Local coordination and symmetry breaking

Jukka Suomela
Aalto University, Finland

Zürich, 26 August 2015

Running example:
Maximal matching

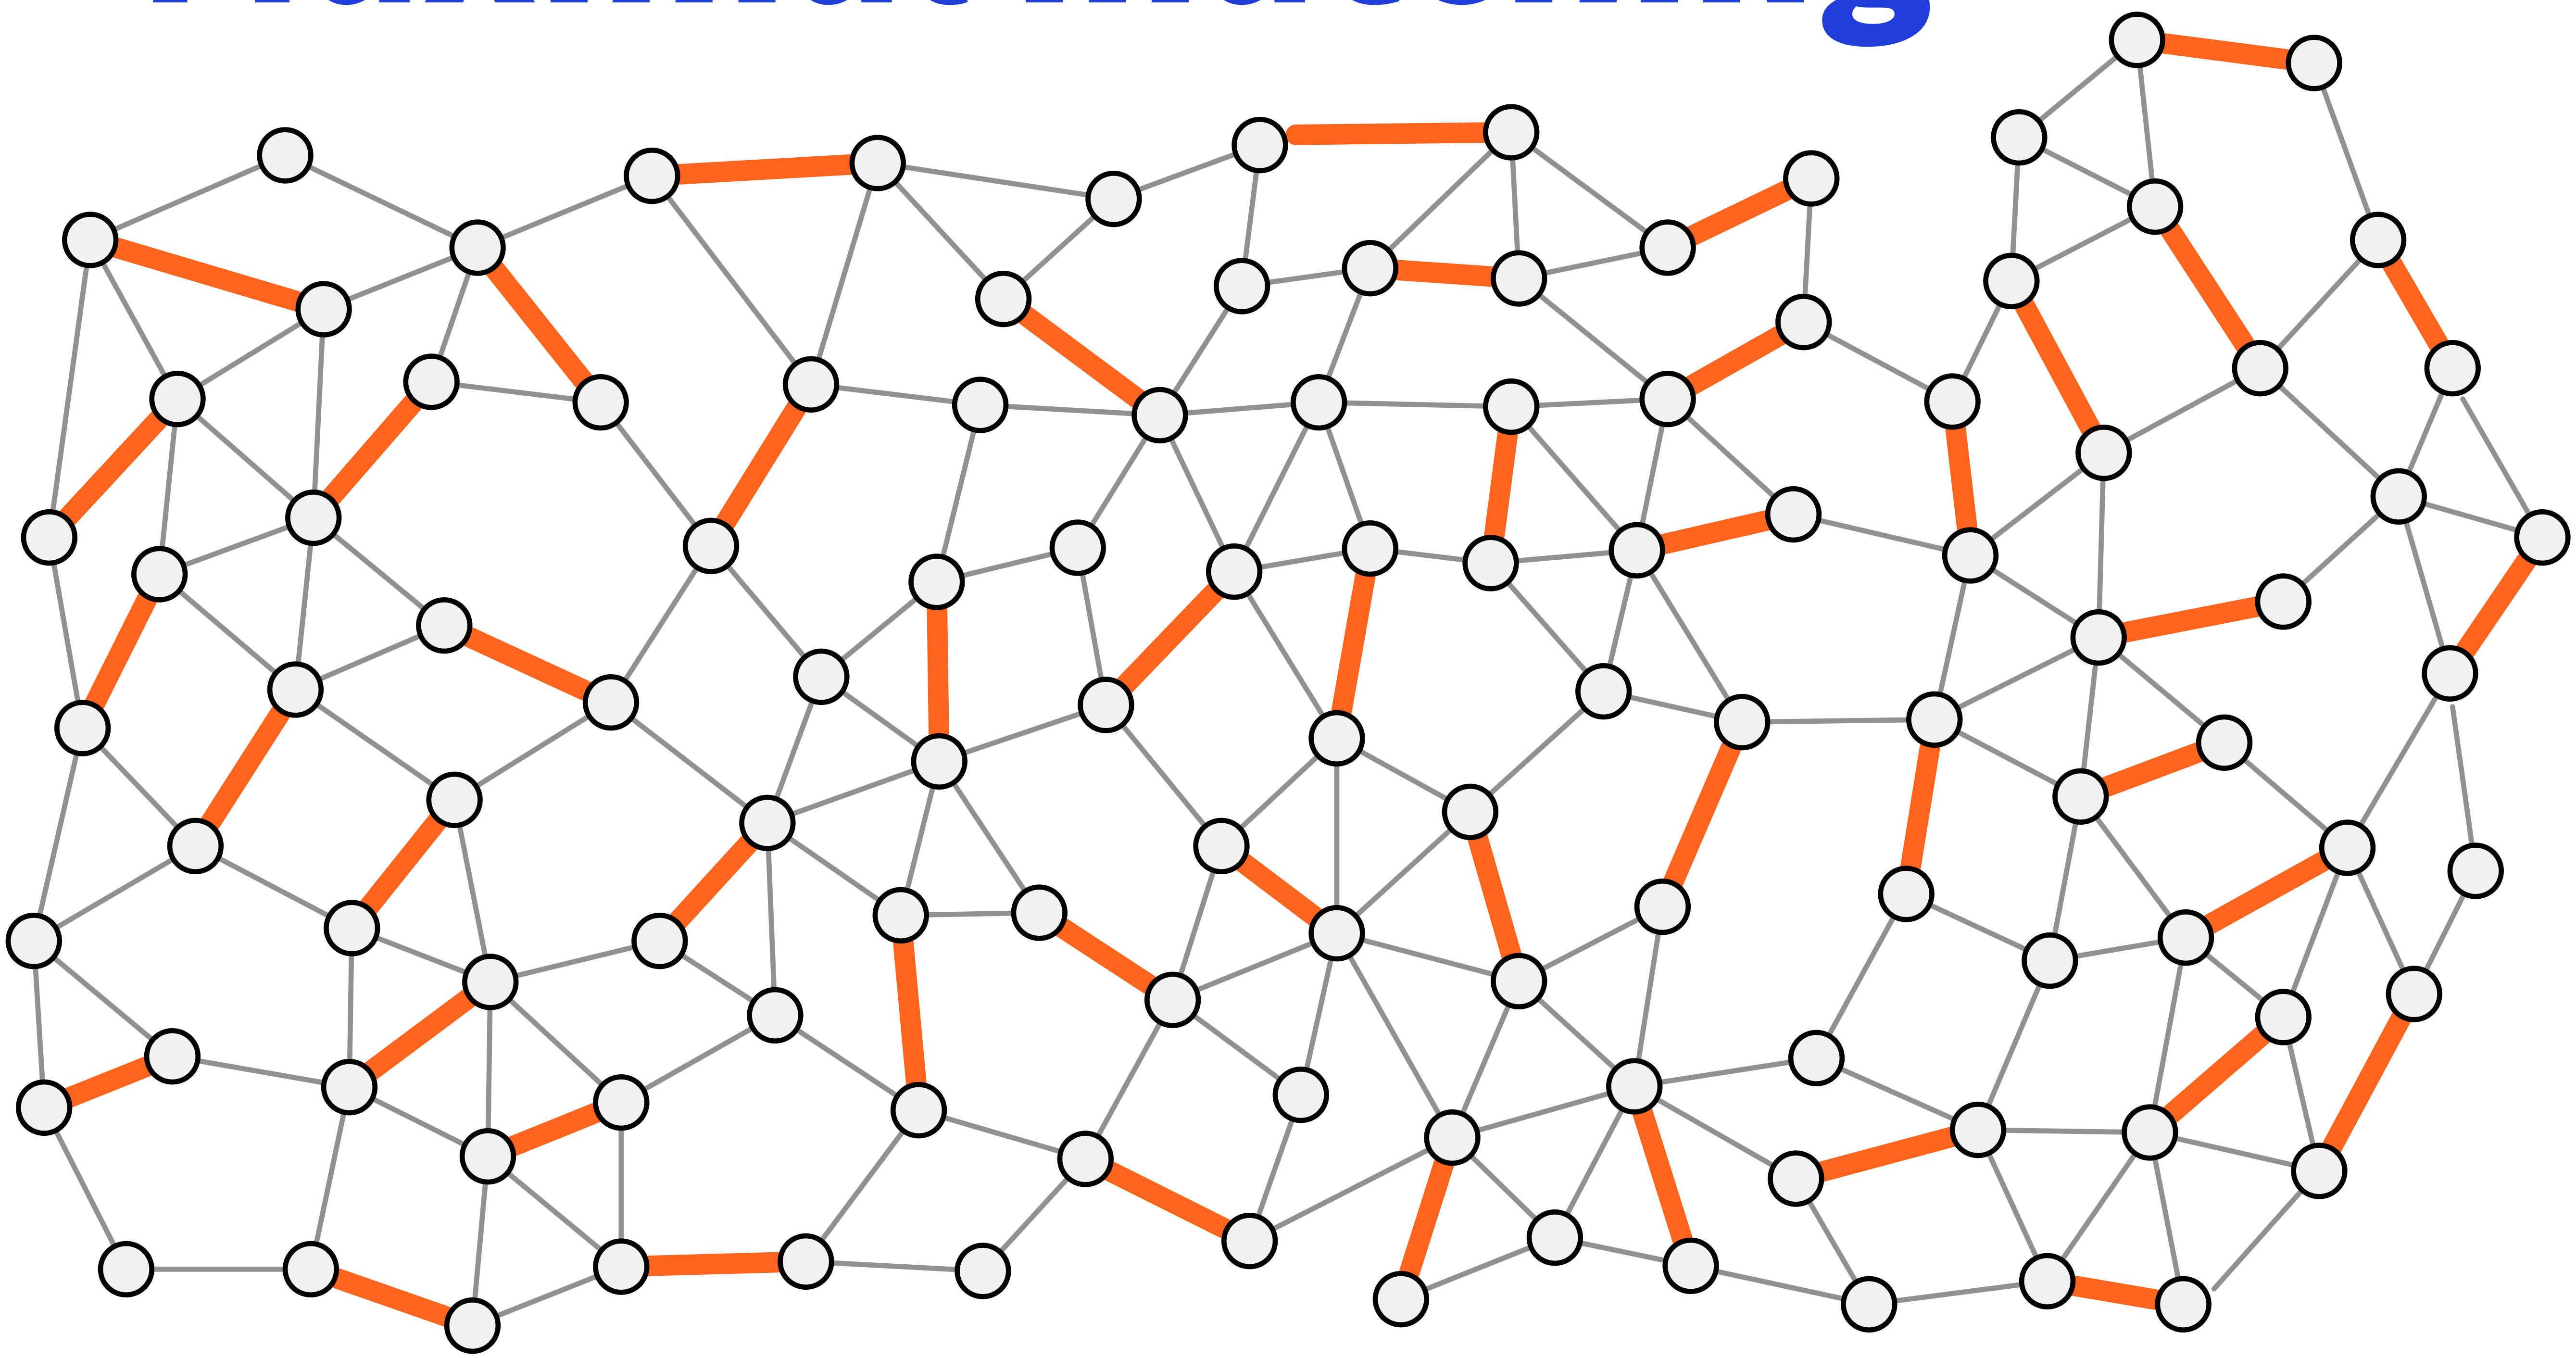

LOCAL model

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

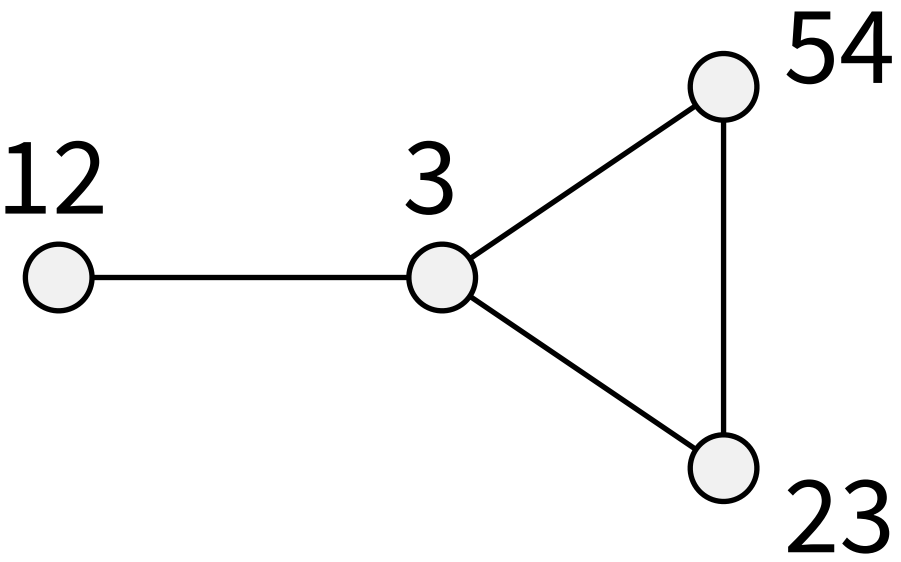

LOCAL model

- **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

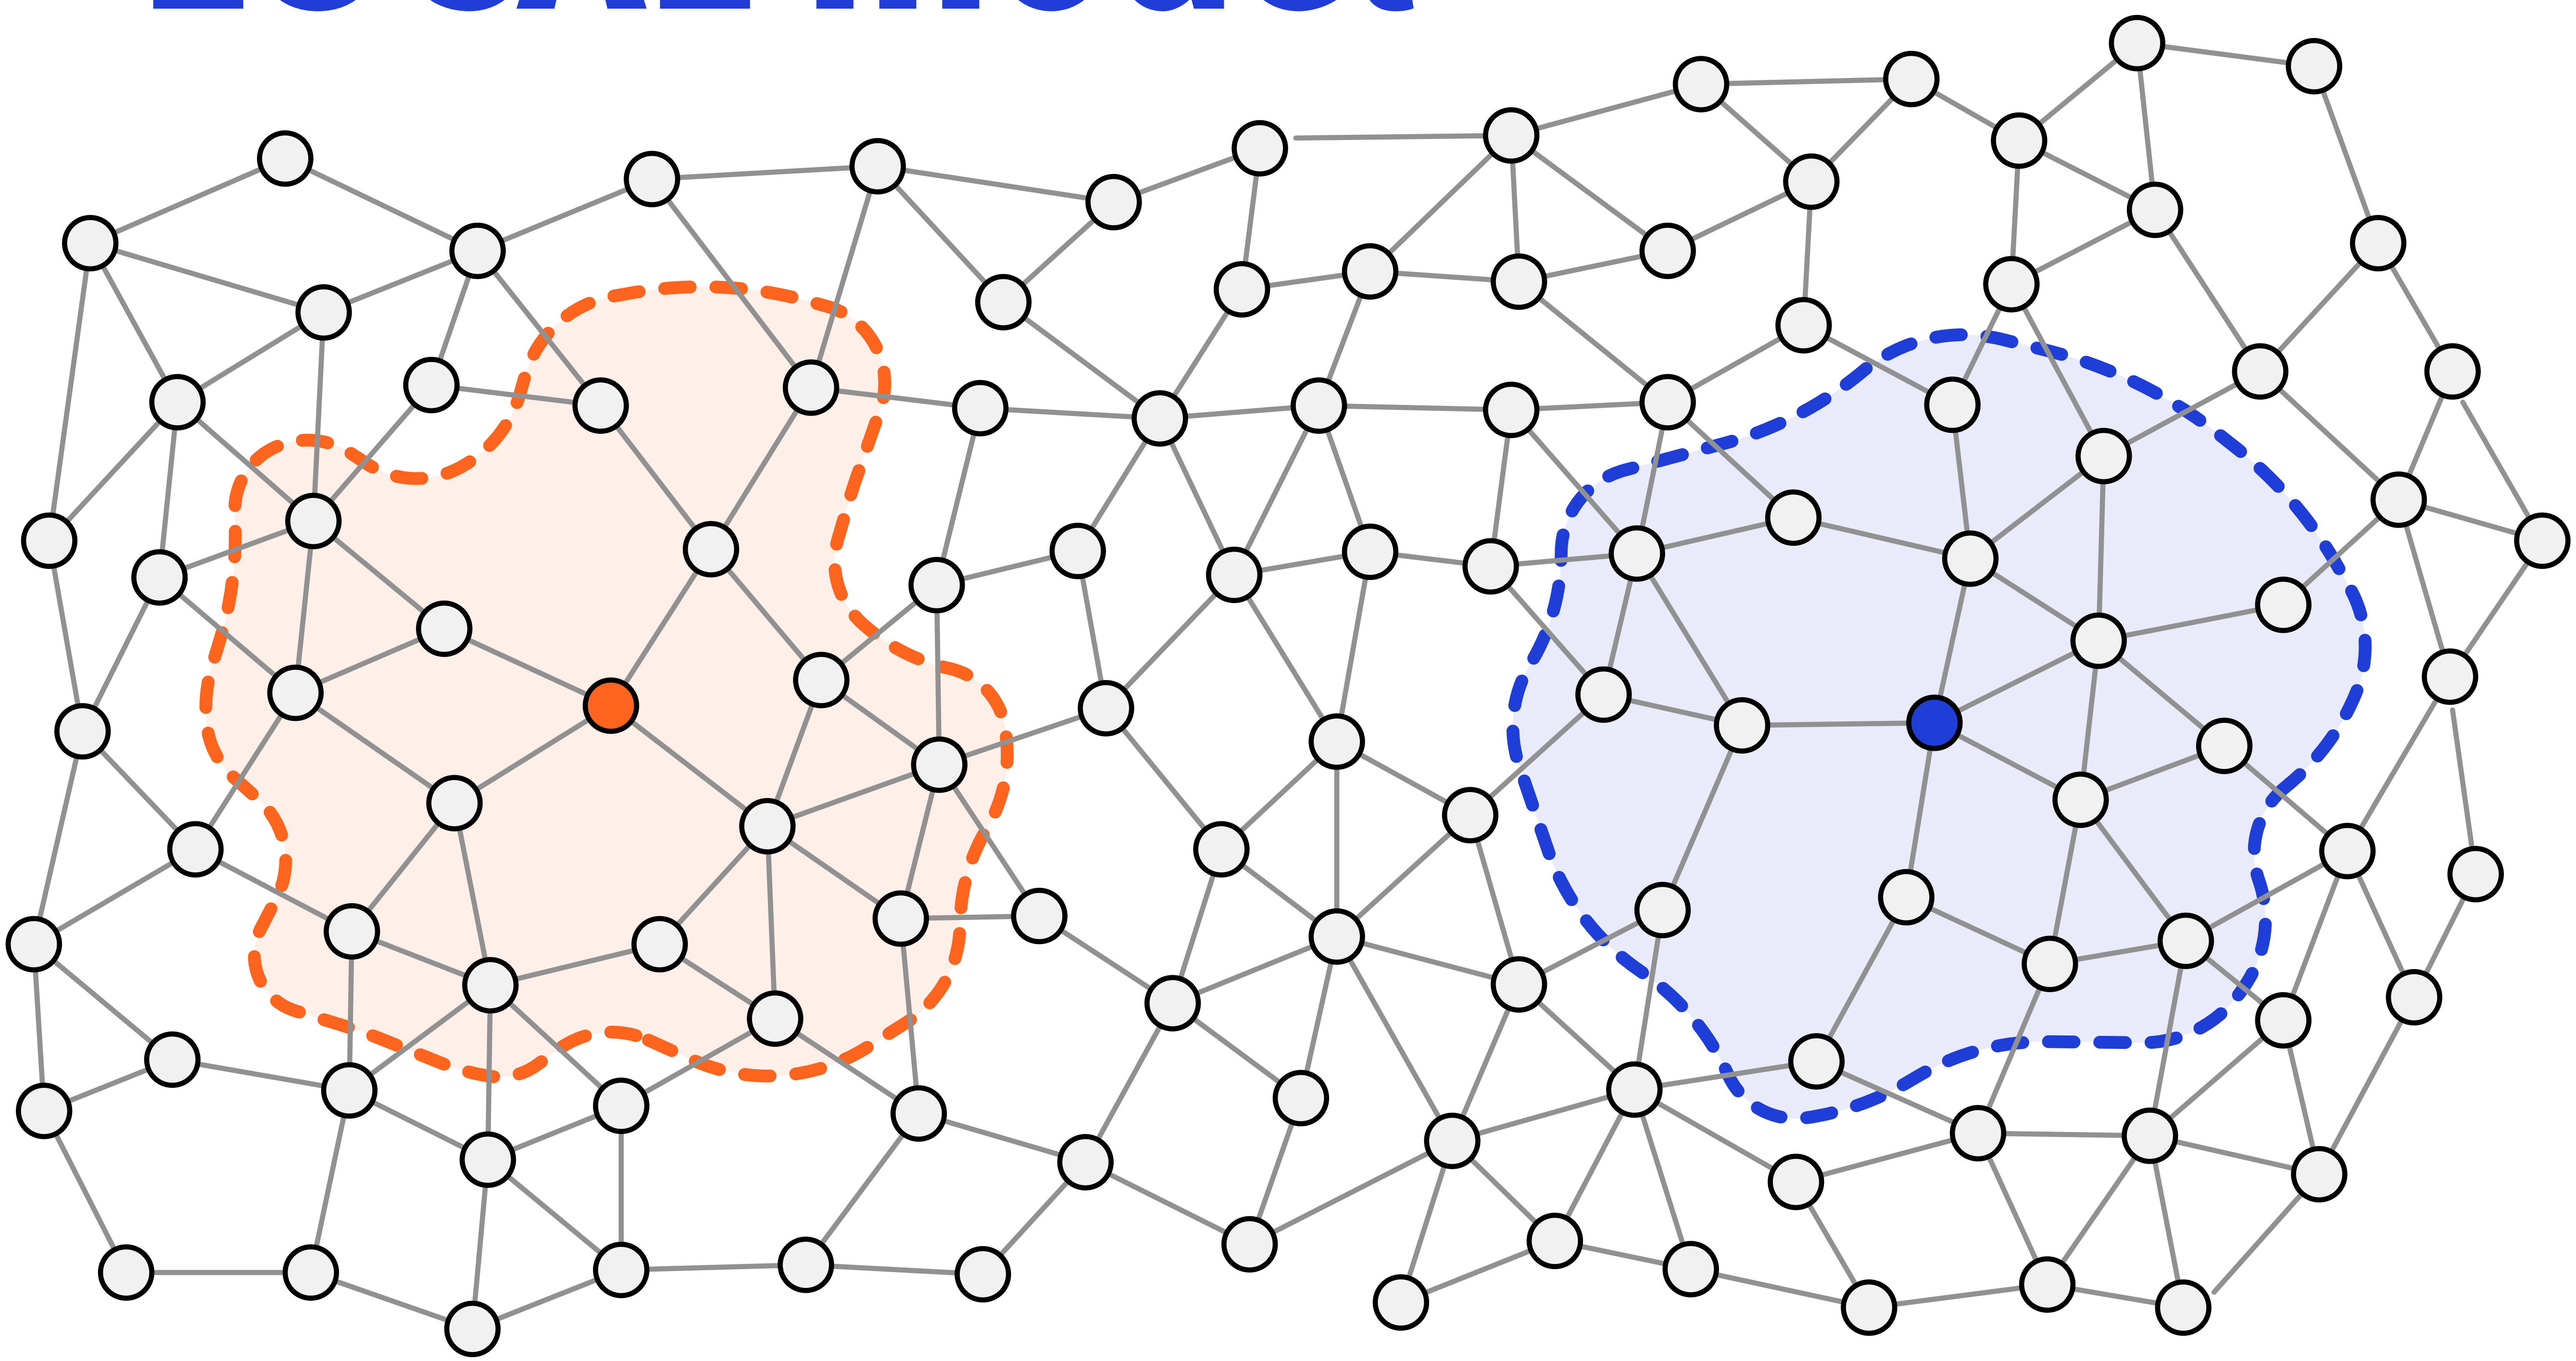

LOCAL model

A:

