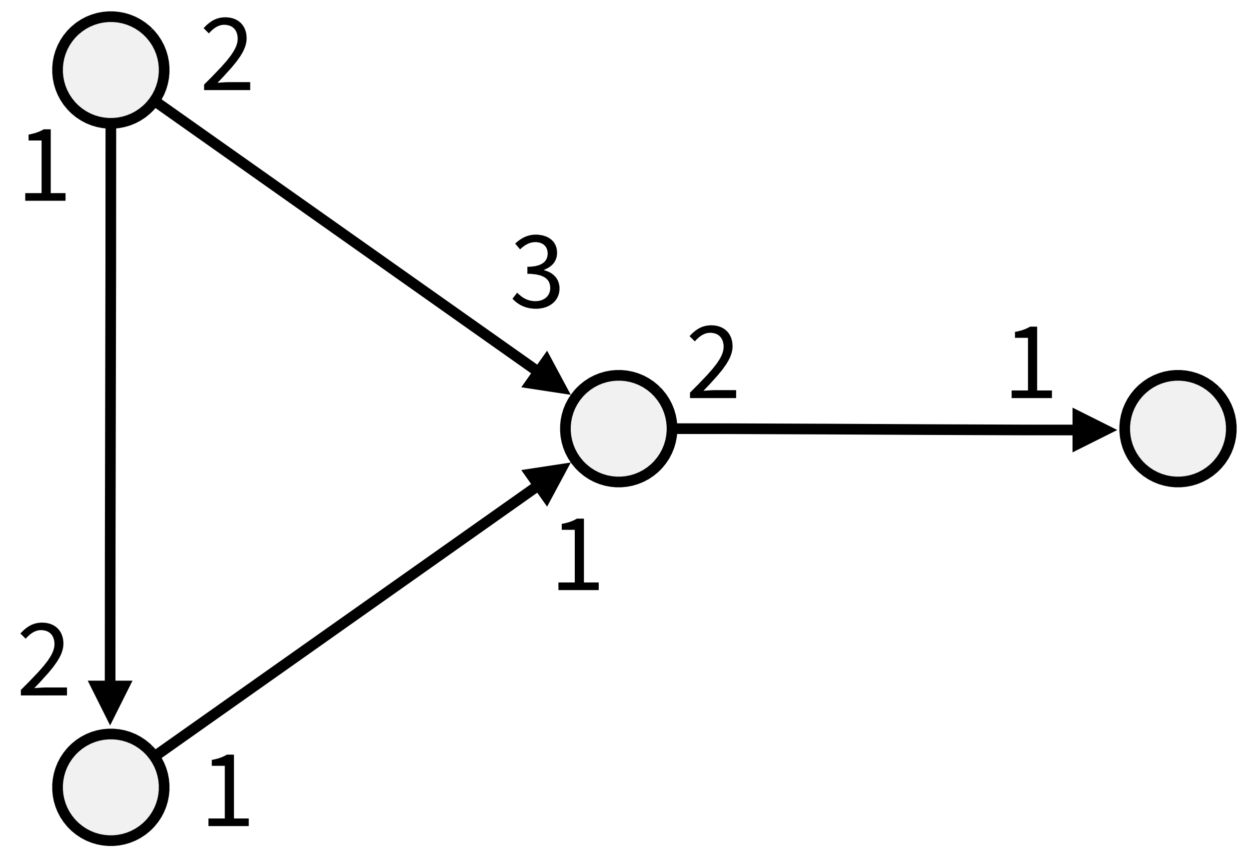

EC: edge colouring

No identifiers

No orientations

Edges coloured