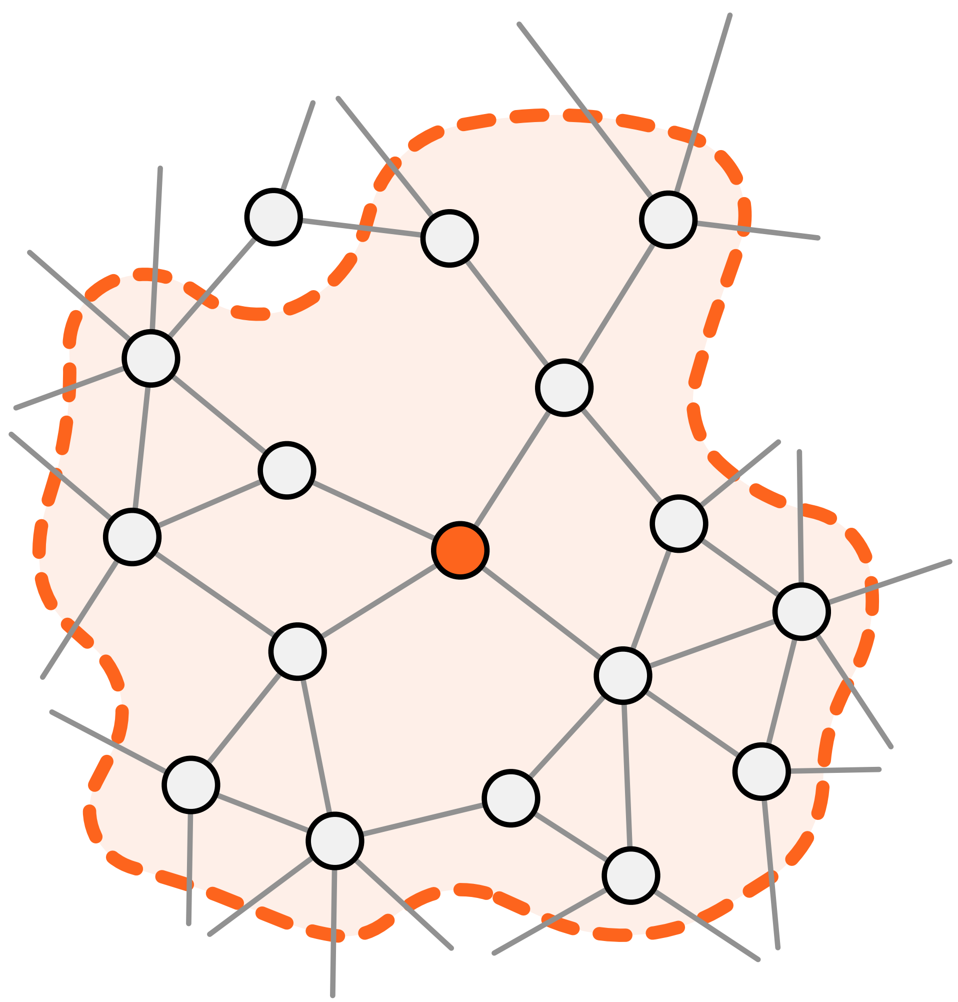

$\mapsto 1$

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

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

Landscape

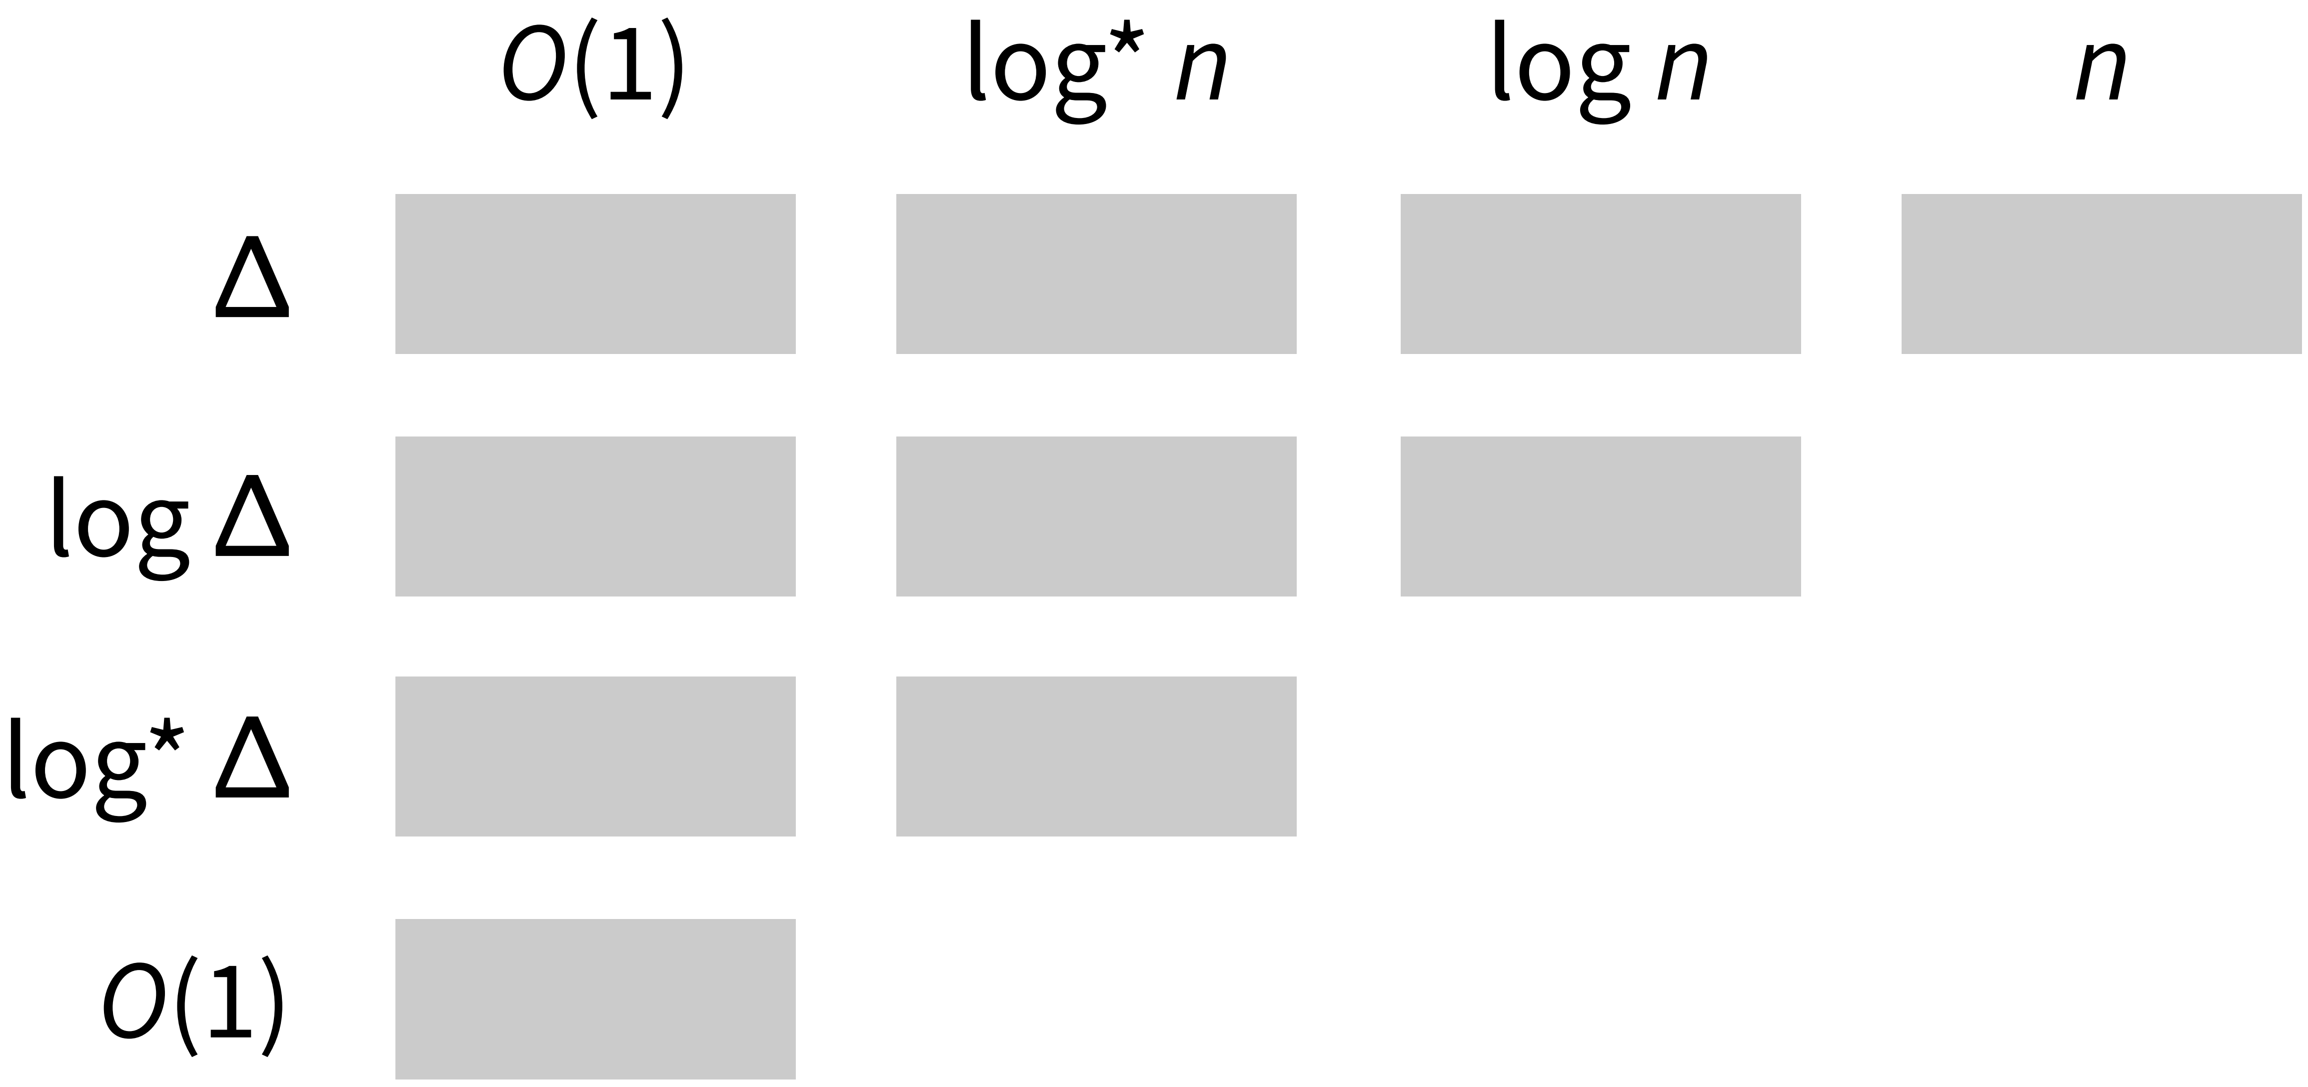

Landscape