with $O(\Delta)$ colours

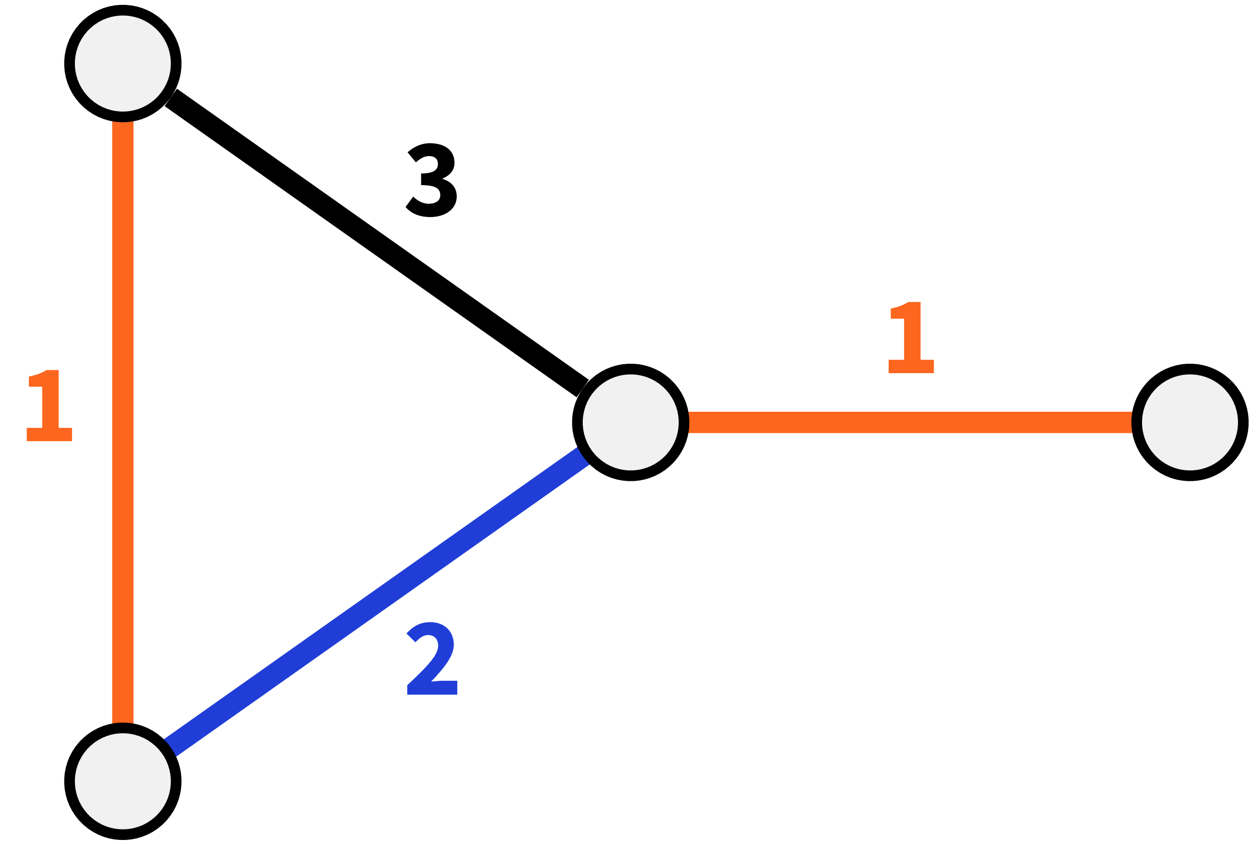

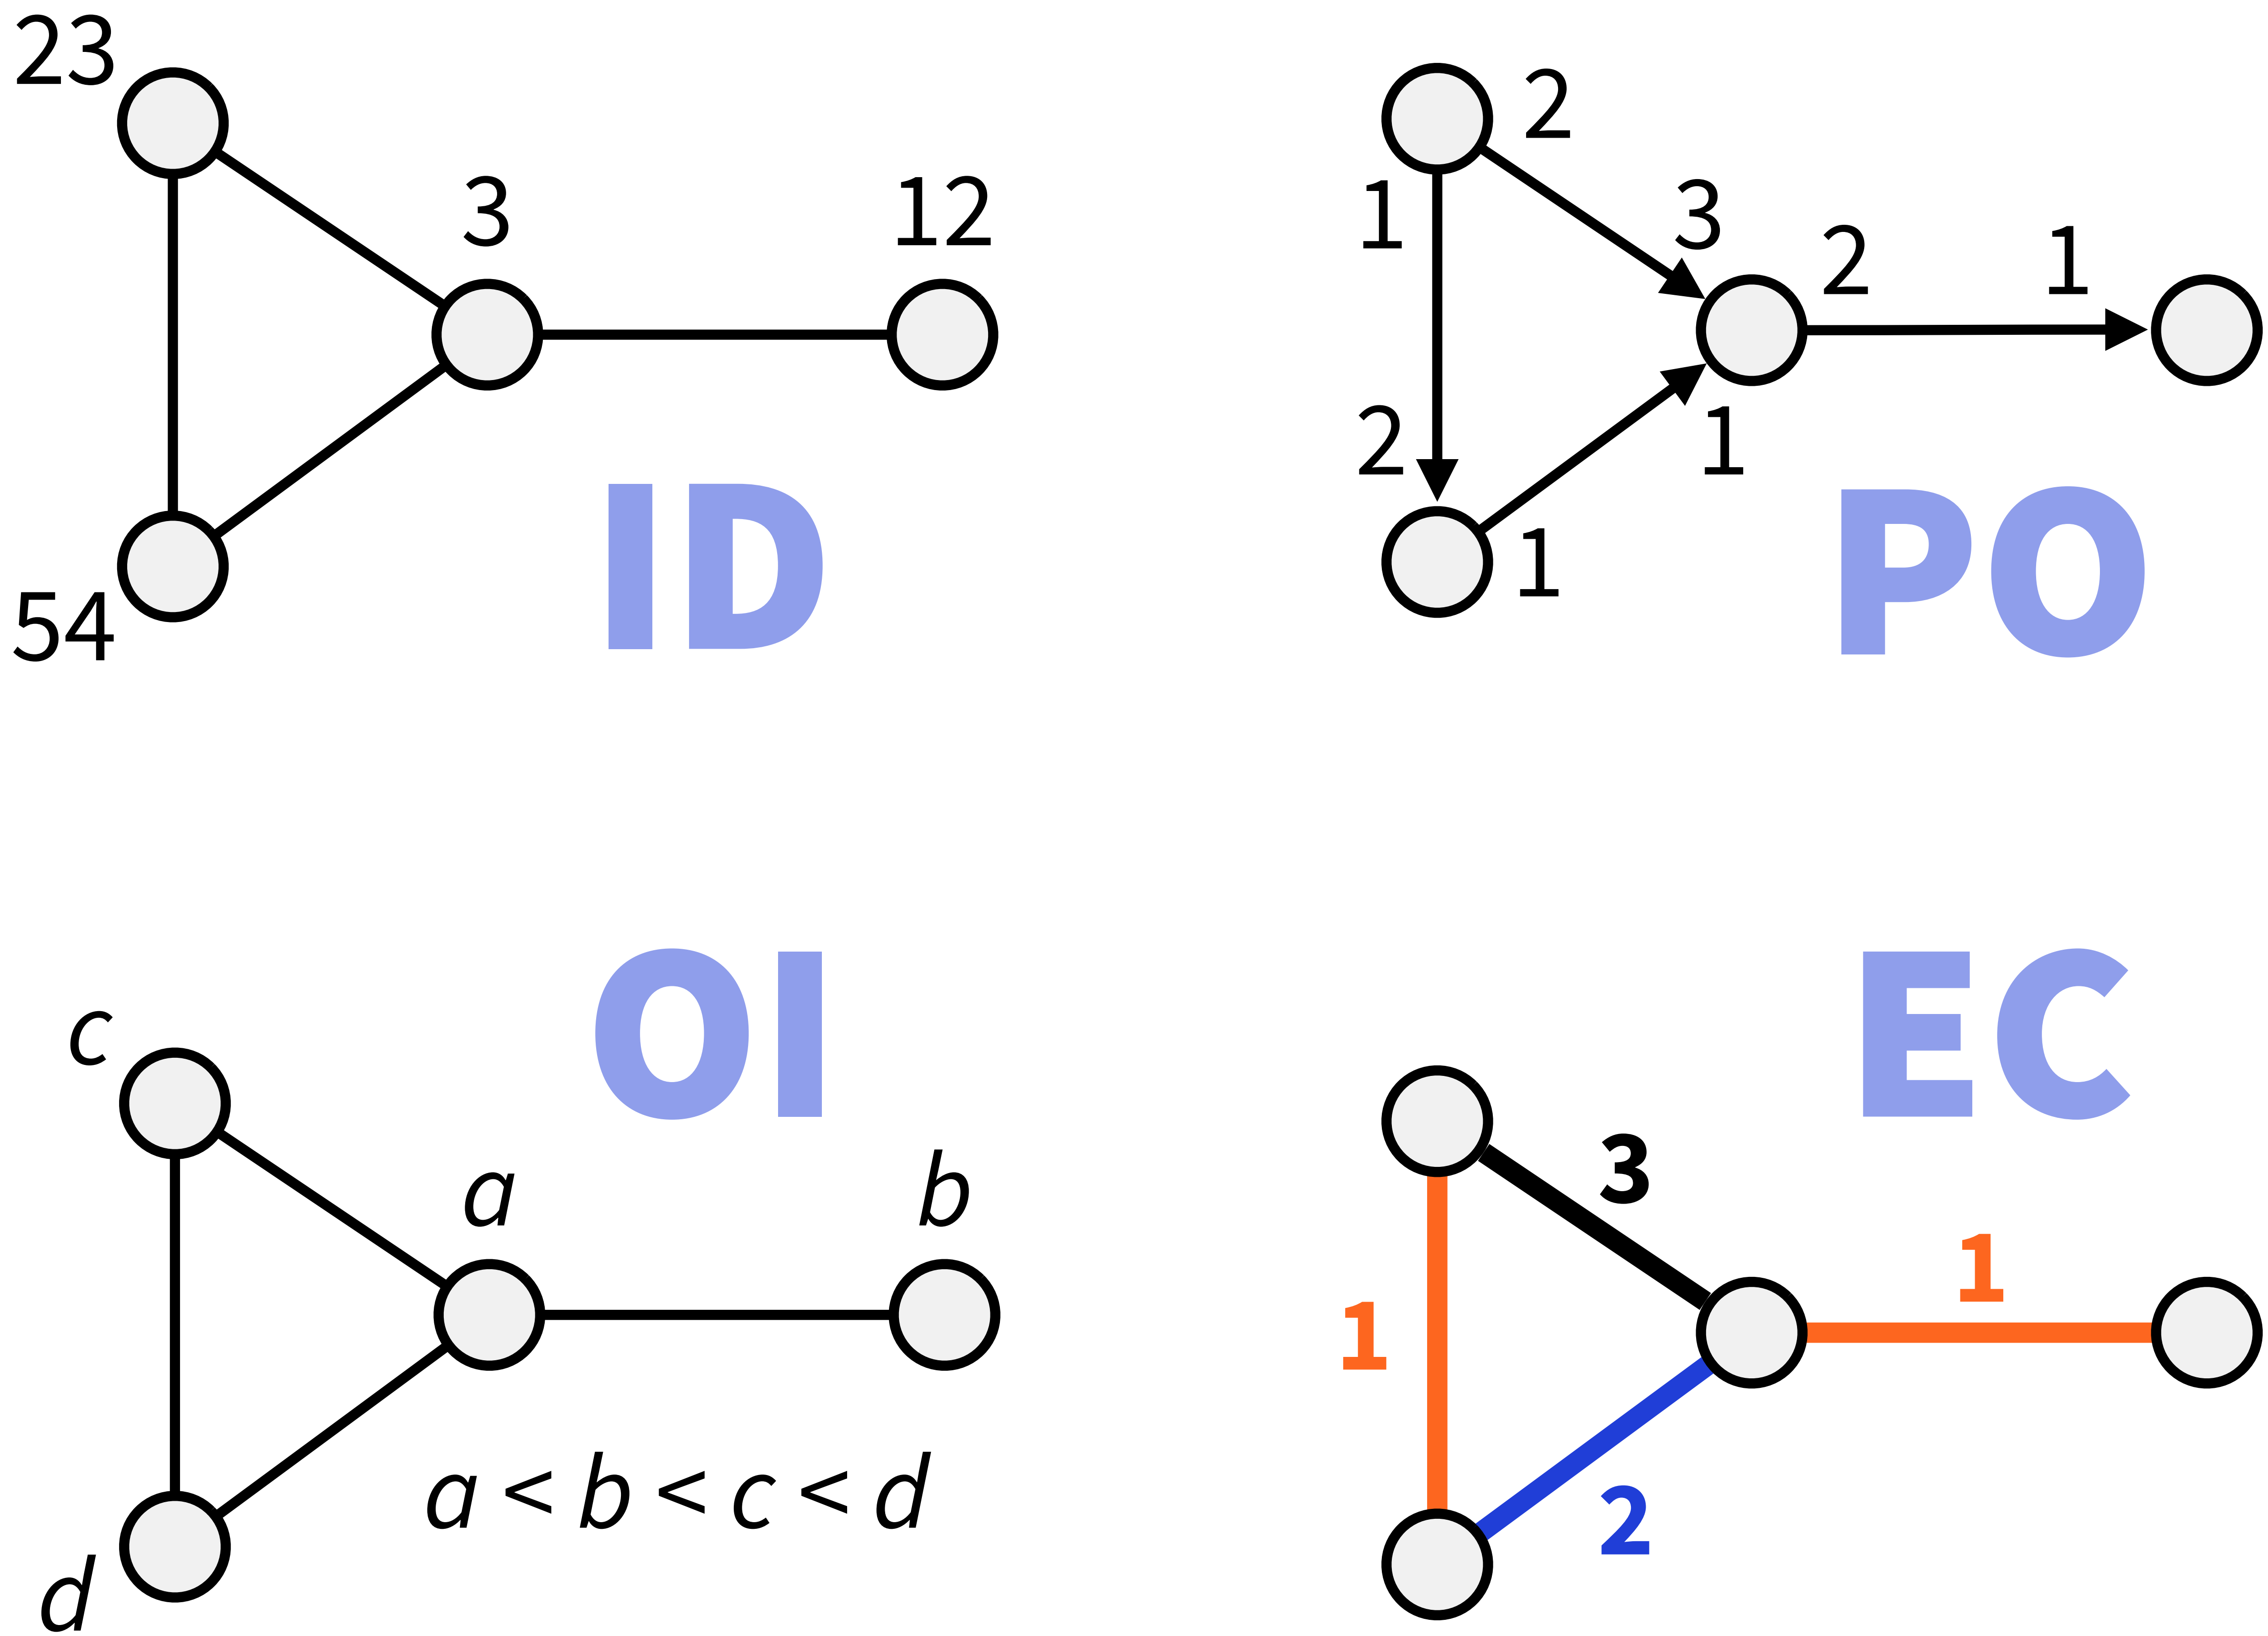

23