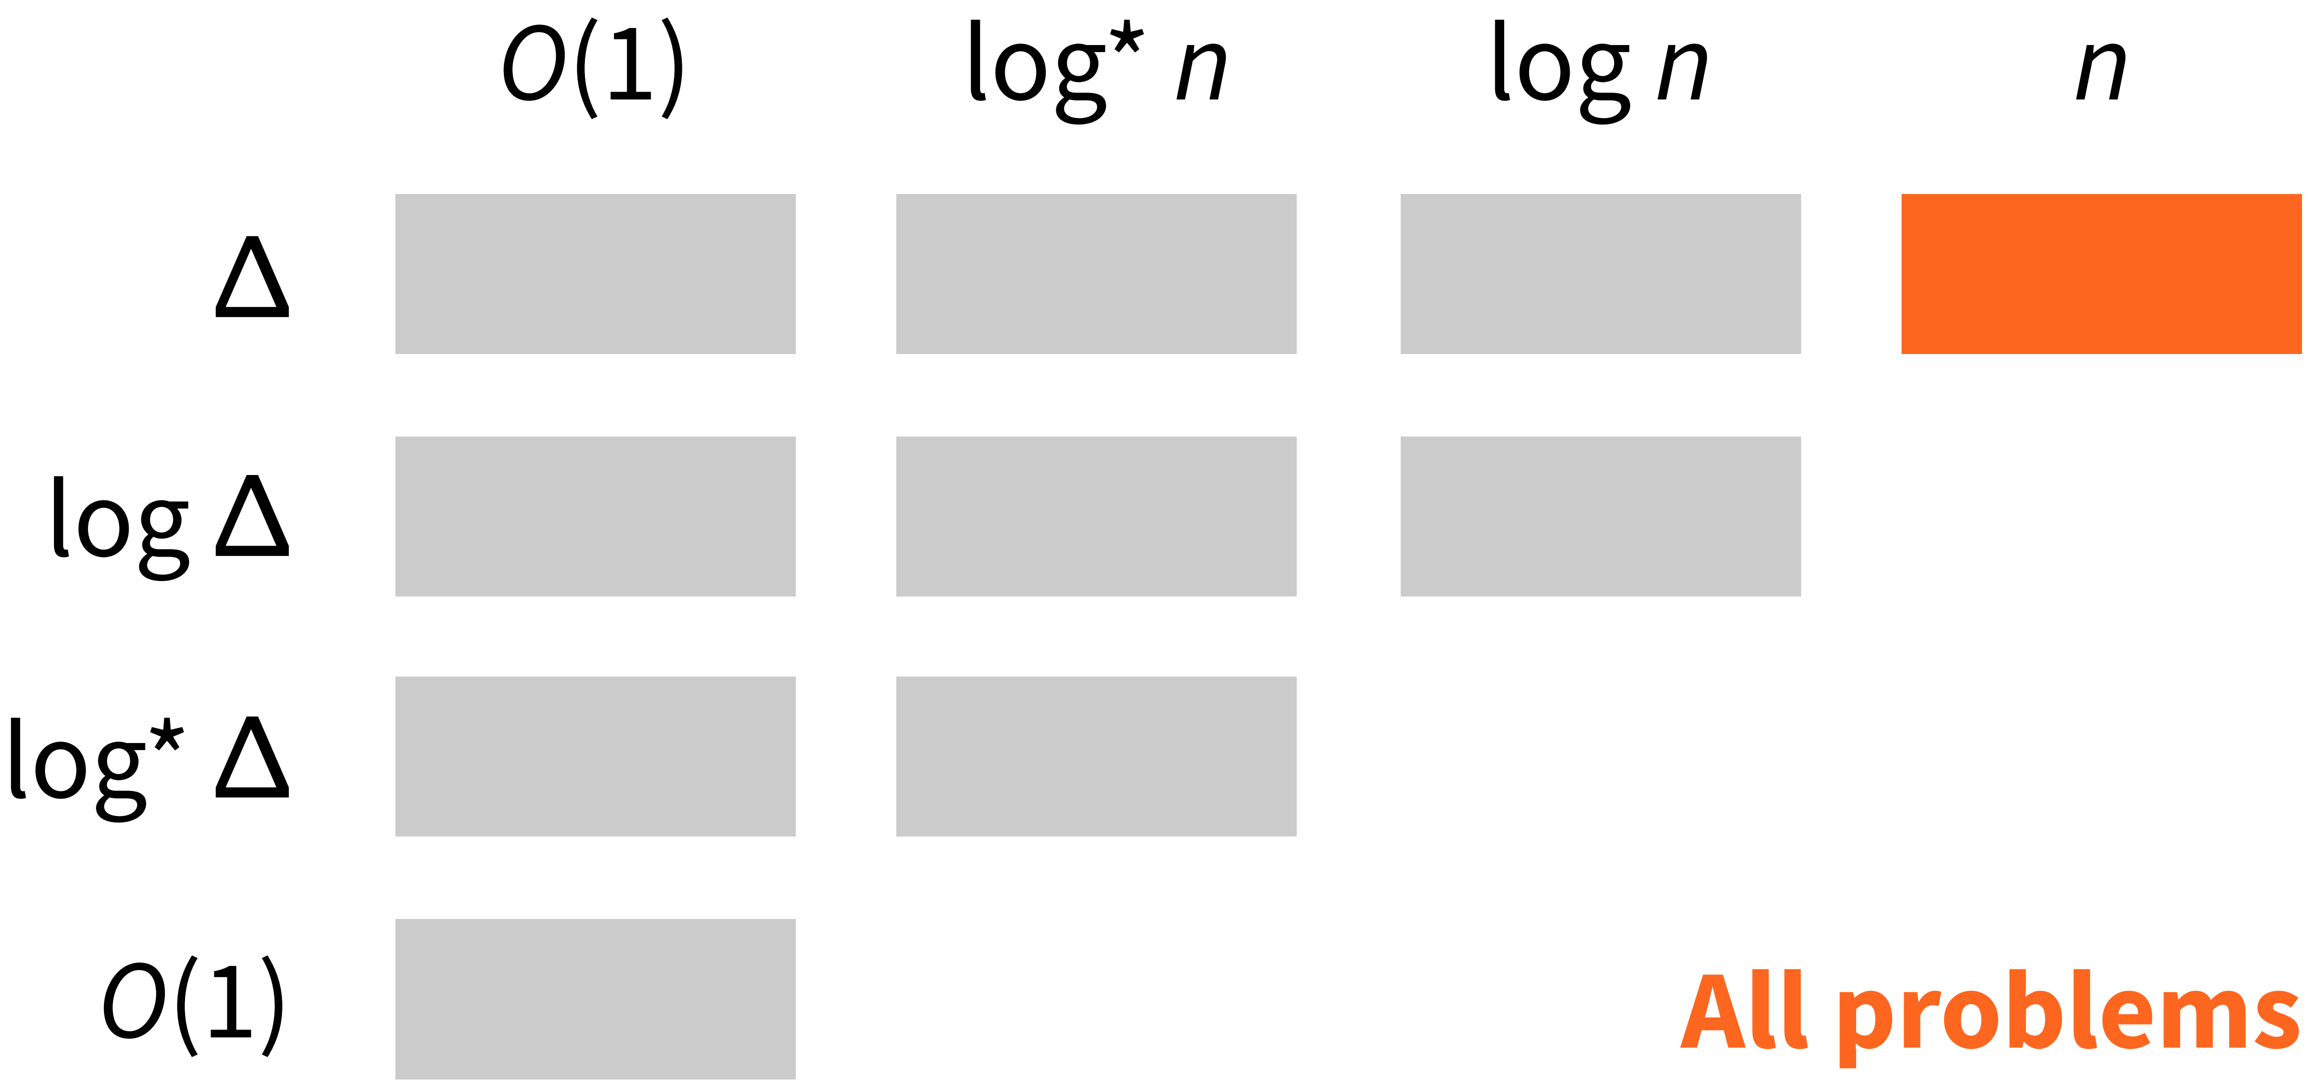

Landscape

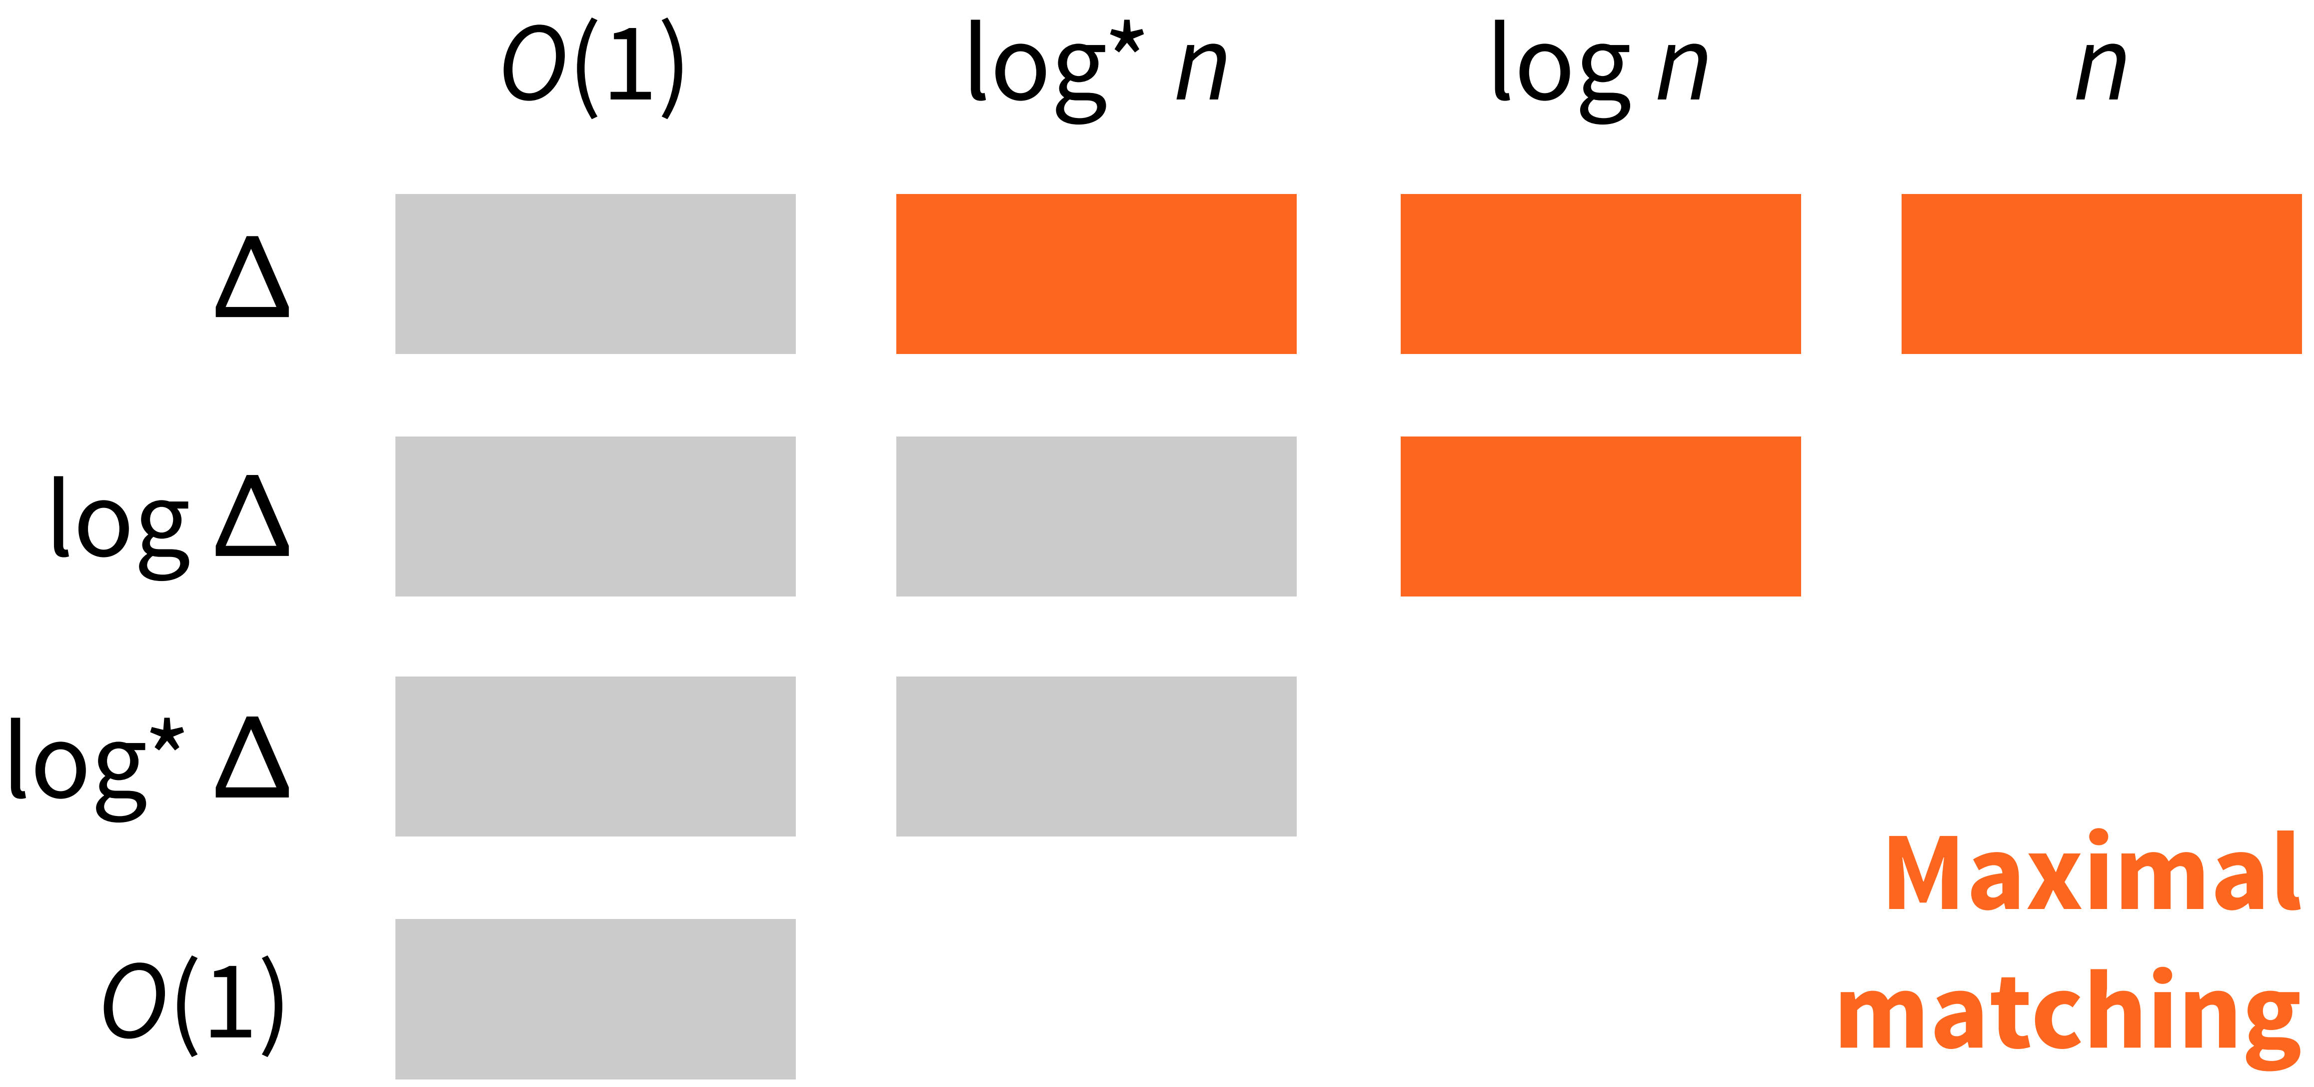

Landscape

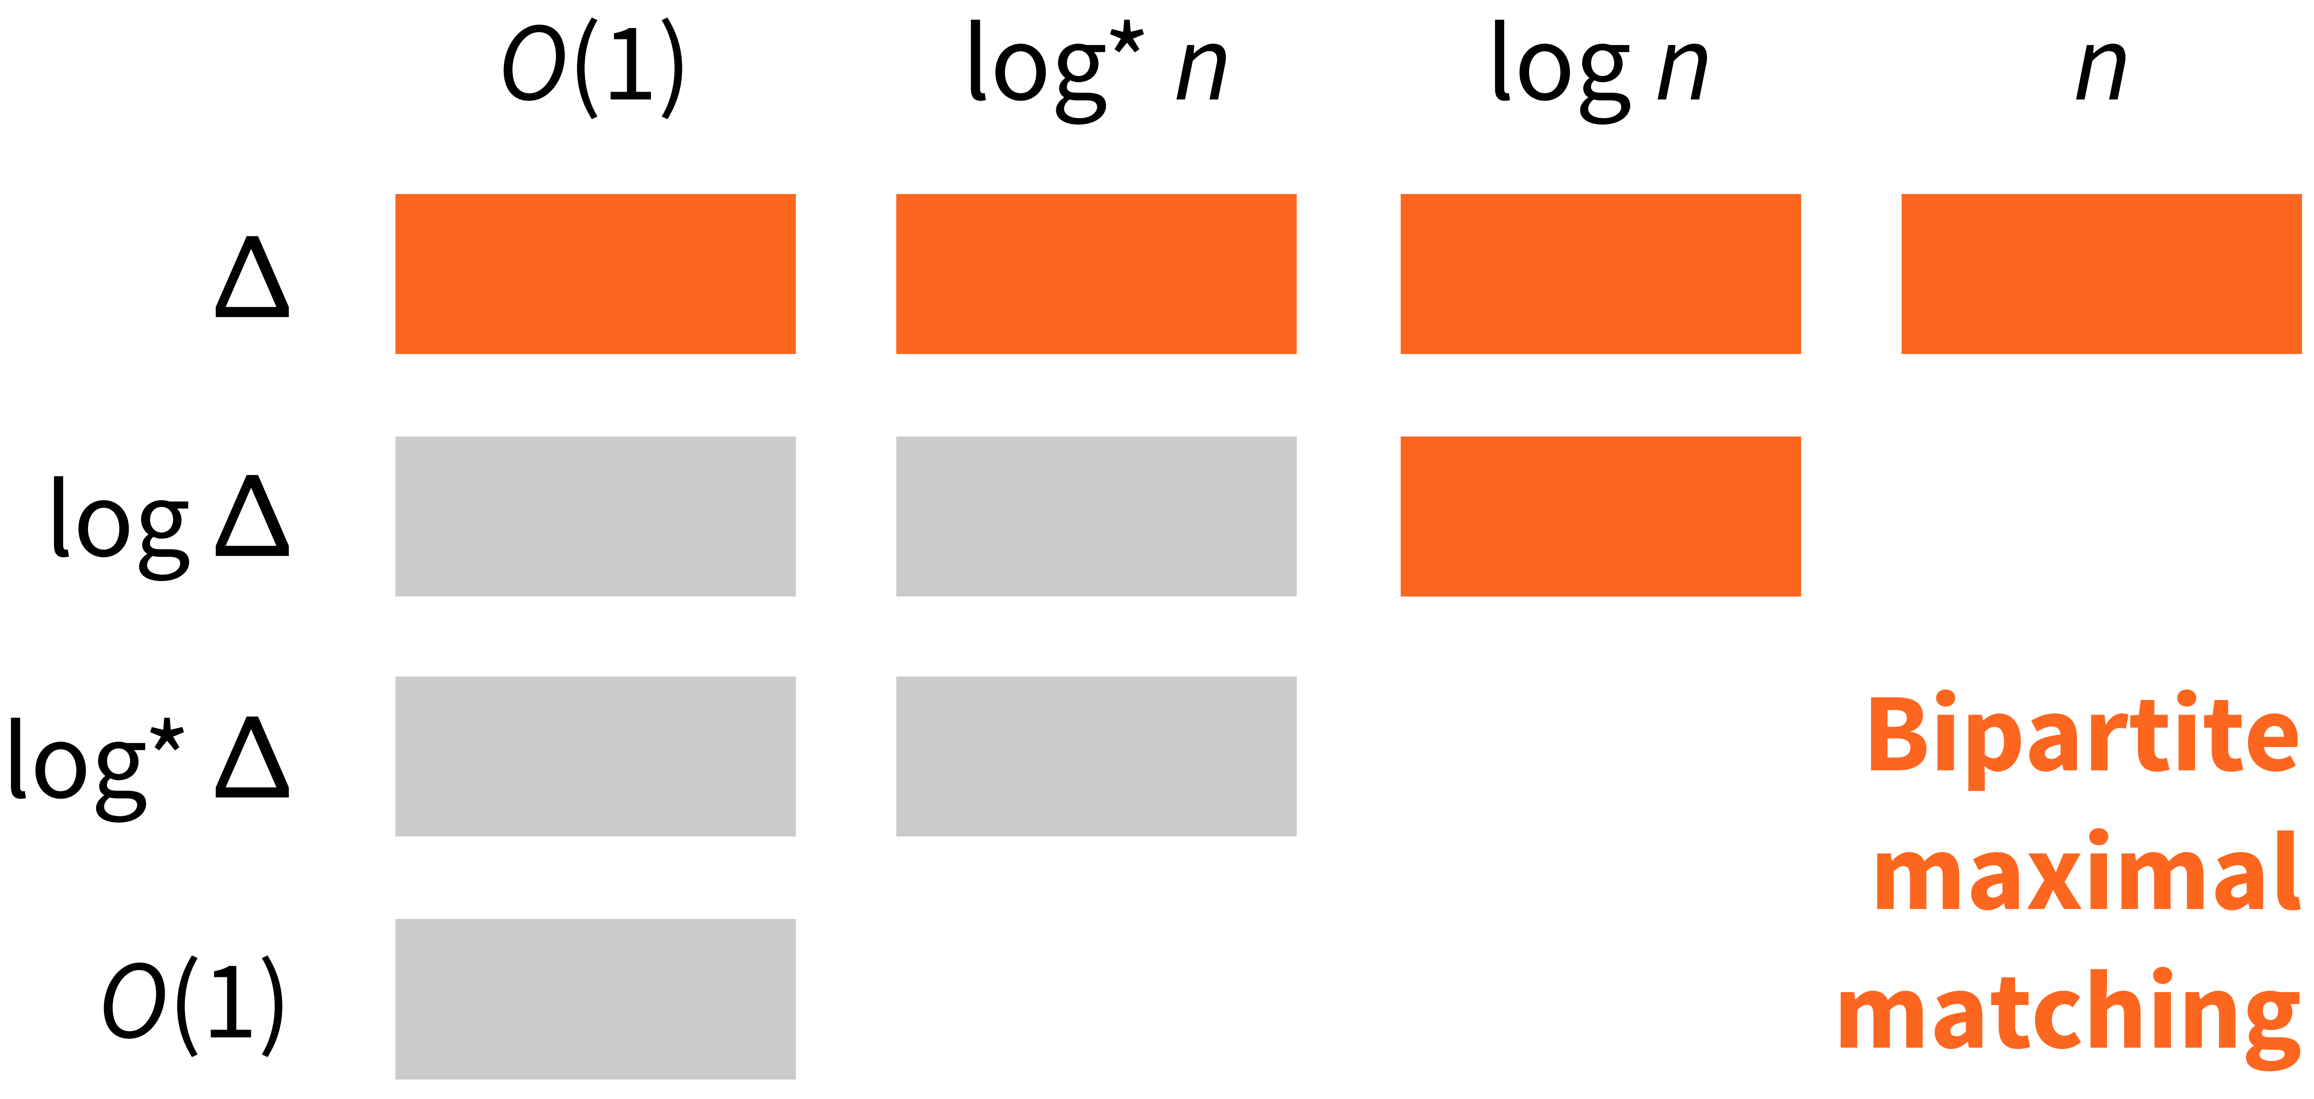
$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
Bipartite maximal matching