3
12
54
ID
2
1
3
2
1
2
1
1
PO
OI
c
a
b
d
a < b < c < d
EC
3
1
1
2

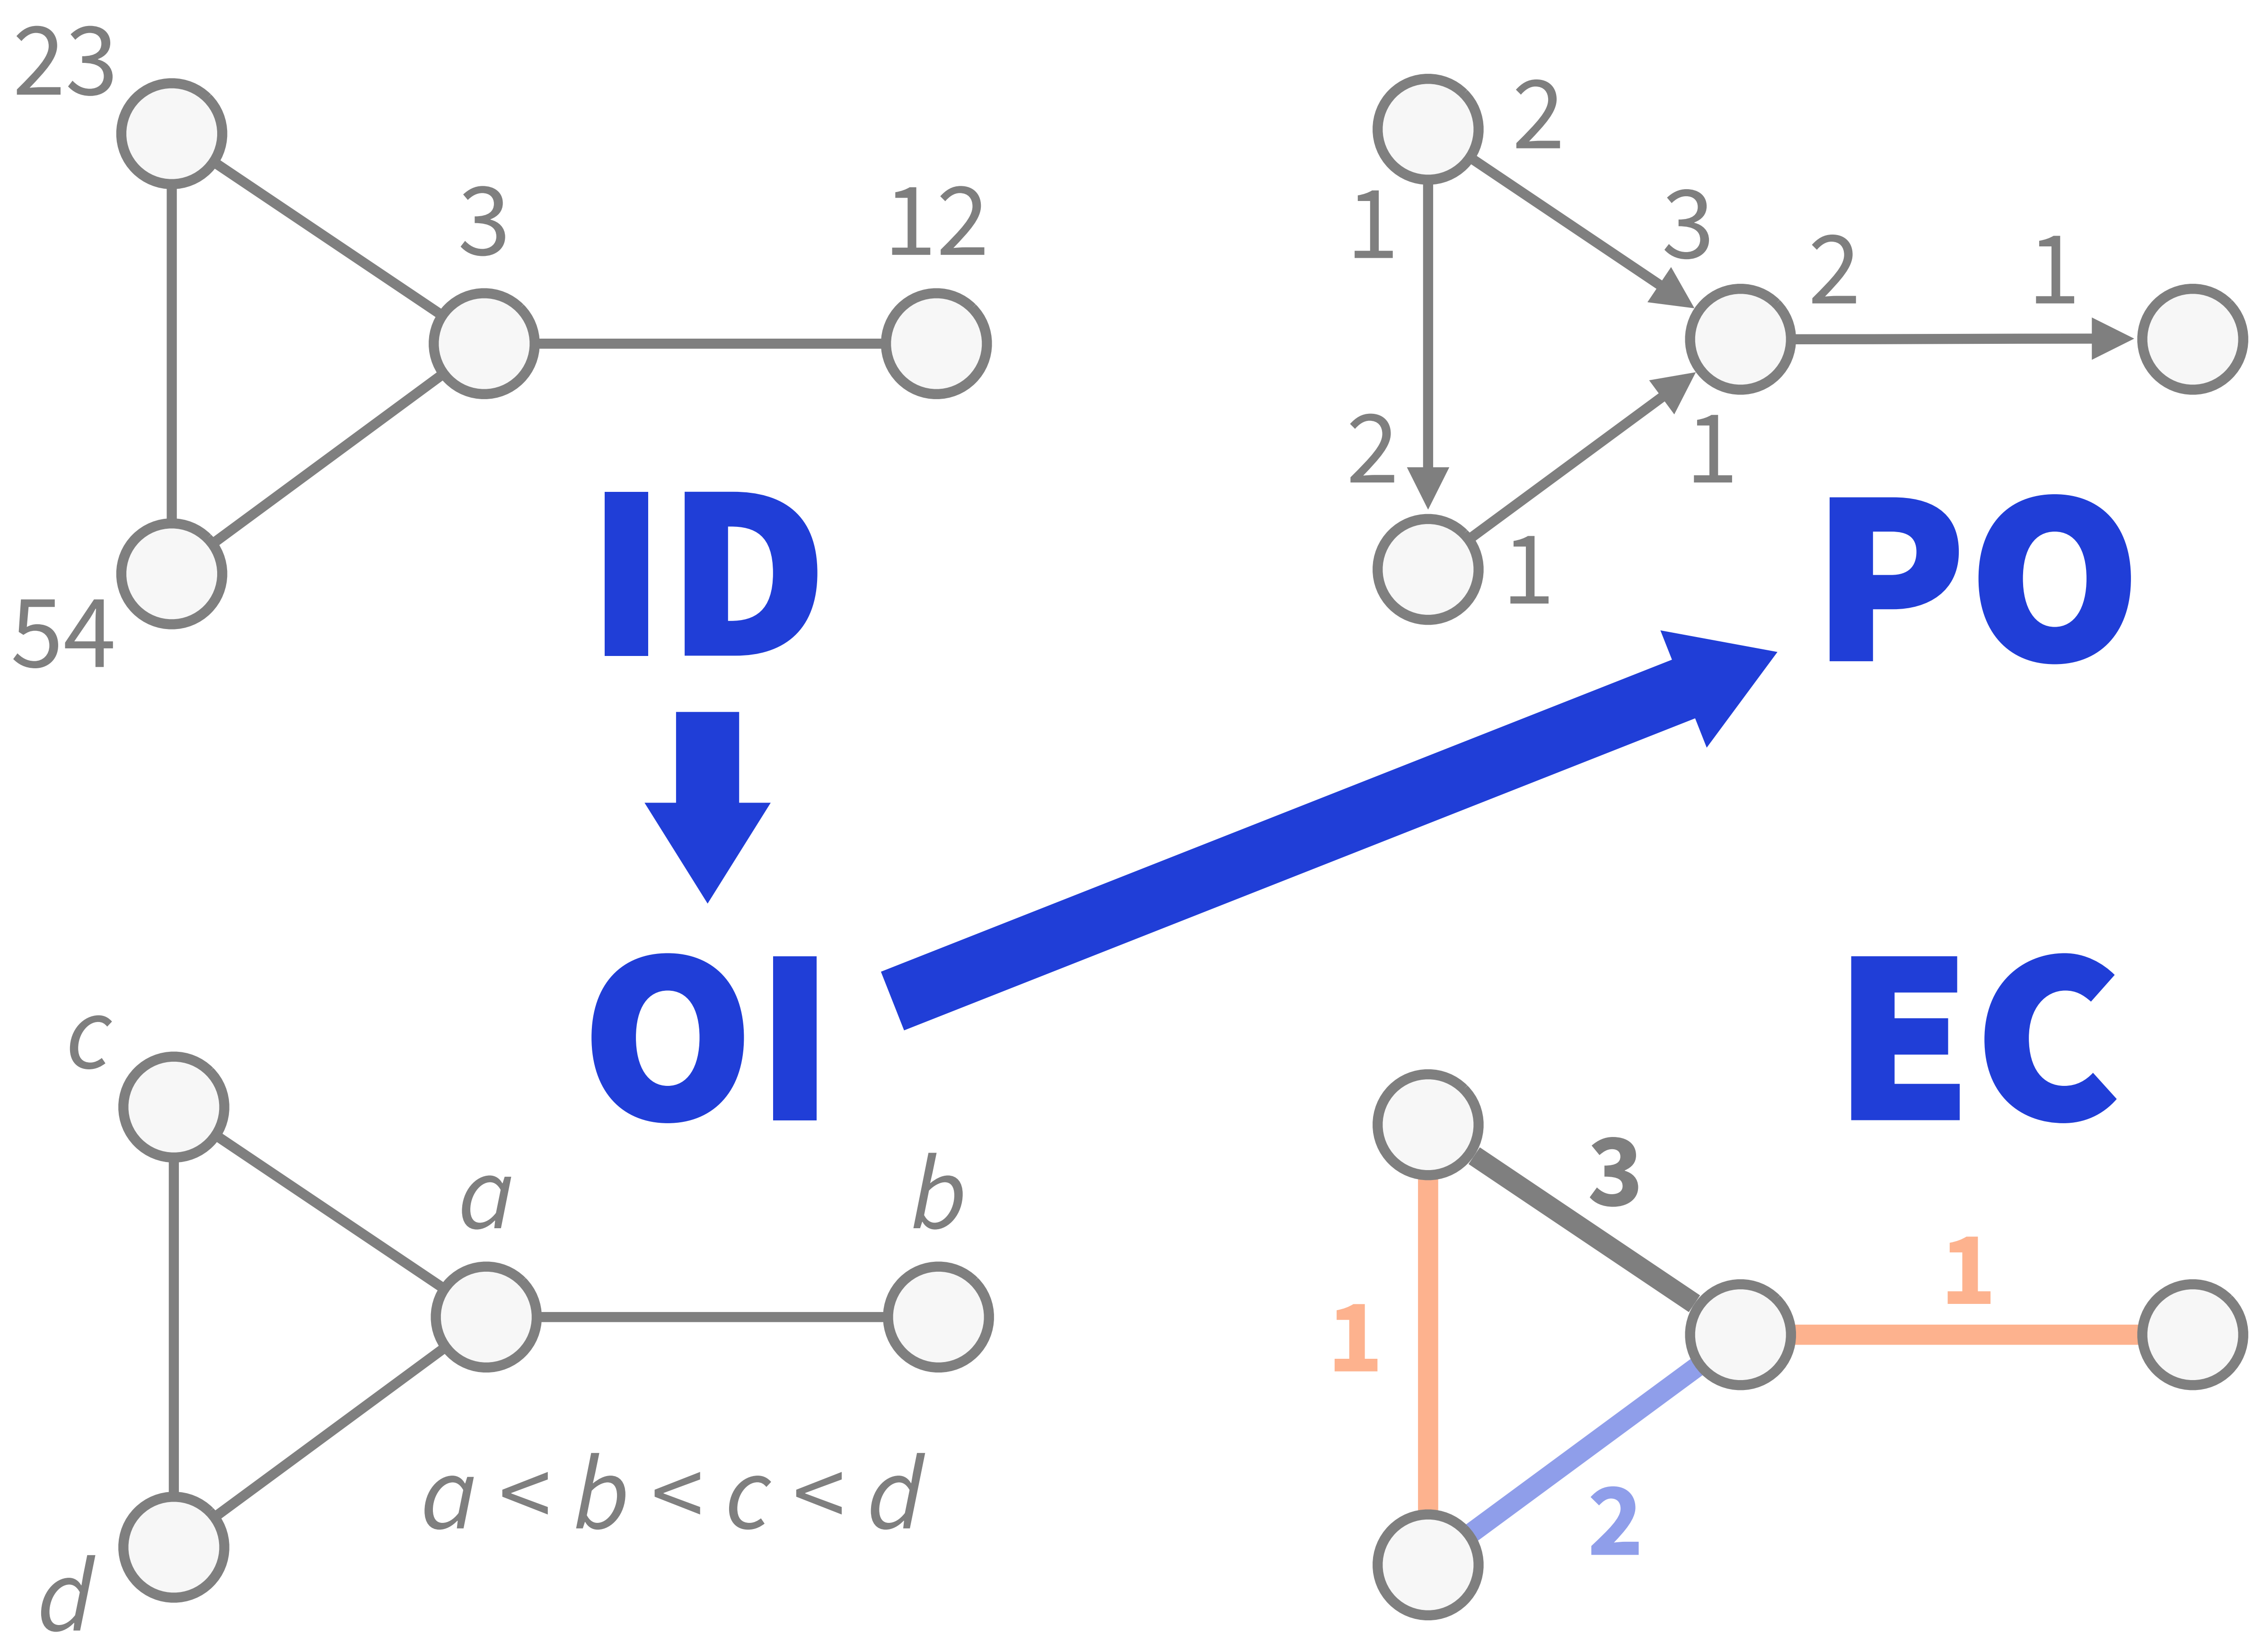

23
3
12
54
ID
OI
PO
EC
2
1
3
2
1
2
1
1
c
a
b
d
a < b < c < d
3
1
1
2

Simulation argument

- **Trivial: ID → OI → PO**
  - for any problem
- **We show: EC → PO → OI → ID**
  - for maximal fractional matching in “loopy graphs”

Proof overview

- **EC model is very limited**
  - maximal fractional matching requires $\Omega(\Delta)$ time in EC, even for **“loopy graphs”**
- **Simulation argument: EC → PO → OI → ID**
  - maximal fractional matching requires $\Omega(\Delta)$ time in ID, at least for **“loopy graphs”**

EC

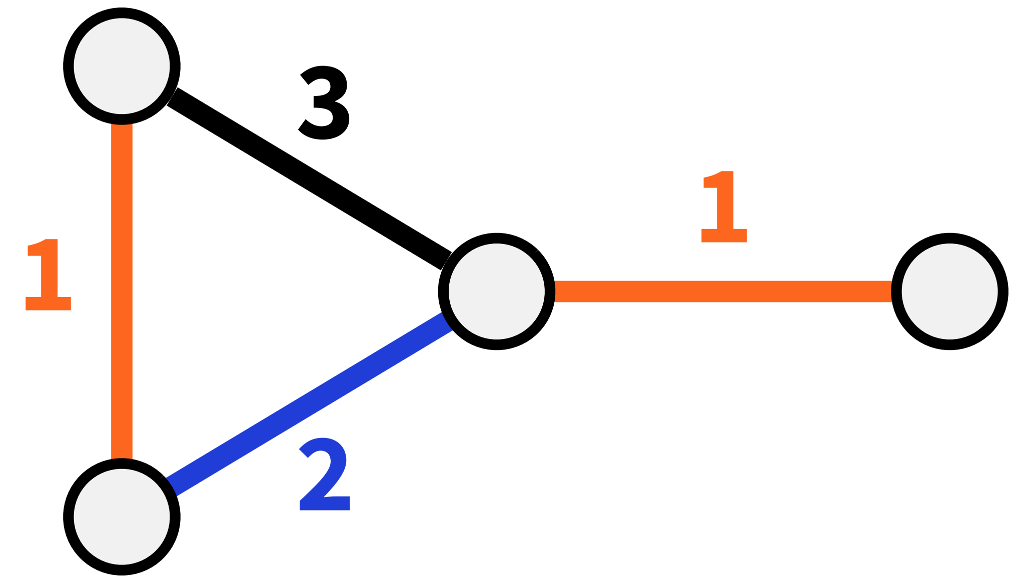

- **Recursively construct a degree-*i* graph where this algorithm takes time *i***
- **Focus on “loopy graphs”**
  - highly symmetric
  - forces algorithm to produce “tight” outputs (all nodes saturated, “perfect matching”)