Landscape

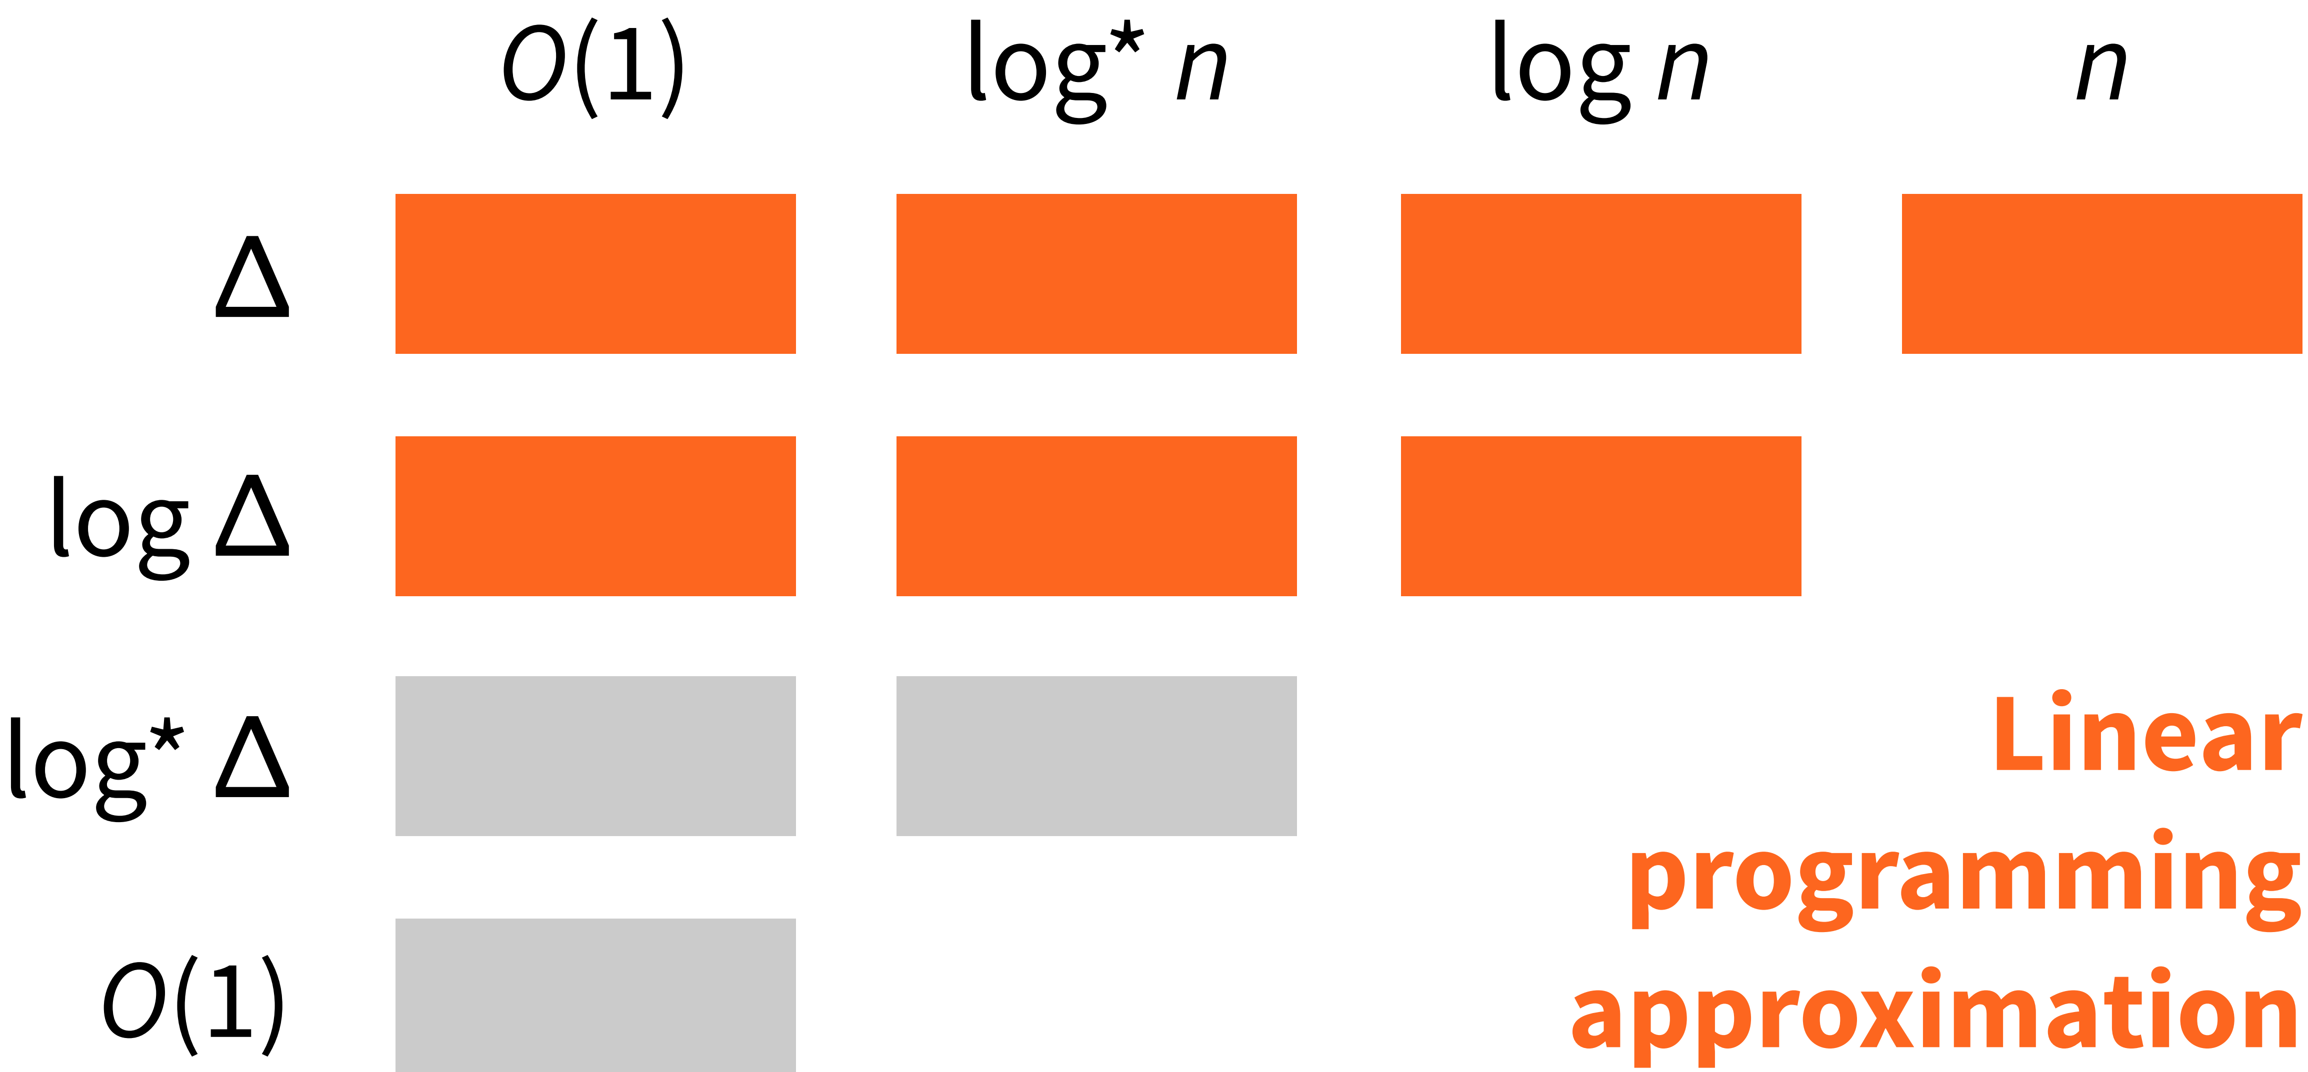
$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
Linear programming approximation

Landscape

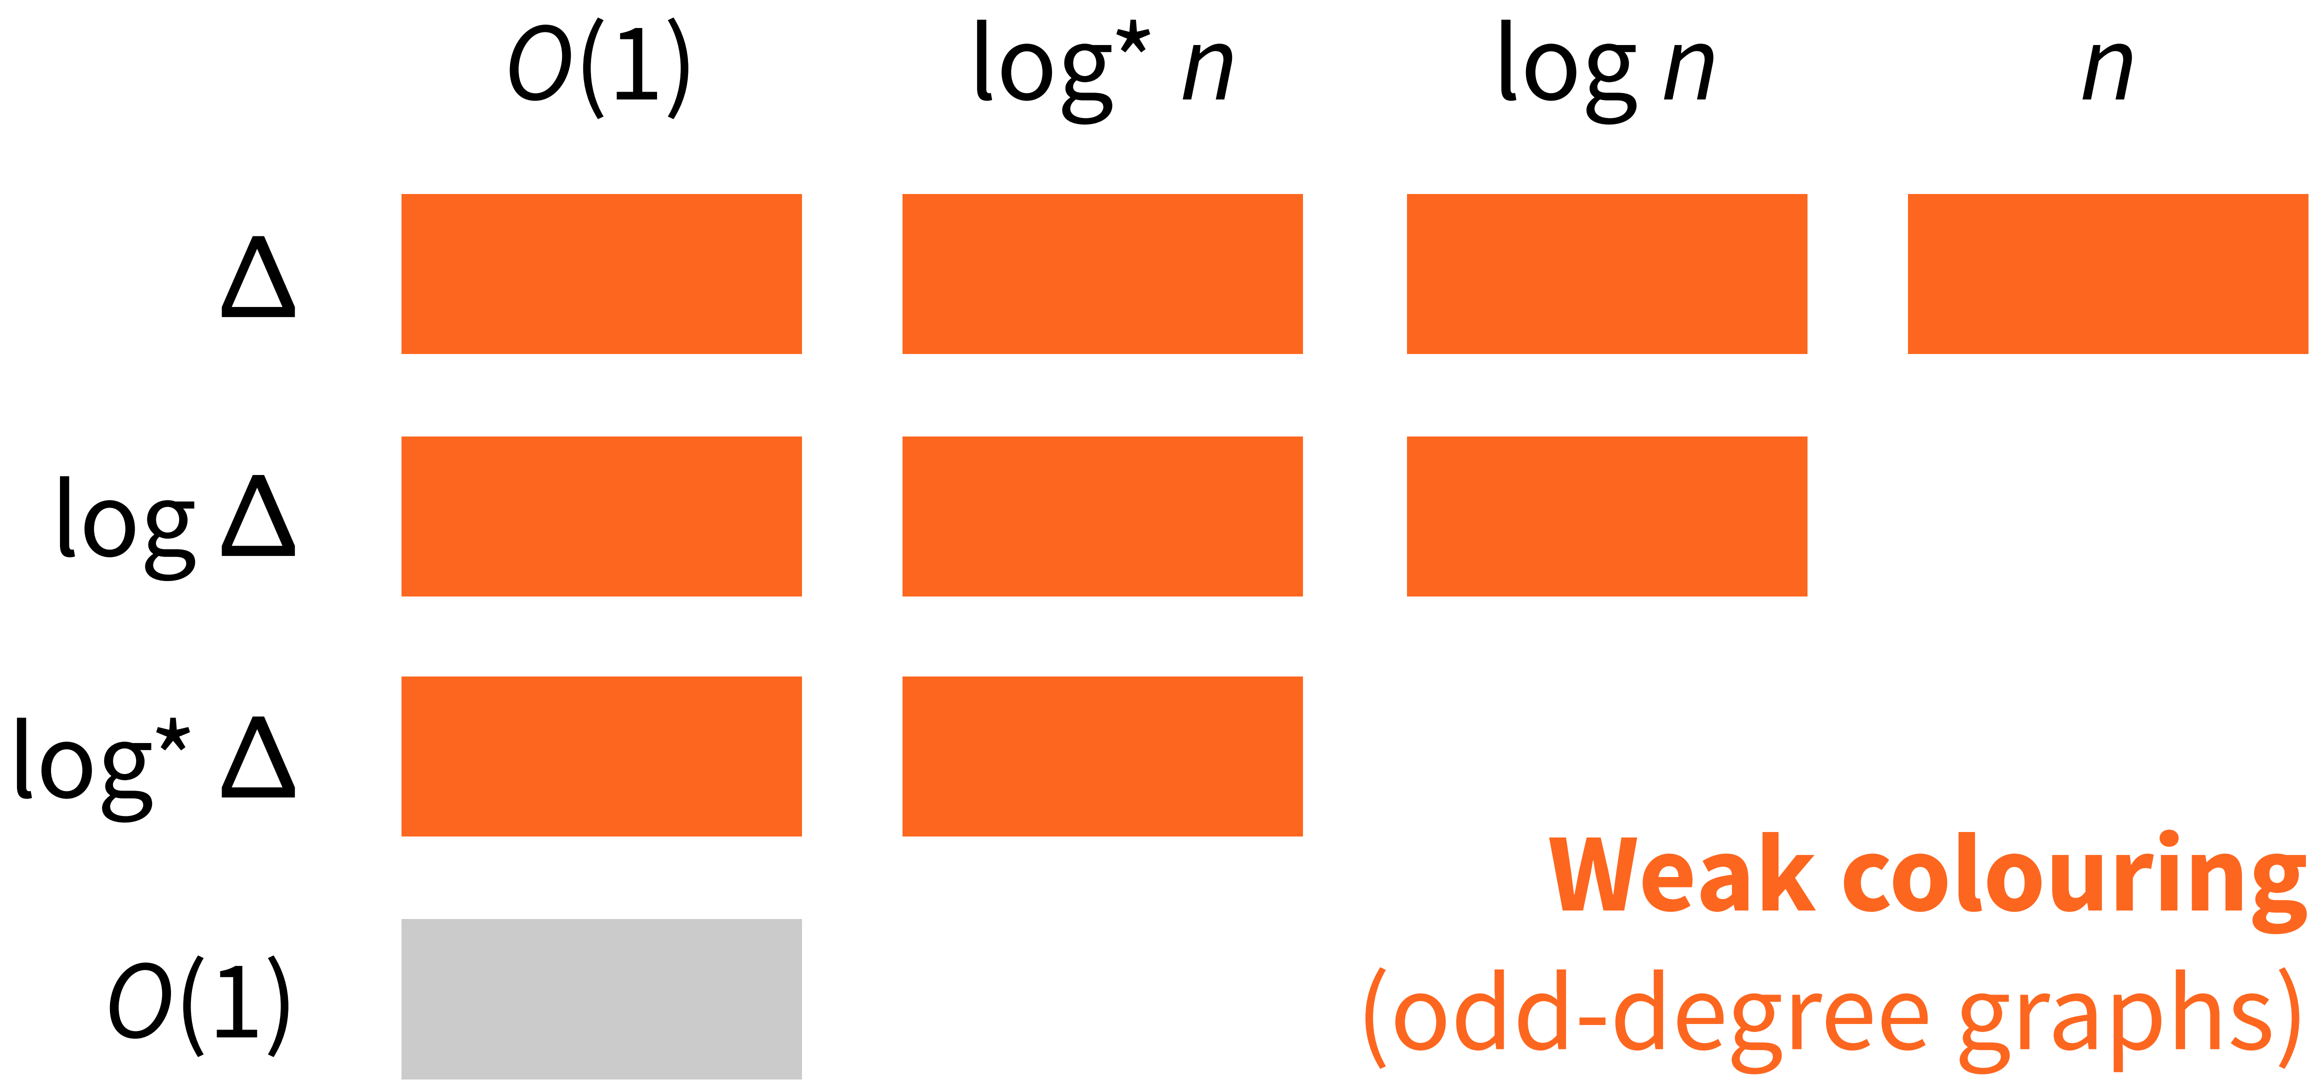
O(1)
log* n
log n
n
Δ
log Δ
log* Δ
O(1)
Weak colouring
(odd-degree graphs)

Landscape

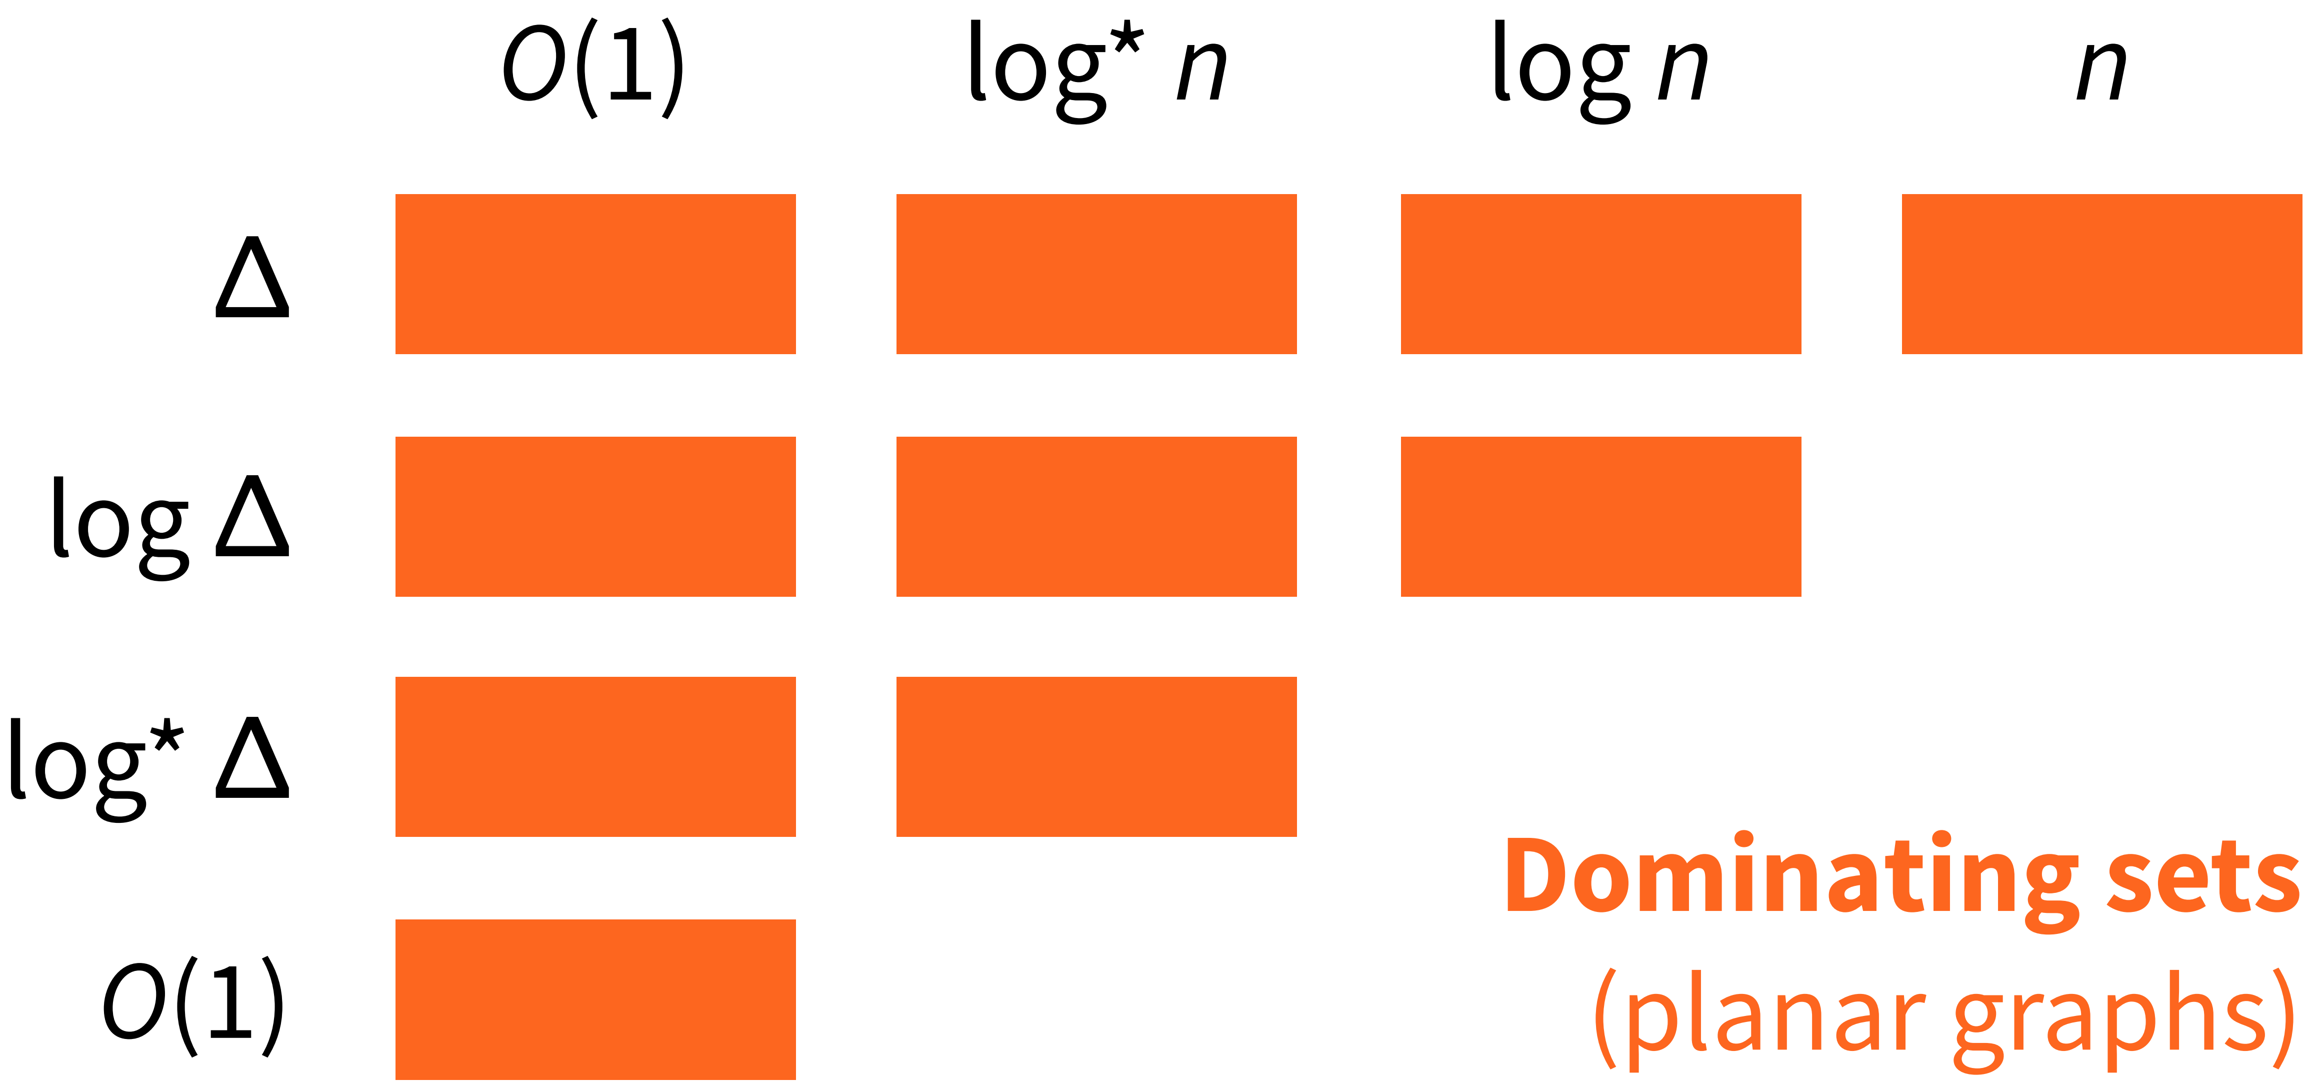

$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
Dominating sets
(planar graphs)