EC → PO

"Unhelpful" port numbering & orientation

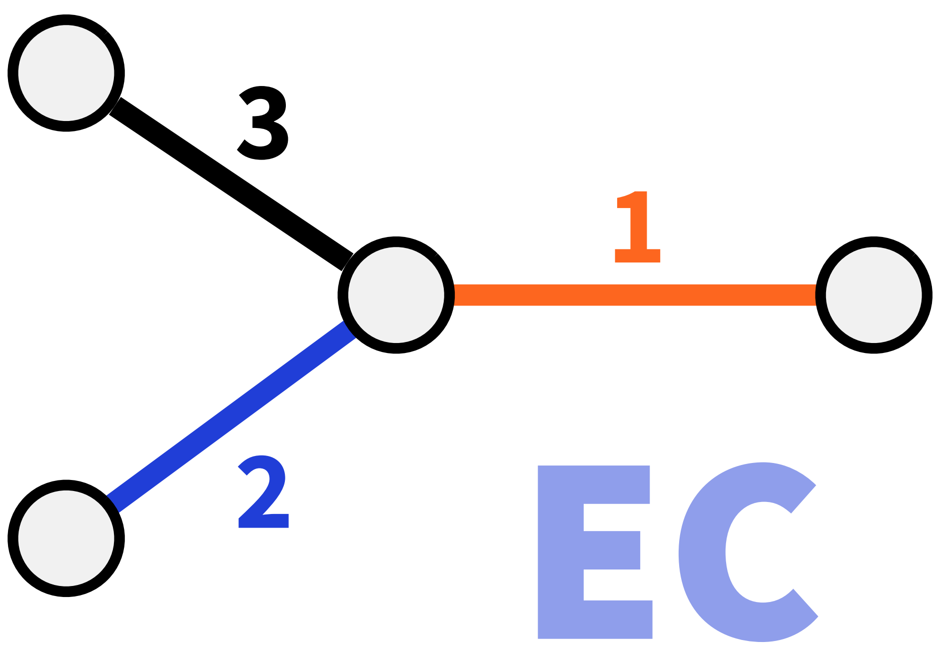

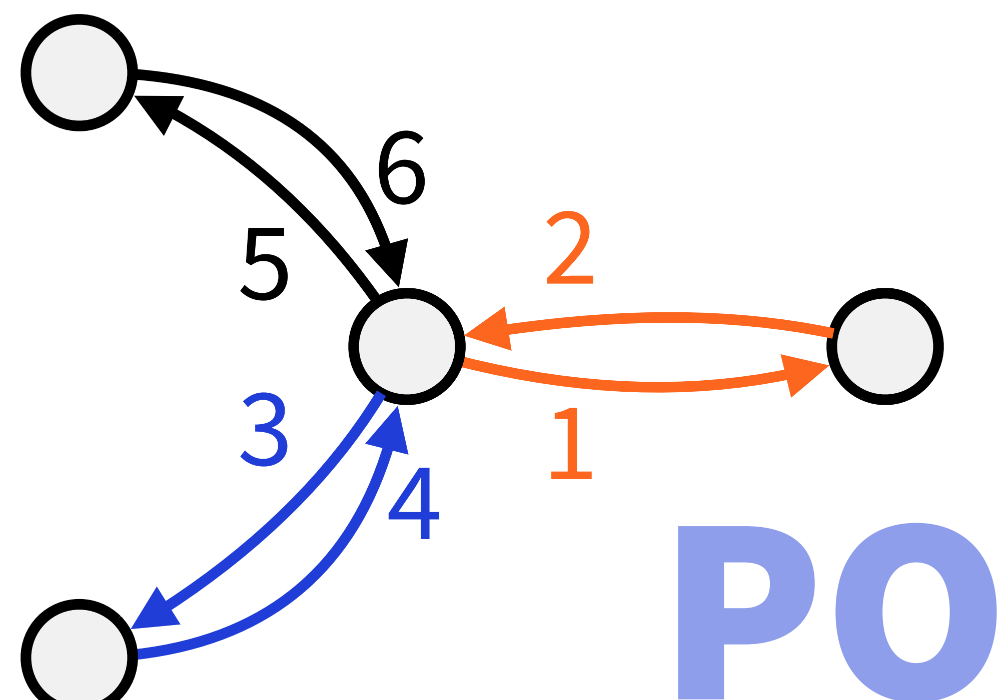

PO → OI

“Unhelpful” total order

can be easily constructed given a port numbering and orientation

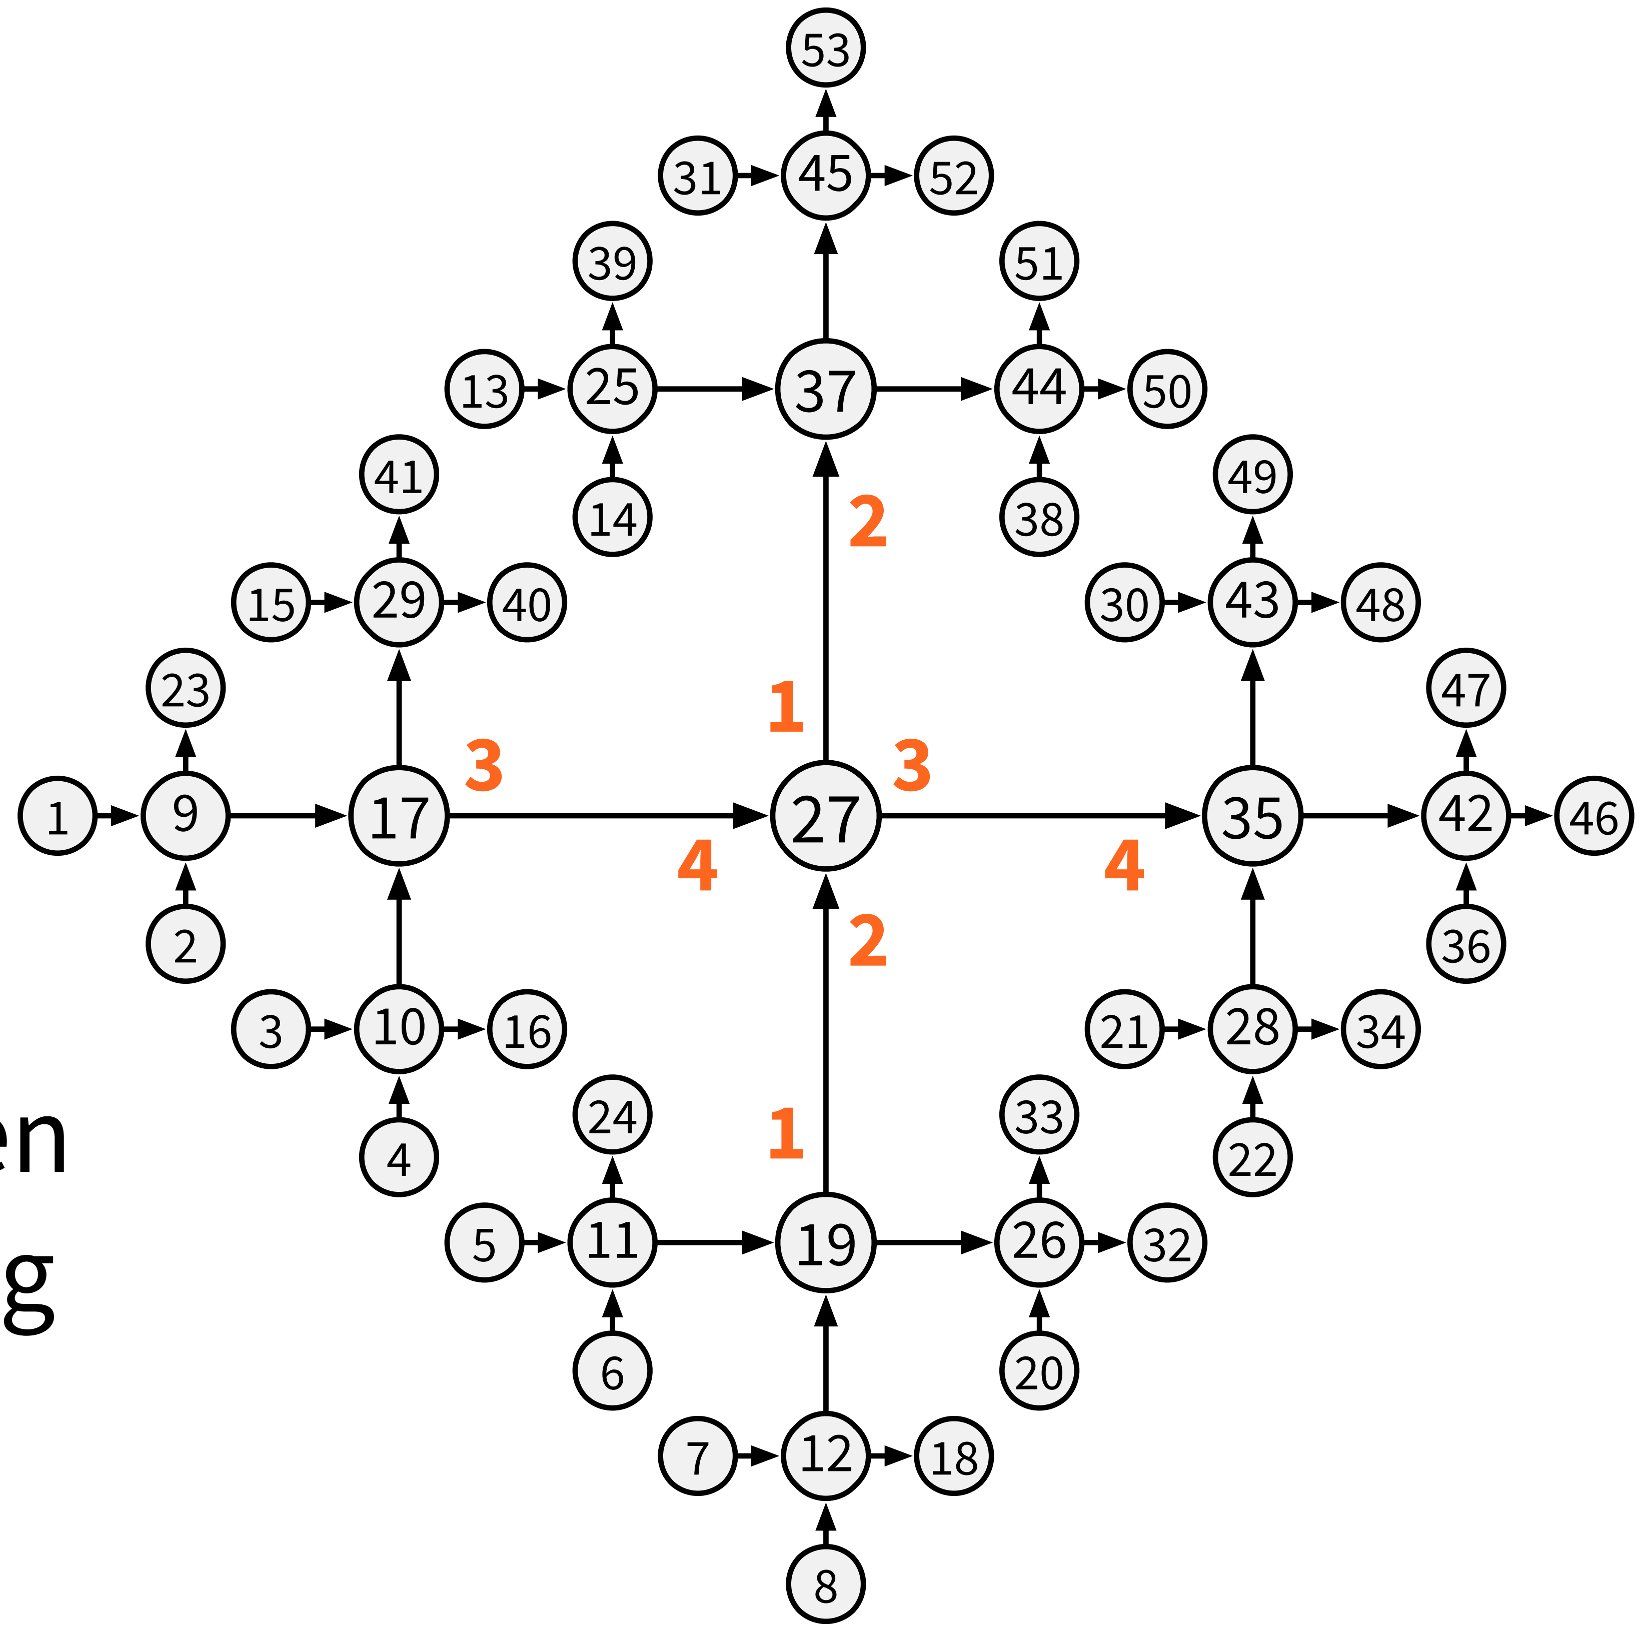

OI → ID

“Unhelpful” unique identifiers

Ramsey-like argument:
for any algorithm we can find unique identifiers that do not help in comparison with total order

EC → PO → OI → ID

- **In general:** stronger models help
- **In our case:** we can always come up with situations in which ID model is not any better than EC model

What about other problems?

- **Now we have a linear-in-Δ lower bound for maximal fractional matching**
- **Can we use the same techniques to prove lower bounds for other problems?**
  - e.g., maximal matching?

General recipe

1. **Find a suitable “simple model”**
2. **Prove a lower bound for the simple model**
   - keep input “symmetric”
   - keep output “tight” and “fragile”
   - local changes have non-local consequences

General recipe

1. **Find a suitable "simple model"**
2. **Prove a lower bound for the simple model**
3. **Amplify the lower bound**
   - simple model → OI (some thinking required)
   - OI → ID (standard techniques)

What about maximal matchings?

- **Could we use the same techniques to show that $o(\Delta) + O(\log^* n)$ is not sufficient for maximal matching?**
- **Two obstacles…**

What about maximal matchings?

- **Obstacle 1 — final step:**
  - final step OI → ID based on a Ramsey argument
  - works great for t independent of n
  - fails if $t \approx \log^* n$

What about maximal matchings?

- **Obstacle 2 — starting point:**
  - $O(\log^* n)$ time enough to find e.g. graph colouring
  - cannot assume “symmetric” input
  - difficult to force “tight” and “fragile” output

What about maximal matchings?

- **Two hard, interlinked obstacles**
- **How to proceed:**
  - get rid of obstacle 1 — $\log^* n$
  - focus on obstacle 2 — asymmetry
- **Start with bipartite maximal matchings**

Maximal matching in 2-coloured graphs

- **Can be solved in time $O(\Delta)$ independently of n**
- **Can focus on just one obstacle: asymmetry**
- **Most of the other machinery already exists!**
  - we just need tight bounds for simple models
  - should be easy to generalise to LOCAL model

Maximal matching in 2-coloured graphs

- **Until we have lower bounds: reductions, conditional lower bounds**
  - many other problems are at least as hard as bipartite maximal matching
  - locally optimal semi-matching in time T → bipartite maximal matching in time T

Summary

- **Distributed time complexity, LOCAL model**
- **$O(\log^* n)$: “symmetry breaking”, OK**
- **$O(\Delta)$: “local coordination”, poorly understood**
- **Maximal *fractional* matching solved, next step: *bipartite* maximal matching**