Landscape

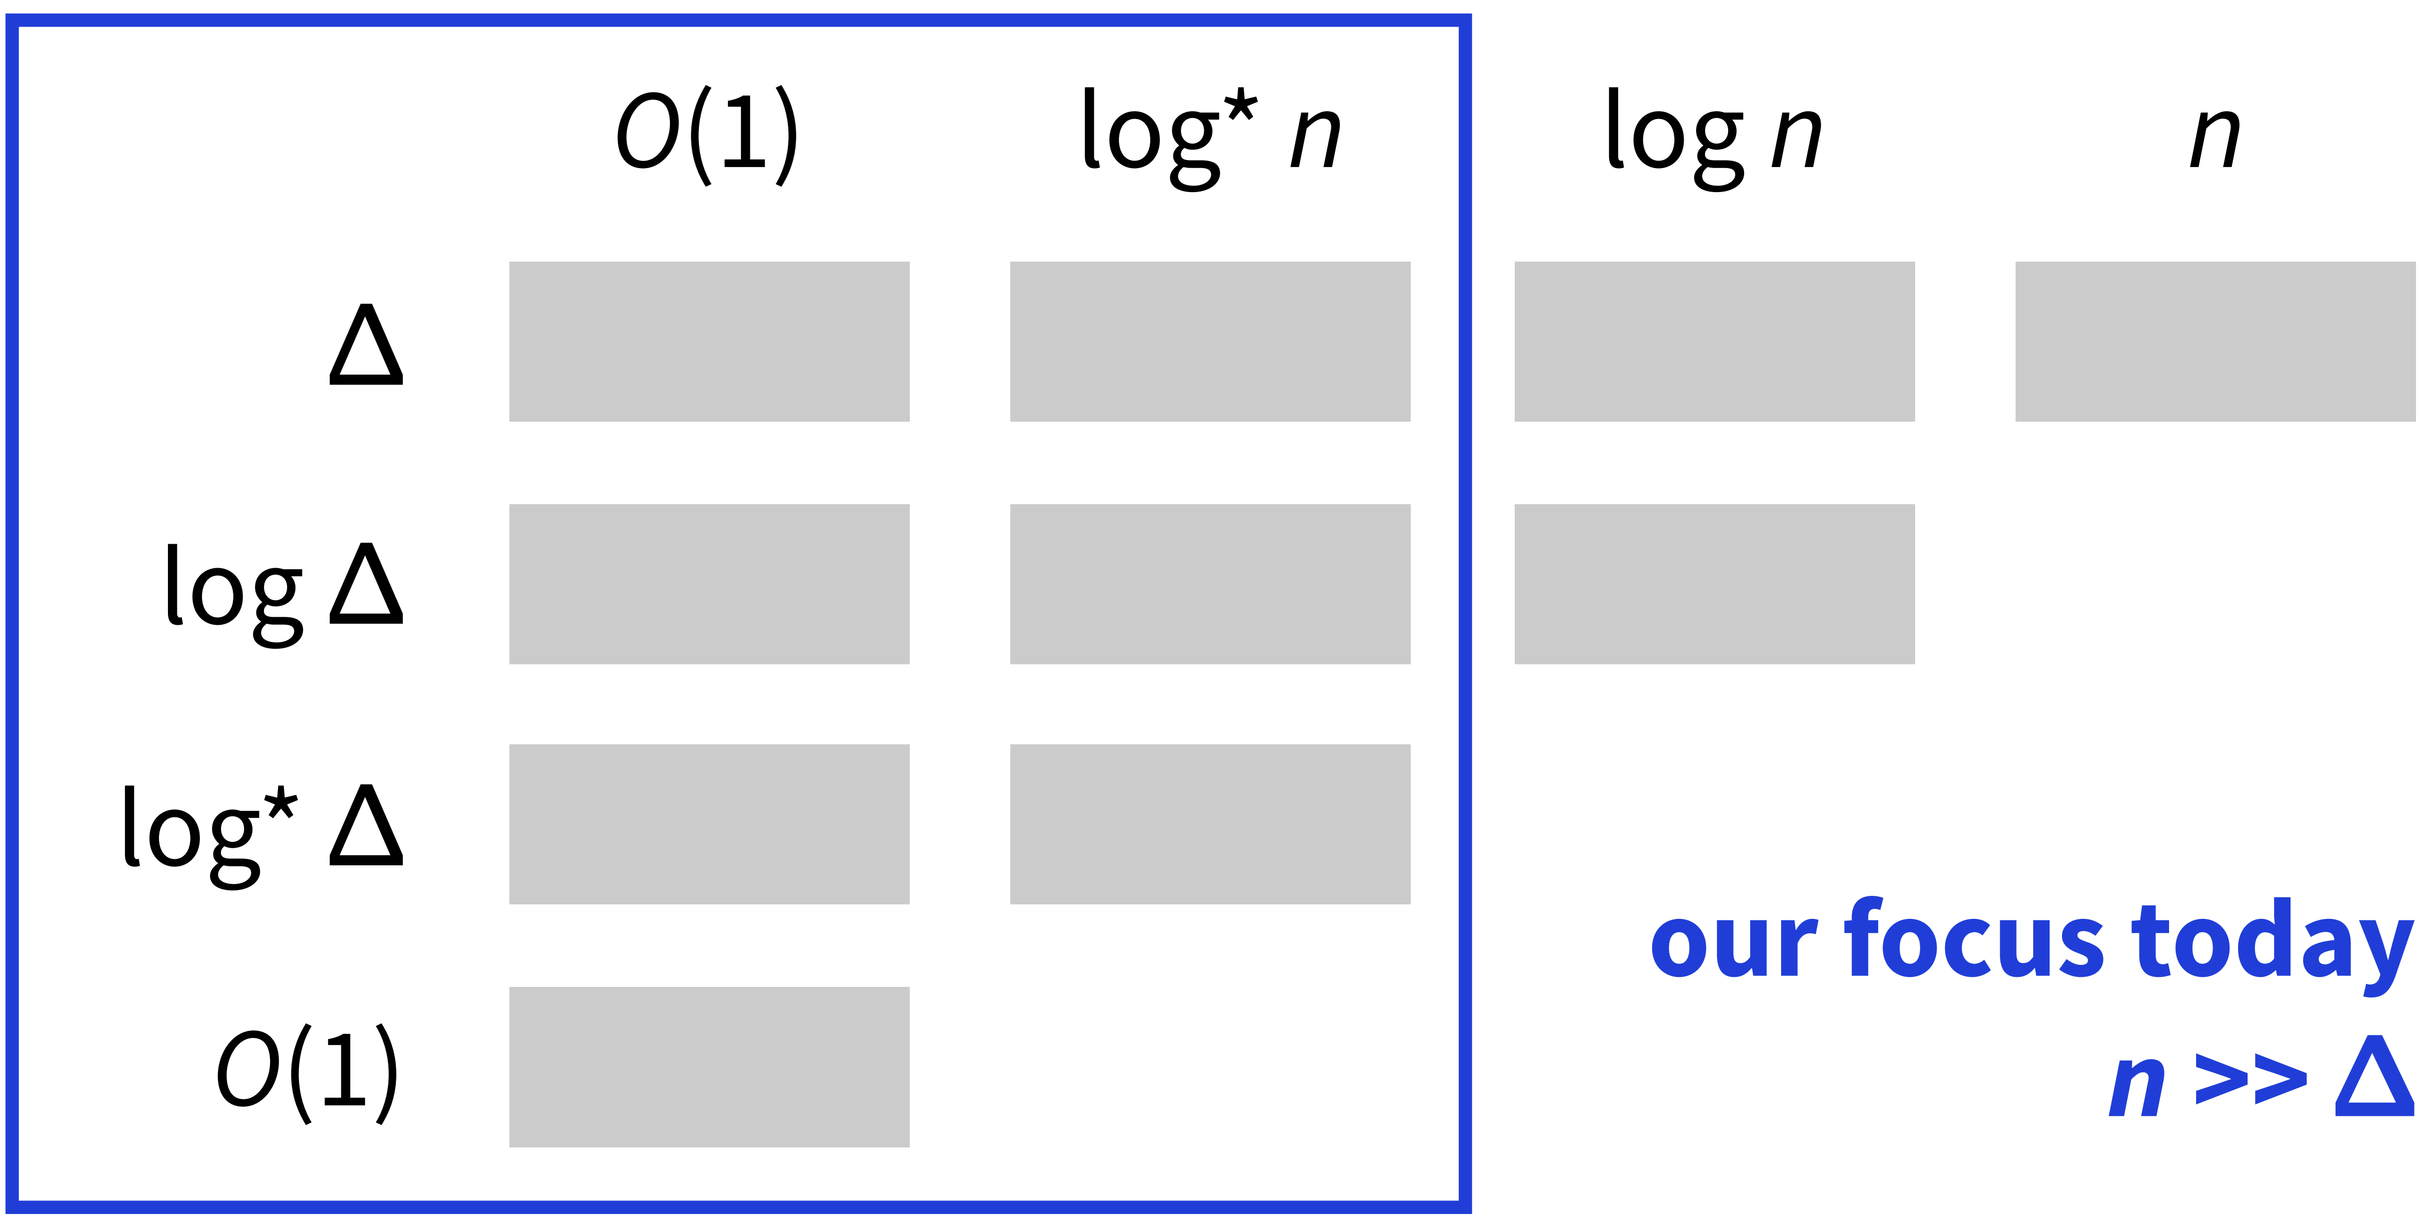

$O(1)$
$\log^* n$
$\log n$
n
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
our focus today
$n >> \Delta$

Typical state of the art

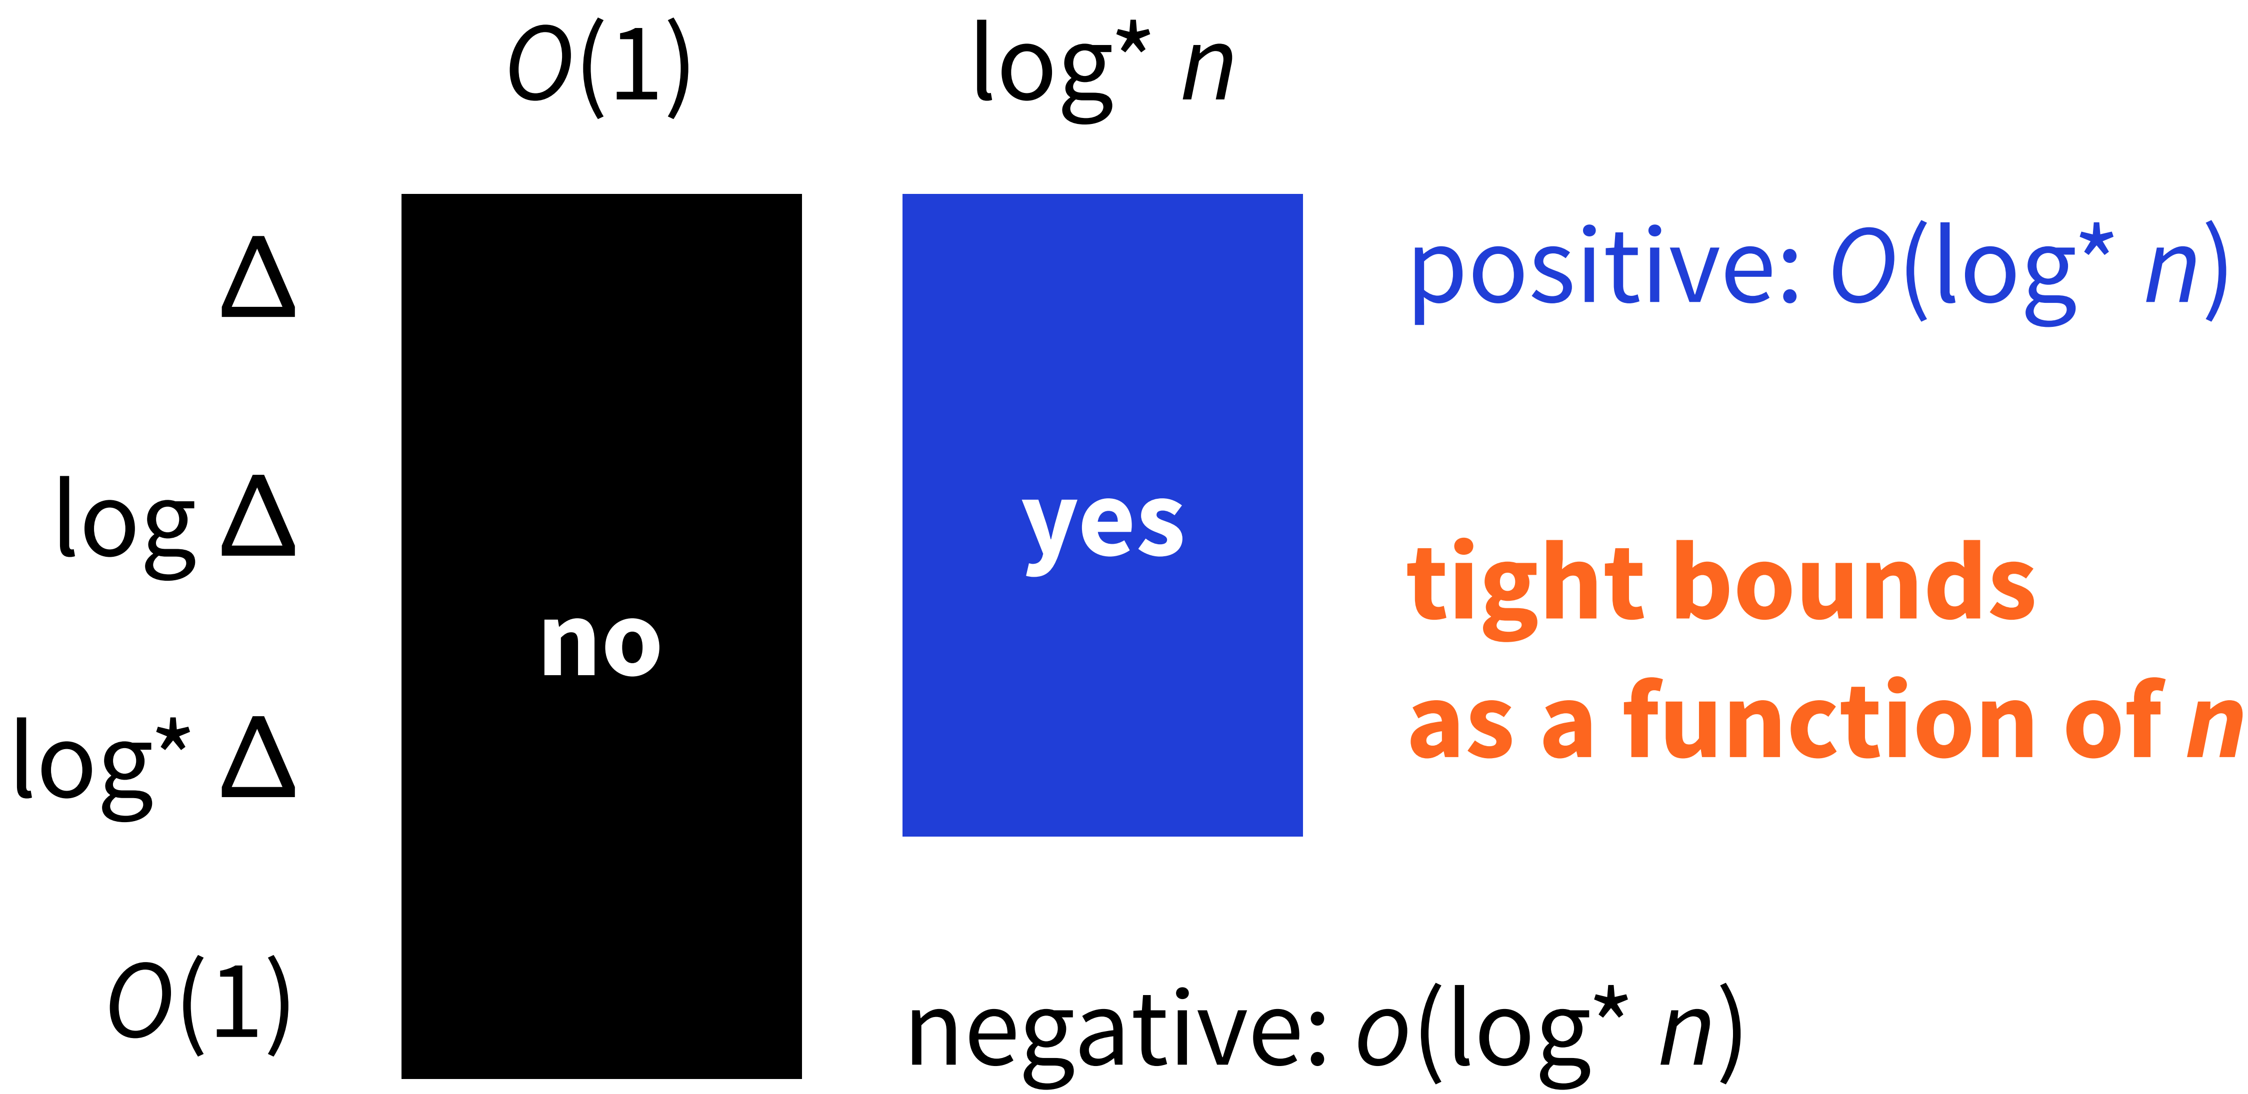

Typical state of the art

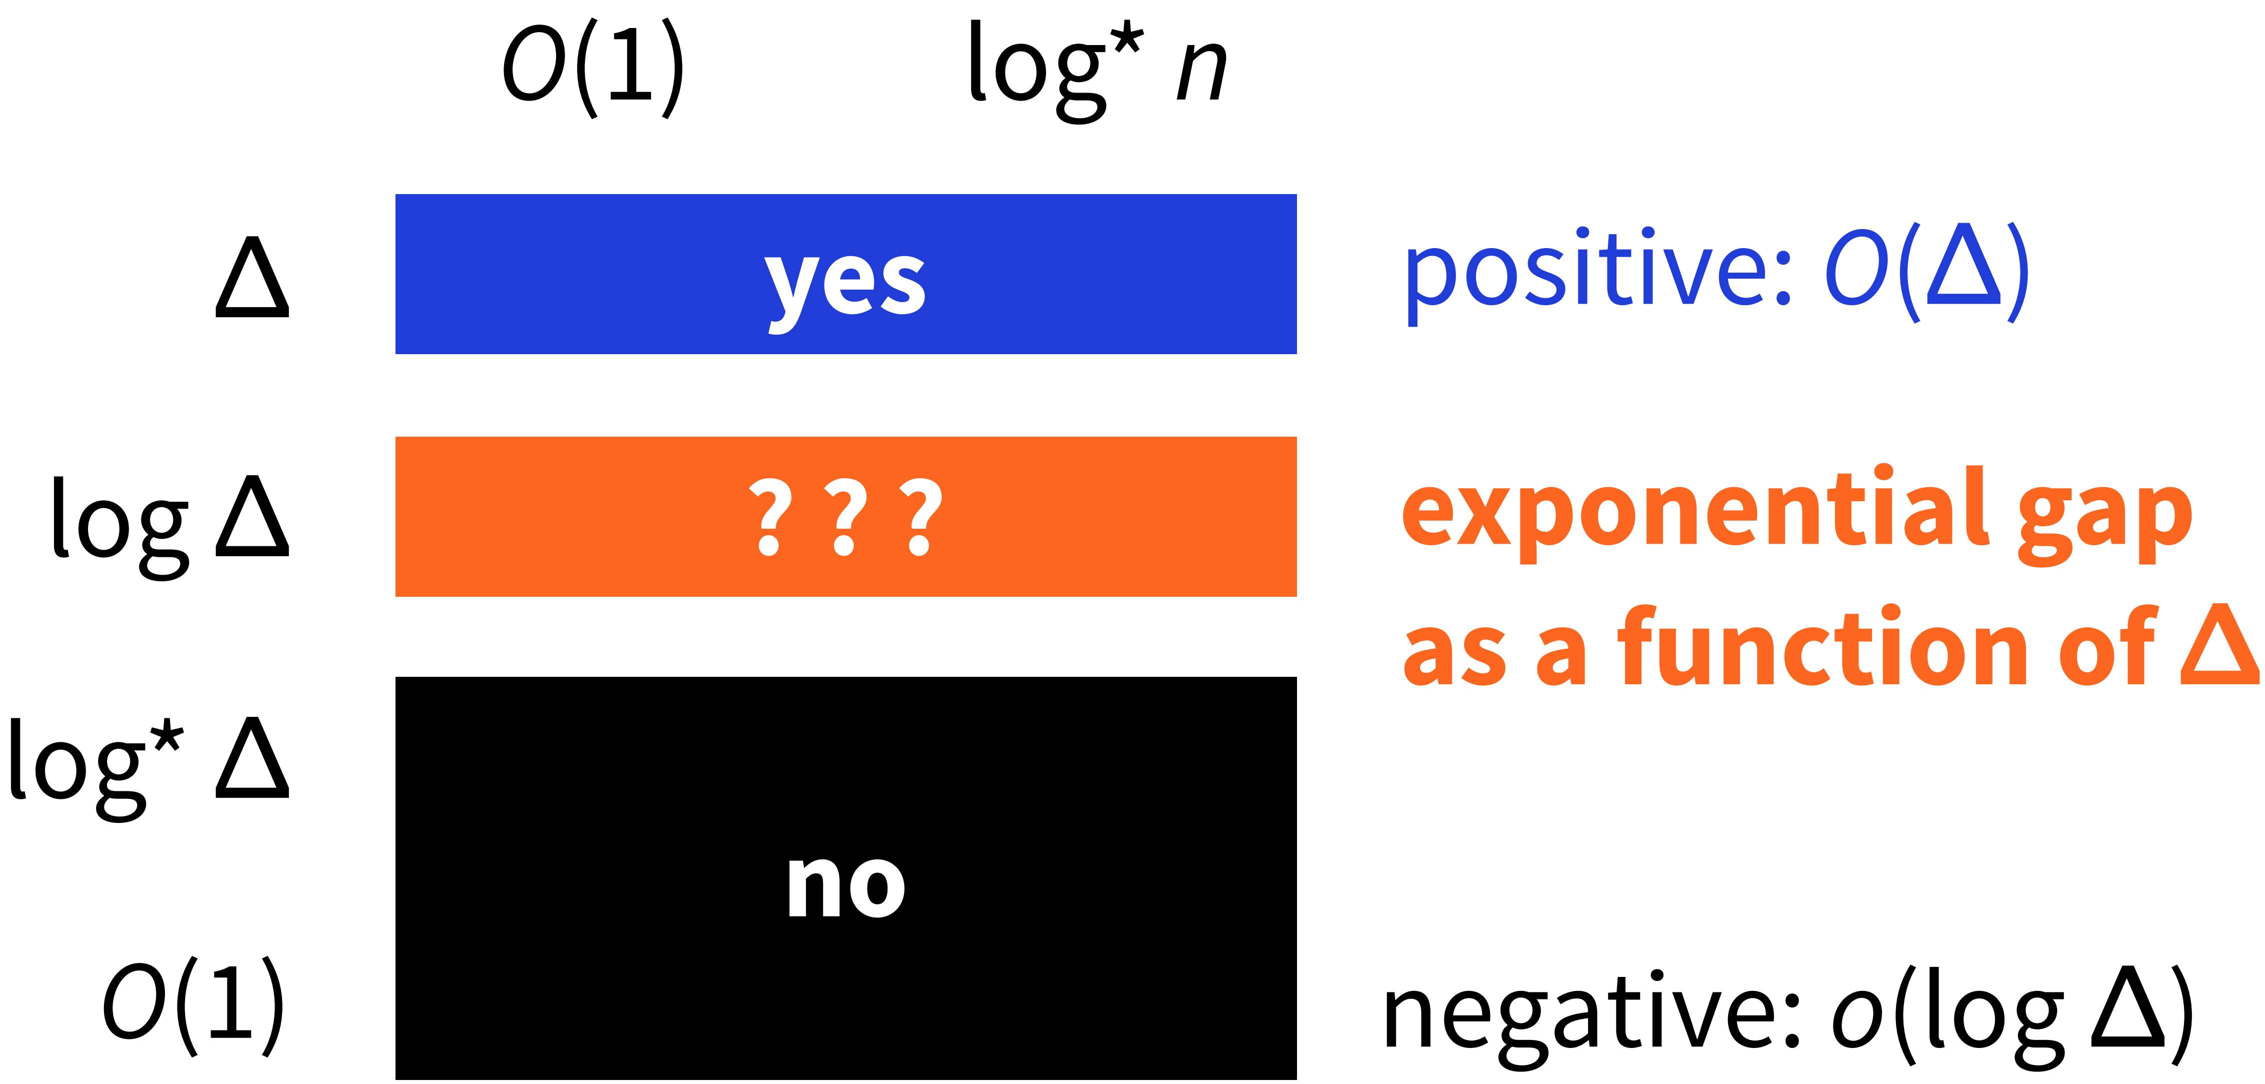

Typical state of the art

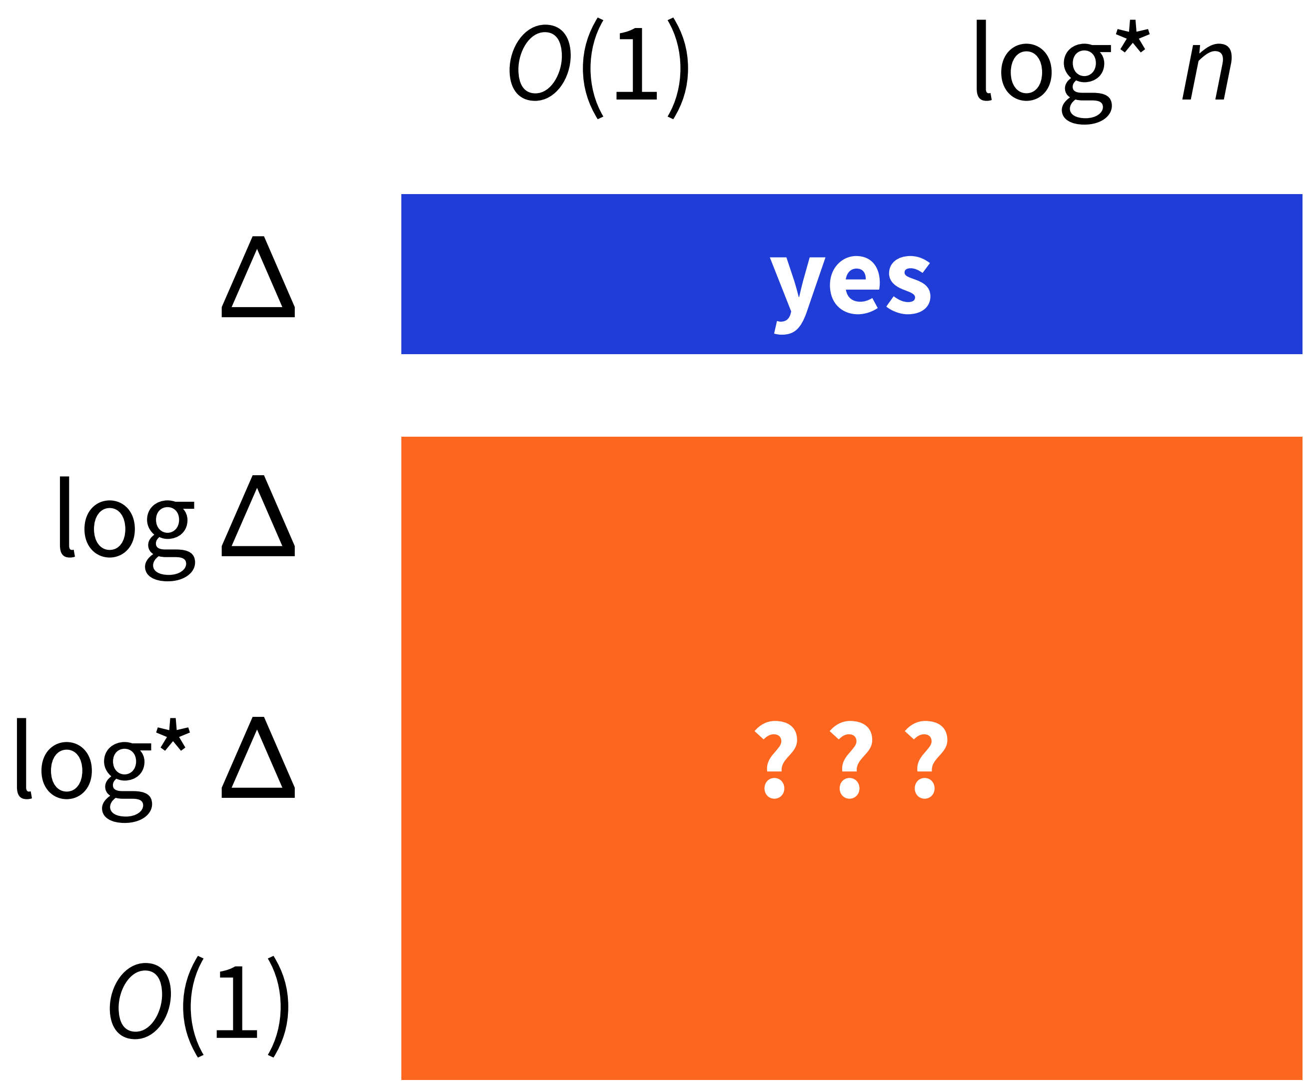

positive: $O(\Delta)$

exponential gap as a function of Δ — or much worse

negative: nothing

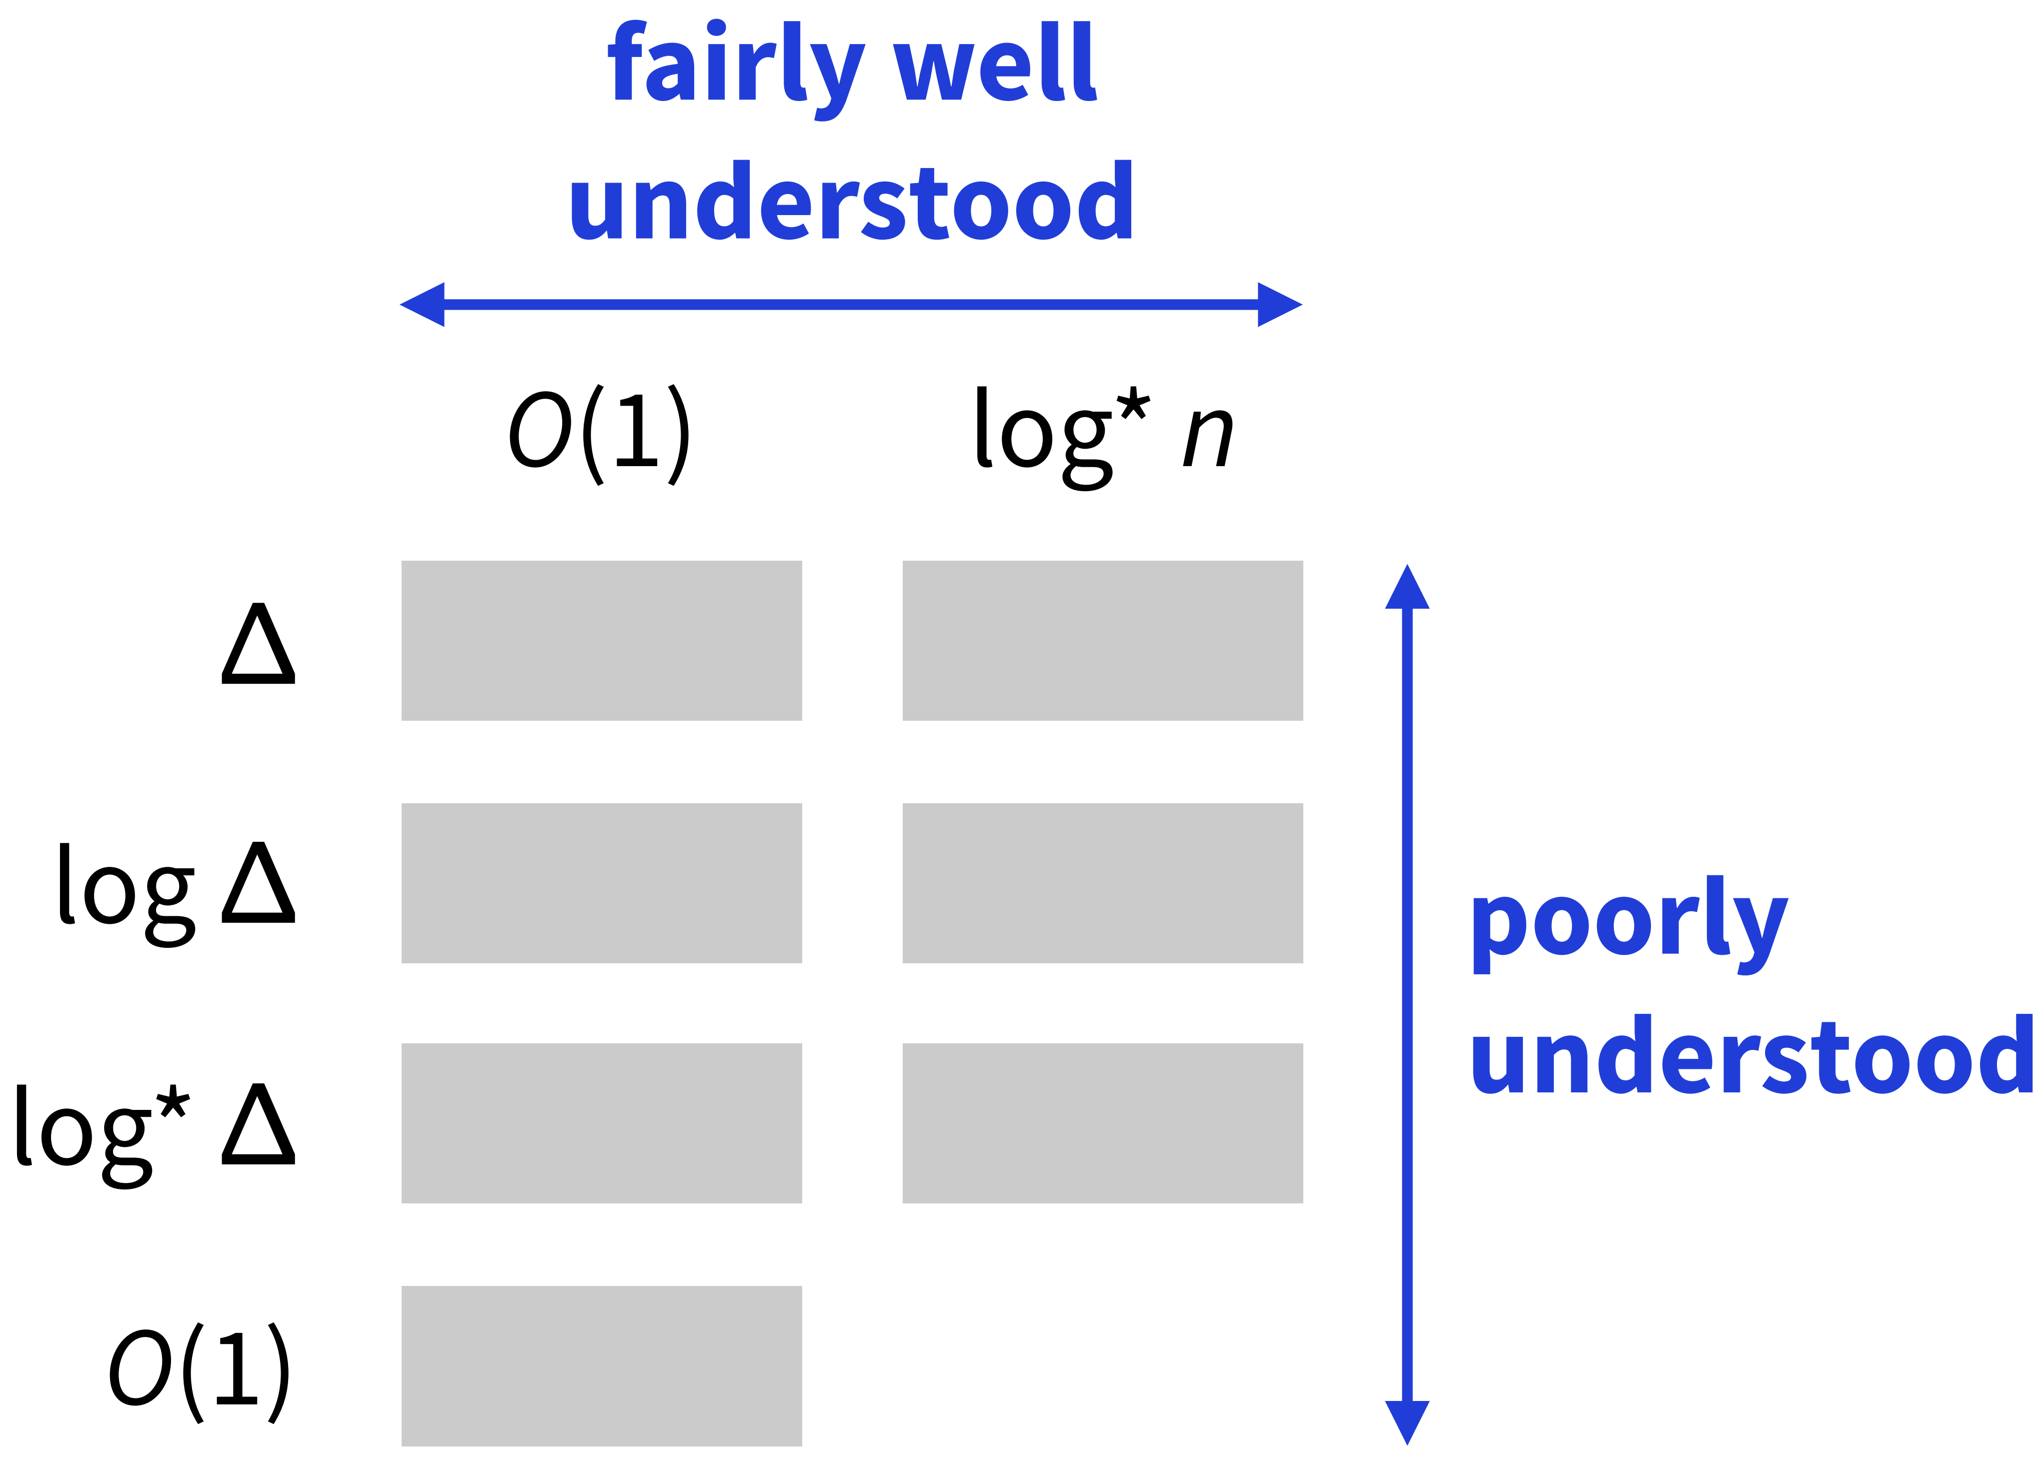

fairly well understood
O(1)
log* n
Δ
log Δ
log* Δ
O(1)
poorly understood

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: **Bipartite maximal matching**

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

Example: **Bipartite maximal matching**

- ***O*(Δ): trivial for Δ-regular graphs**
  - Hańćkowiak et al. (1998)
- ***O*(1): unknown for Δ-regular graphs**

Example:
Semi-matching

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

Example: **Semi-matching**

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

Example:

Weak colouring

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

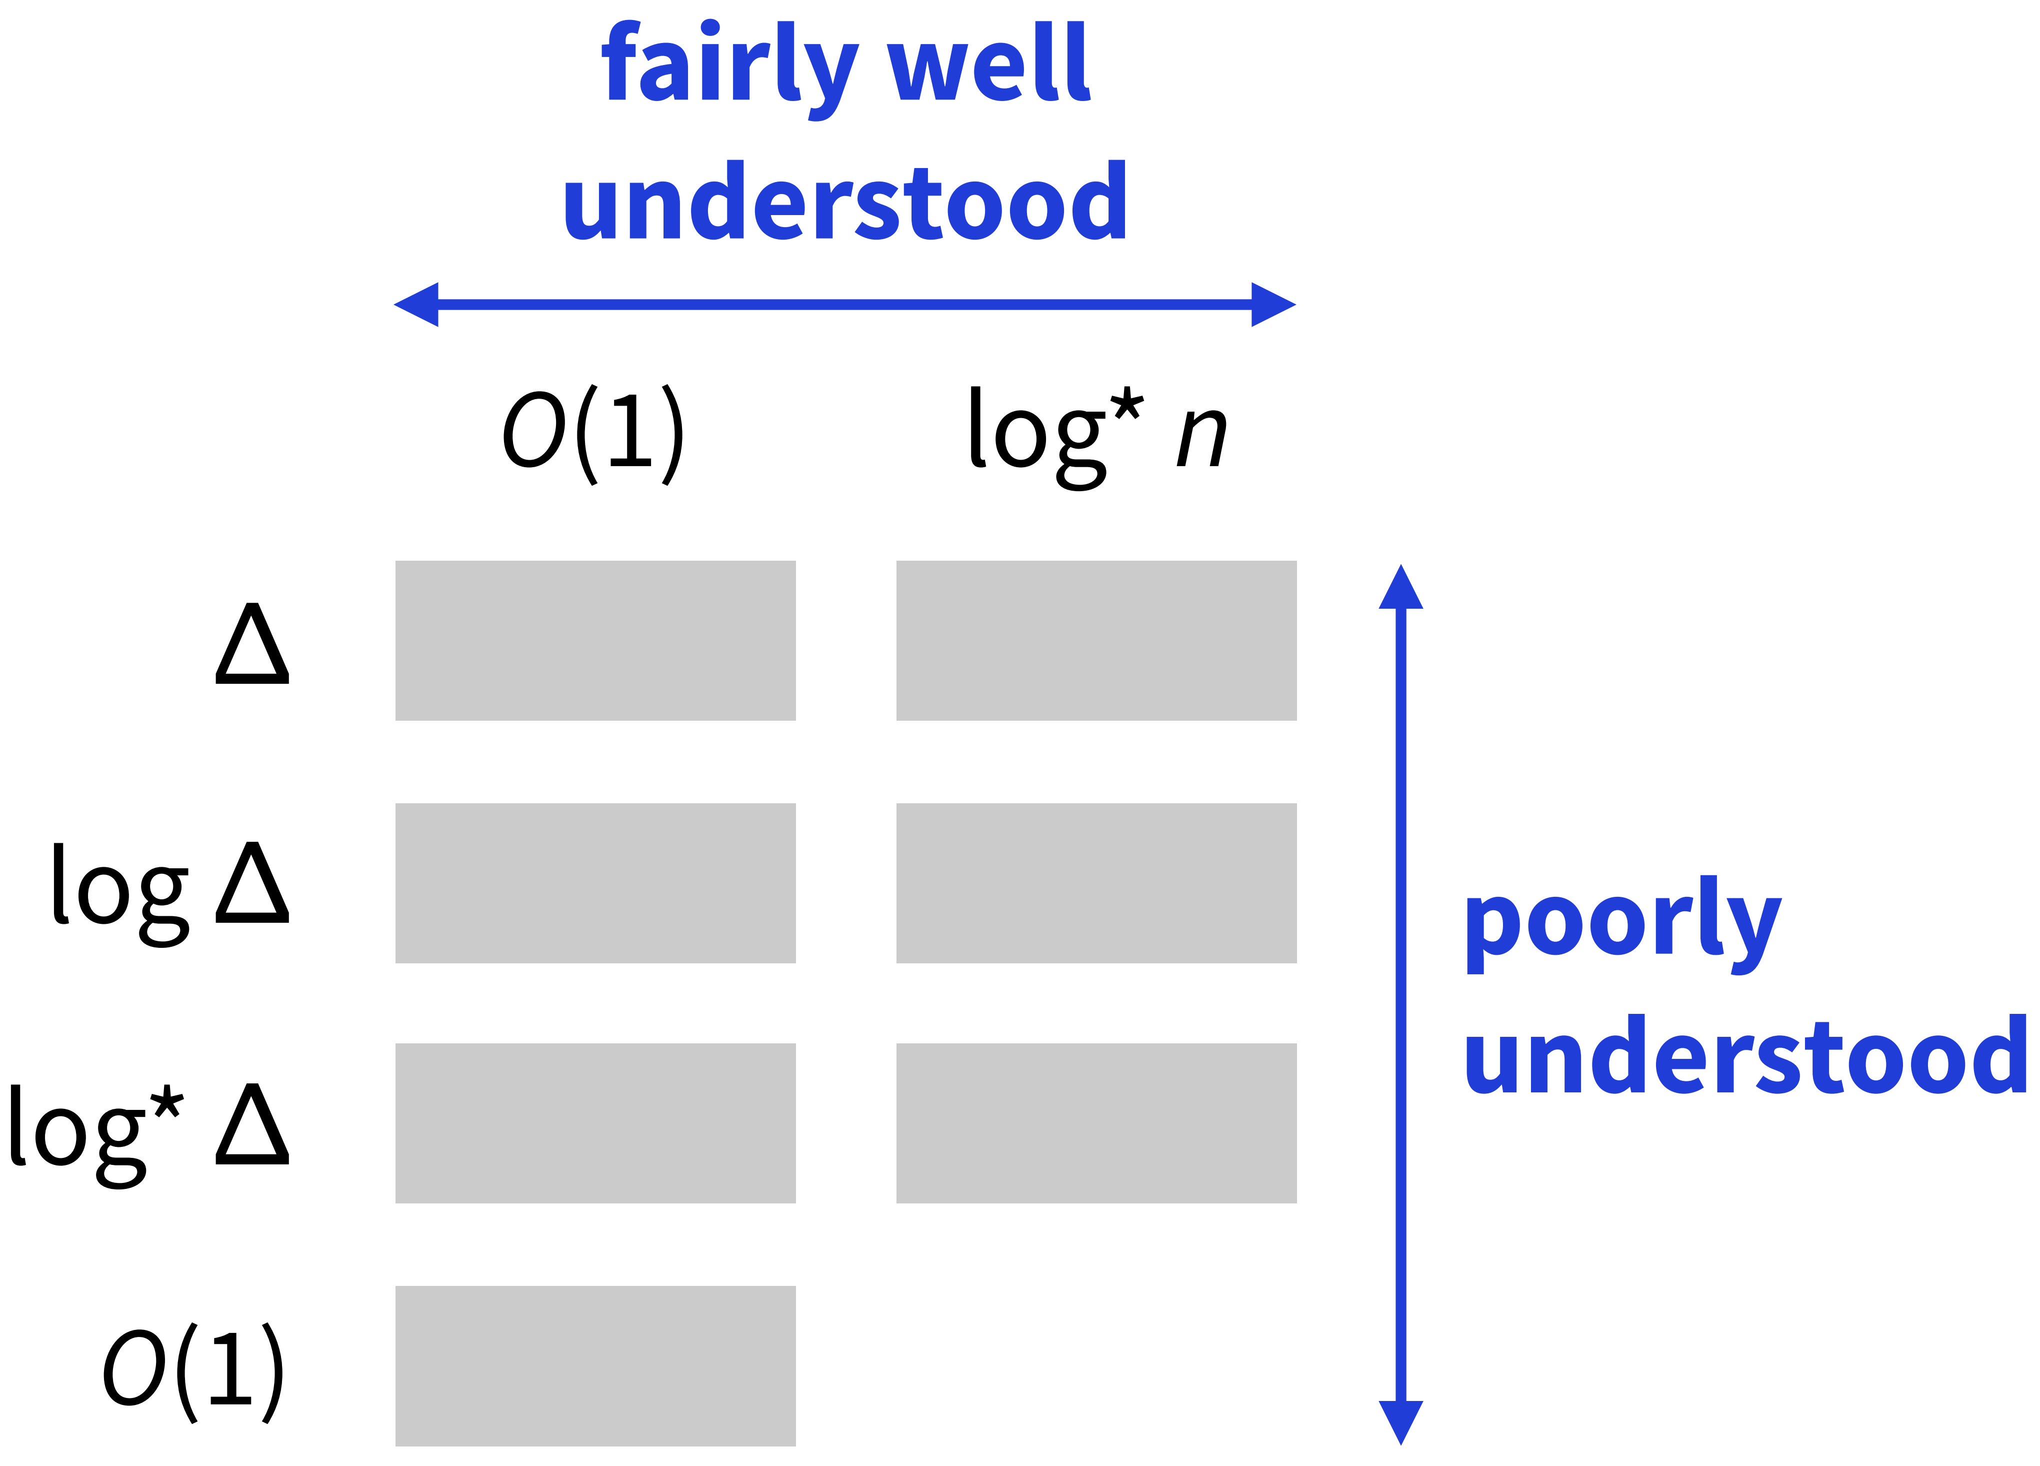

fairly well understood
O(1)
log* n
Δ
log Δ
log* Δ
O(1)
poorly understood

Orthogonal challenges?

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

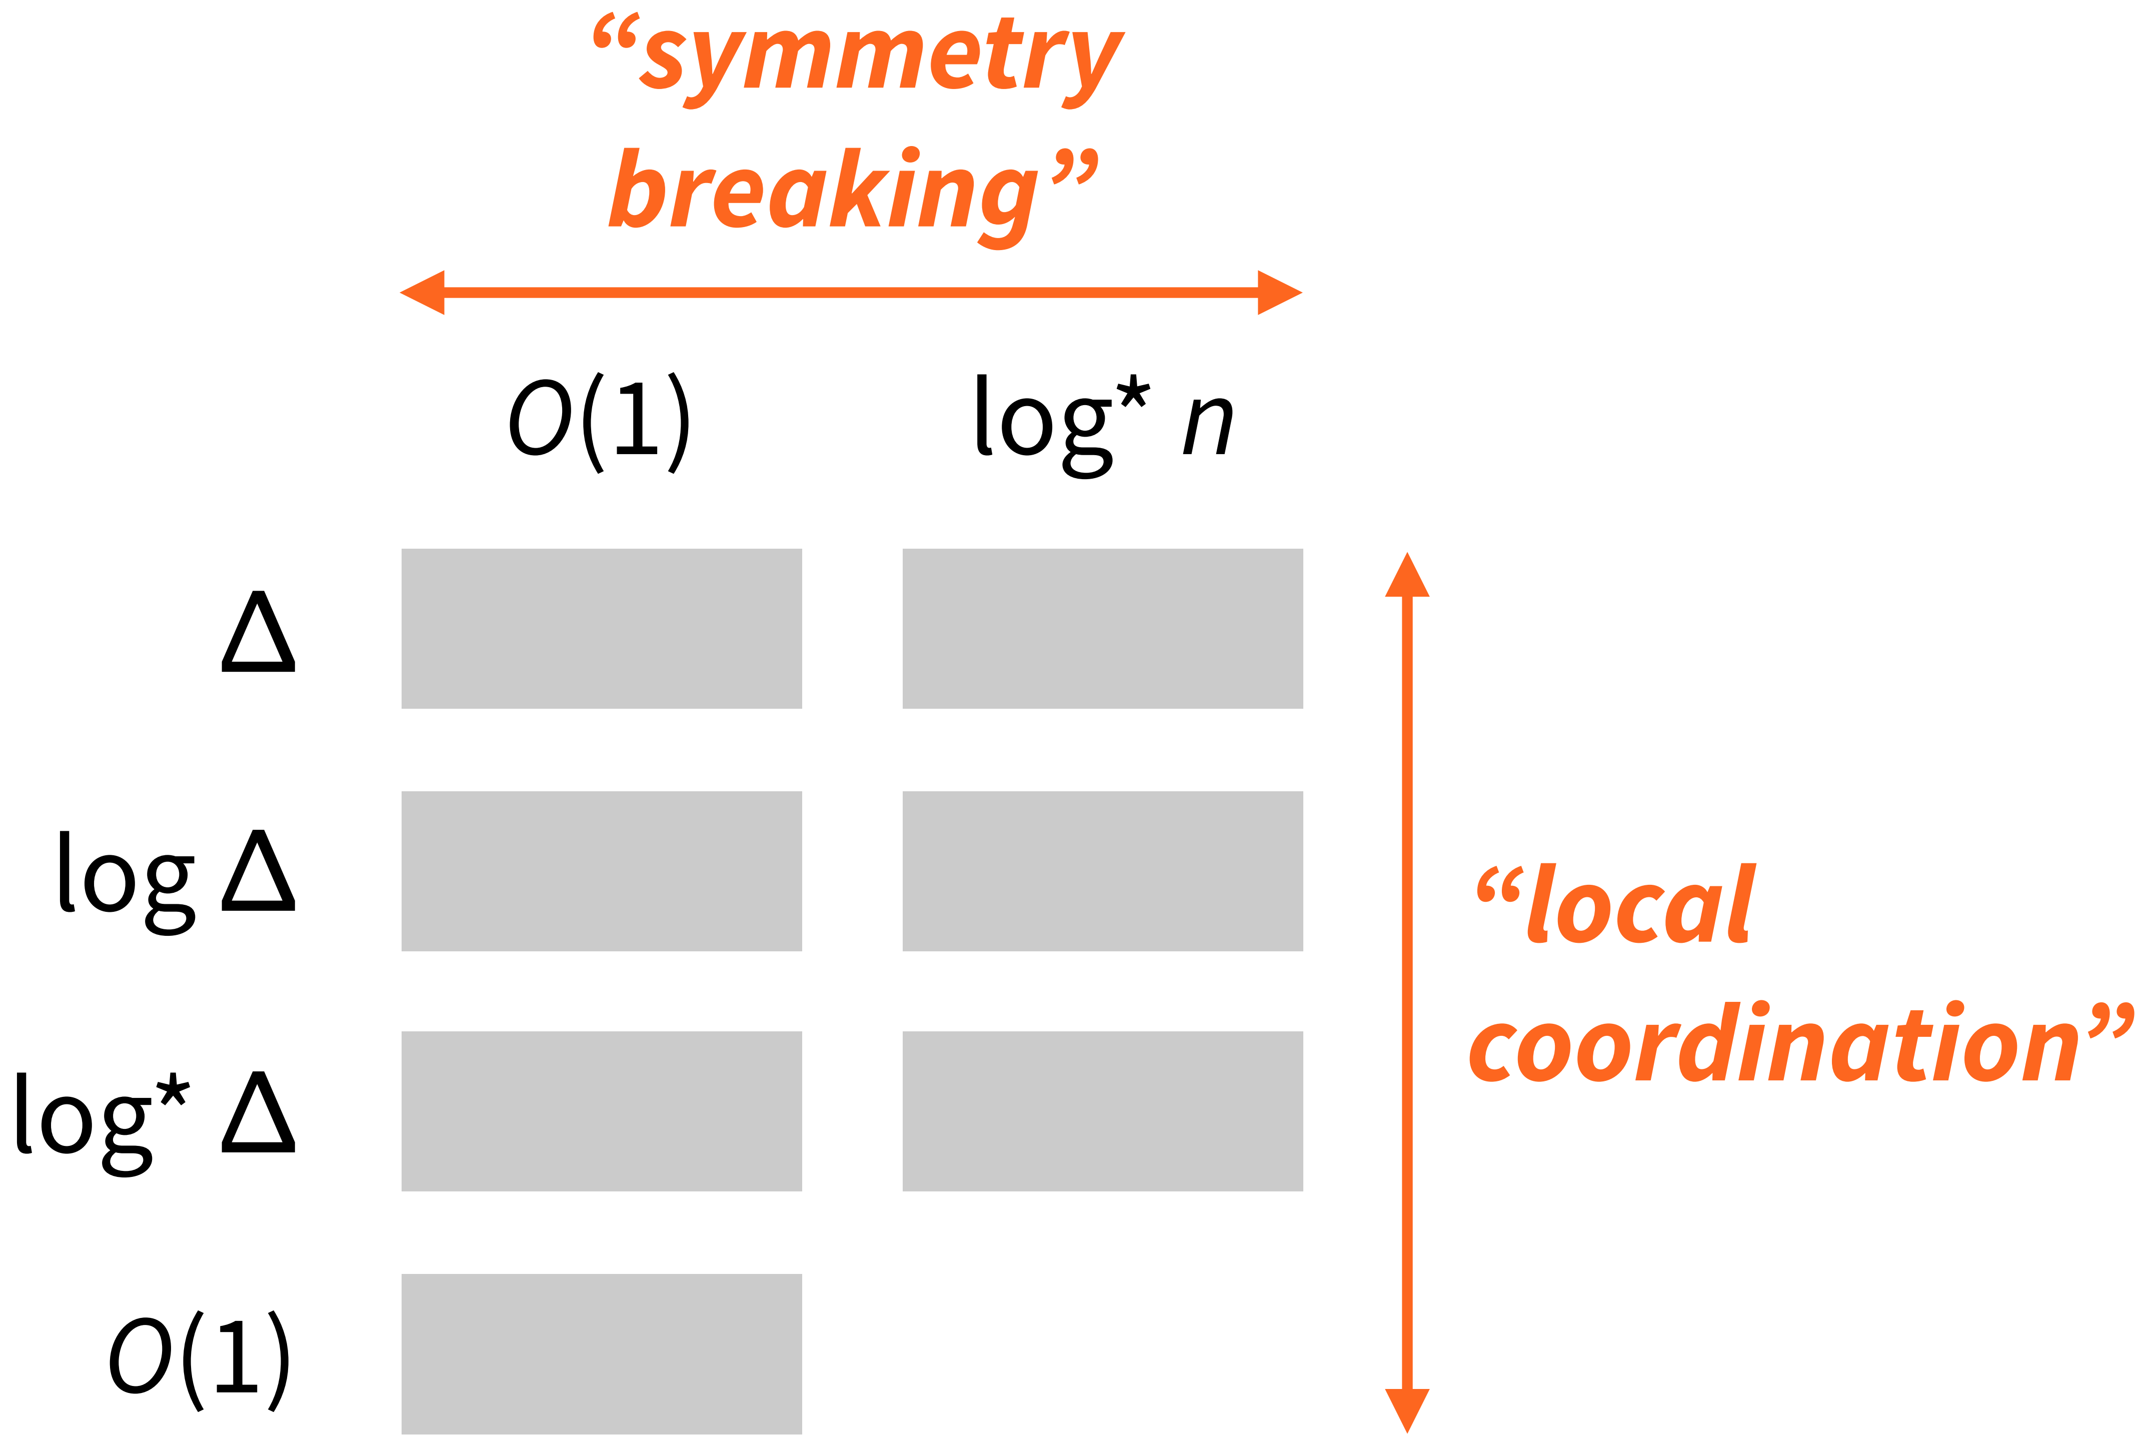
"symmetry breaking"
$O(1)$
$\log^* n$
Δ
$\log \Delta$
$\log^* \Delta$
$O(1)$
"local coordination"

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

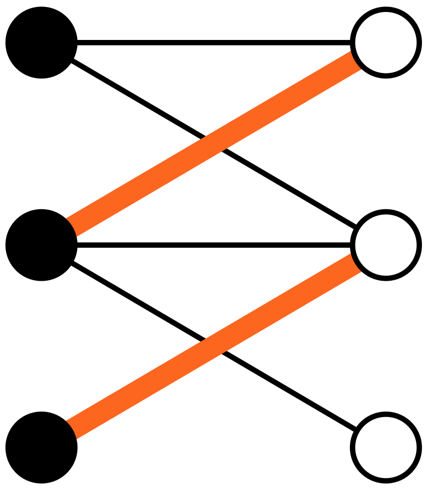

- **"Coordination" ≈ one by one traversal**

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

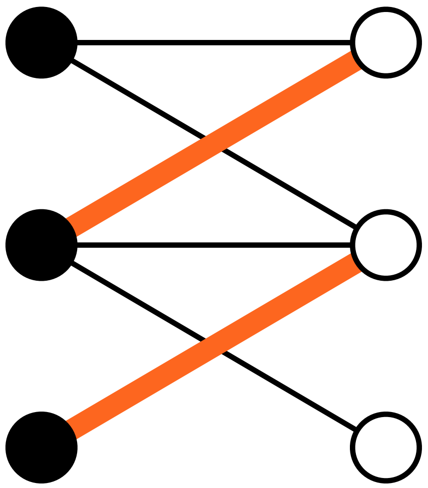

- **Clearly $O(\Delta)$, but is this tight?**

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)

Good news

- **We are finally making some progress!**
- **Key problem:** ***maximal matching***
- **Start with a "toy model":** ***edge colouring model***

EC: edge colouring

No identifiers

No orientations

Edges coloured
with $O(\Delta)$ colours

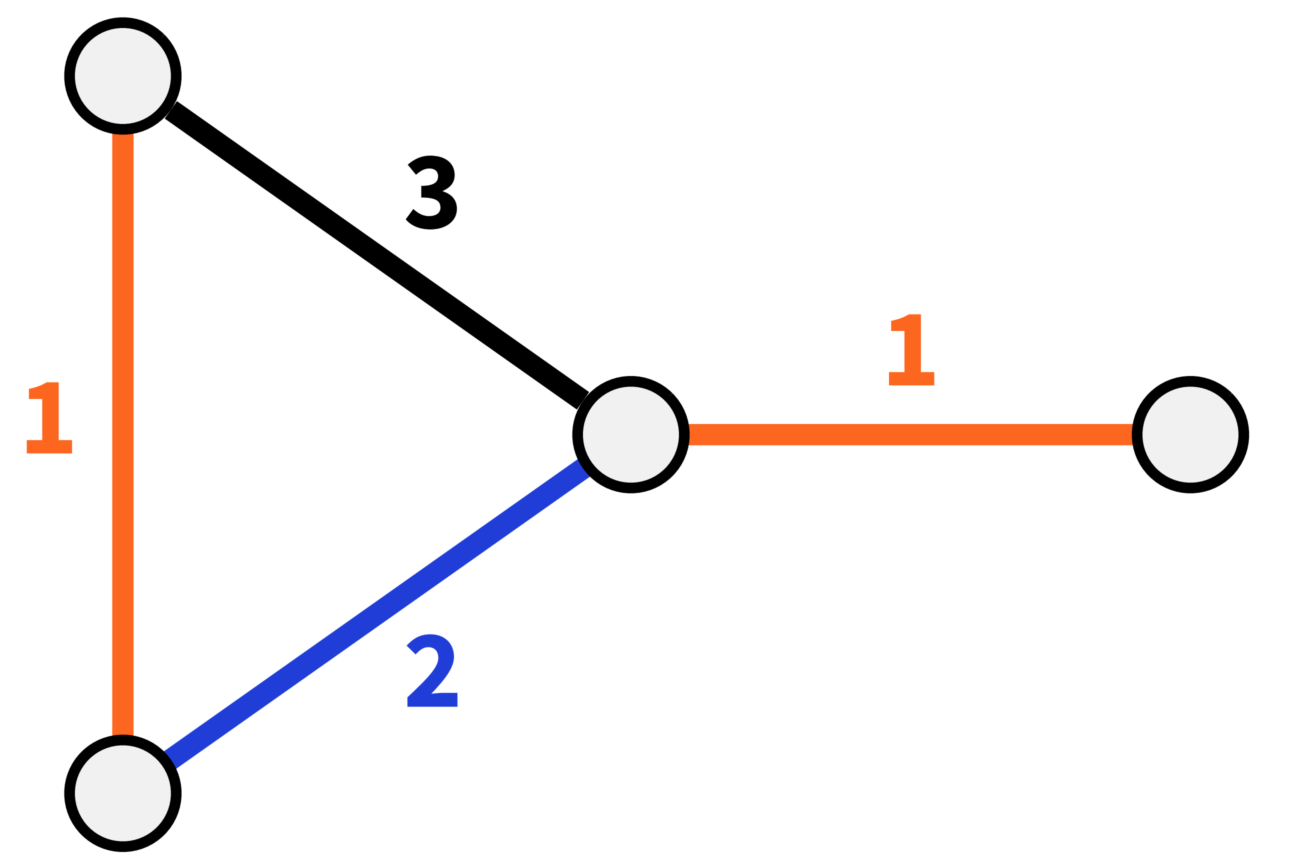

Recent progress

- **Maximal matching in *EC model***
- **$O(\Delta)$: trivial**
  - greedily by colour classes
- **$o(\Delta)$: not possible**
  - PODC 2012

What about the LOCAL model?

- **Not yet there with maximal matchings…**
- **But we can prove lower bounds for maximal *fractional* matchings!**

Matching

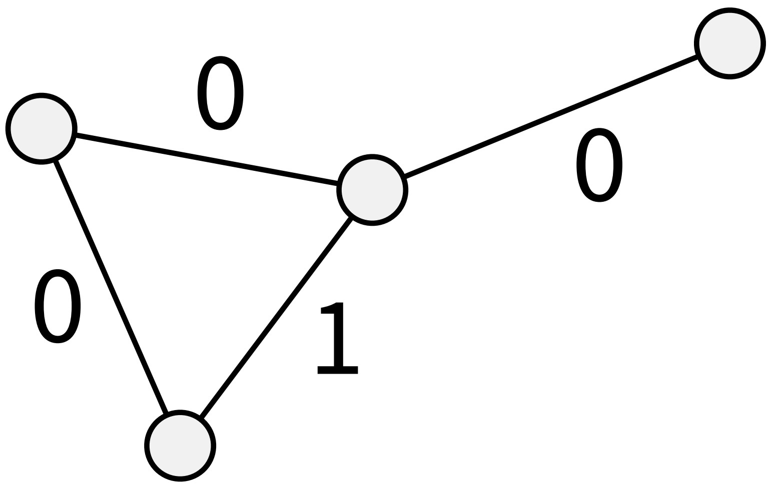

- **Edges labelled with integers {0, 1}**
- **Sum of incident edges at most 1**
- **Maximal matching:**
  cannot increase the value of any label

Fractional matching

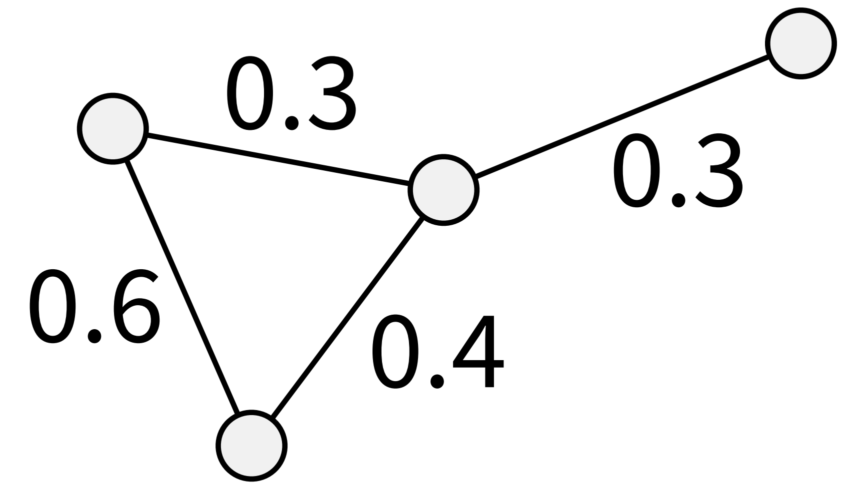

- **Edges labelled with real numbers [0, 1]**
- **Sum of incident edges at most 1**
- **Maximal fractional matching:**
  cannot increase the value of any label

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…

Recent progress

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

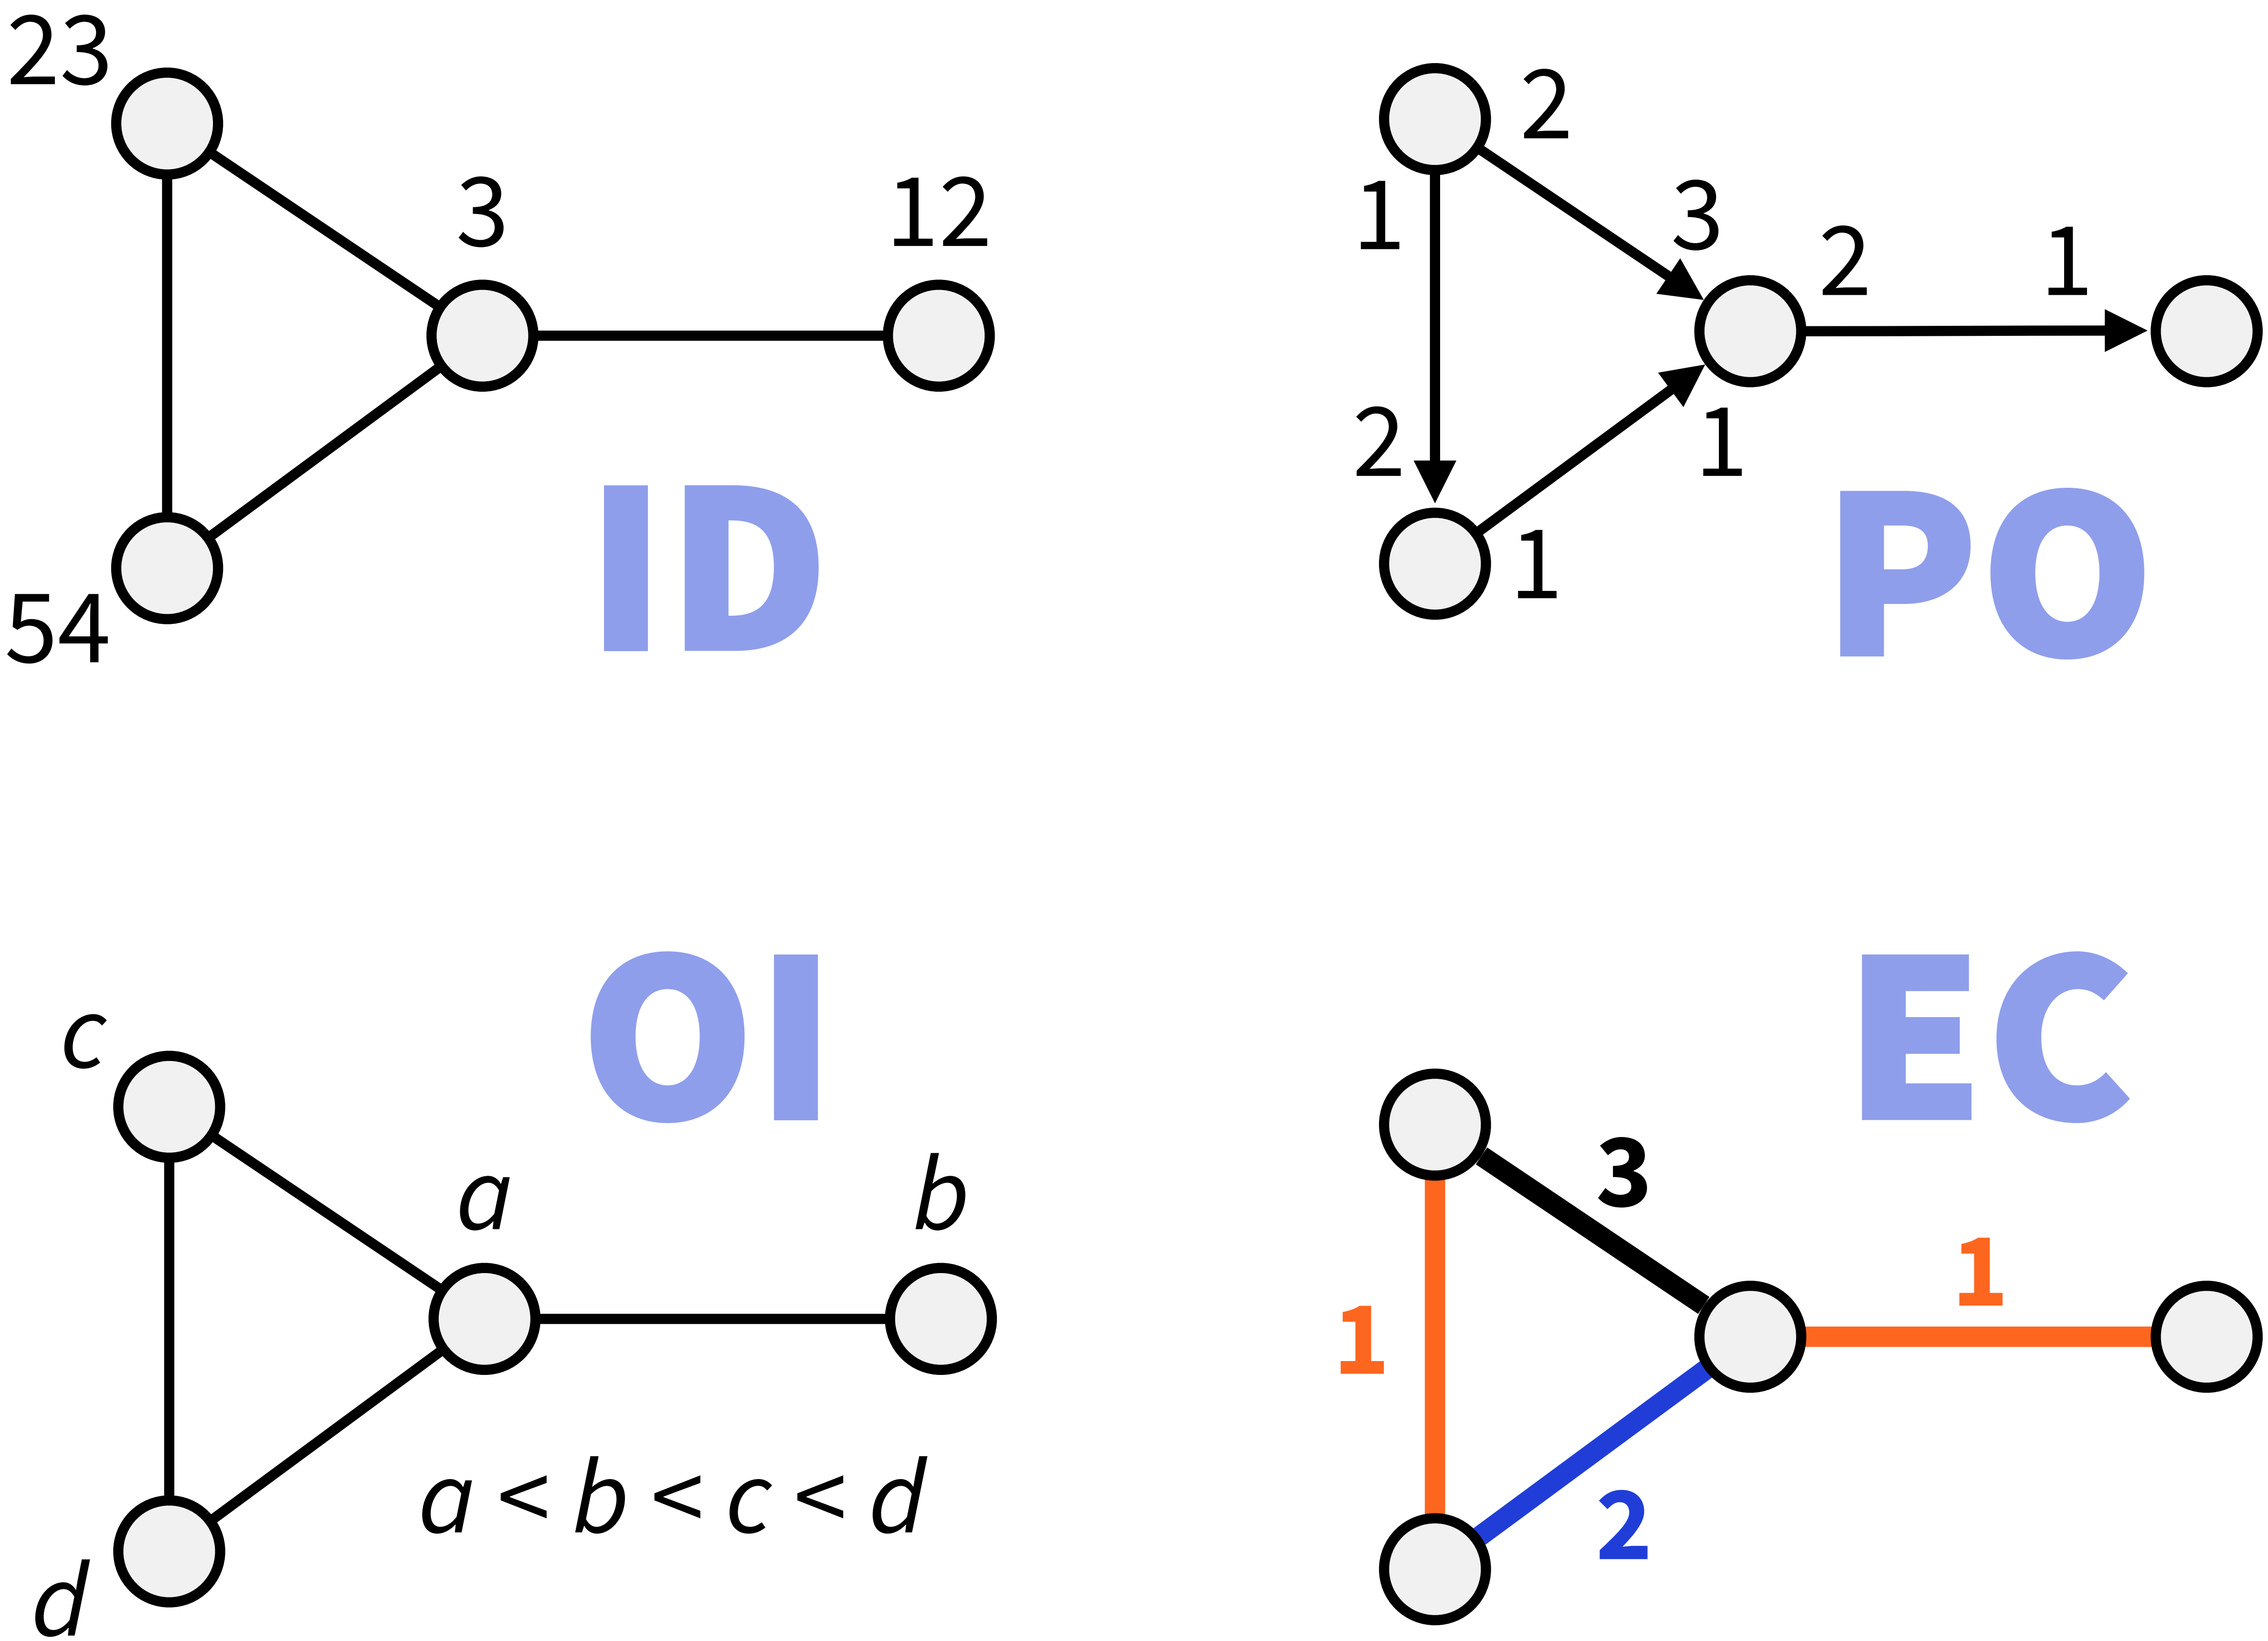
23
3
12
54
ID
2
1
3
2
1
2
1
1
PO
c
a
b
d
a < b < c < d
OI
EC
3
1
1
2

State of the art in 2014

- **Problems with $O(\Delta + \log^* n)$ algorithms:**
  - maximal matching
  - maximal independent set
  - vertex colouring with Δ+1 colours
  - edge colouring with $2\Delta-1$ colours …

State of the art in 2014

- **Problems with $O(\Delta + \log^* n)$ algorithms**
- **Problems with $O(\Delta)$ algorithms:**
  - maximal fractional matching
  - bipartite maximal matching …

State of the art in 2014

- **Problems with $O(\Delta + \log^* n)$ algorithms**
- **Problems with $O(\Delta)$ algorithms**
- **Some linear-in-Δ lower bounds:**
  - maximal matchings, EC model
  - maximal fractional matchings, LOCAL model

State of the art in 2014

- **All these problems characterised as follows:**
    - any partial solution can be **completed**
    - but completion may be **unique**
- “***Completable but tight***” **problems**
    - greedy algorithm works,
      but it may be constrained

State of the art in 2014

- **Conjecture: "*completable but tight*" problems cannot be solved in time $o(\Delta) + O(\log^* n)$**

State of the art in 2015

- **Conjecture: "*completable but tight*" problems cannot be solved in time $o(\Delta) + O(\log^* n)$**
- ***Wrong!***

State of the art in 2015

- **Barenboim (PODC 2015):**
  - *vertex colouring* with Δ+1 colours
  - can be solved in time $o(\Delta) + O(\log^* n)$

We have a separation!

- **Barenboim (PODC 2015):**
  - *edge colouring* with $2\Delta - 1$ colours
  - possible in time $o(\Delta)$ in **EC model**
- **PODC 2012:**
  - *maximal matching*
  - not possible in time $o(\Delta)$ in **EC model**

Next steps?

- **Separation for maximal independent set and (Δ+1)-vertex colouring in weak models**
- **Model:** anonymous vertex-coloured graphs
- **Lower bound:** just take line graphs
- **Upper bound:** adapt Barenboim's idea **??**

Next steps?

- **What is the new conjecture?**
- **Which problems require linear-in-Δ rounds?**
- **(Δ+1)-colouring:** ***not***
- **Greedy colouring:** ***perhaps??***
  - lower bounds: e.g. Gavoille et al. (2009)

Next steps?

- **Linear-in-Δ lower bound for bipartite maximal matching**
- **Good:** pure local coordination, no symmetry-breaking needed
- **Needed:** extend known techniques so that they tolerate 2-coloured inputs

Next steps?

- **Poorly understood:** **optimisation problems**
- **Example:** minimum vertex cover (VC) vs. maximal fractional matchings (MFM)
- **Good:** MFM → 2-approximation of VC
- **Needed:** 2-approximation of VC → MFM **???**

Next steps?

- **Reductions, conditional lower bounds!**
  - hardness, completeness?
- **Problems that are *at least as hard as* bipartite maximal matching**
- **Problems that are *at most as hard as* bipartite maximal matching**

Summary

- **Distributed time complexity, LOCAL model**
- **$O(\log^* n)$: "symmetry breaking", OK**
- **$O(\Delta)$: "local coordination", poorly understood**
- **Next step: *bipartite* maximal matching**