Locality

in distributed, dynamic, online, and quantum settings

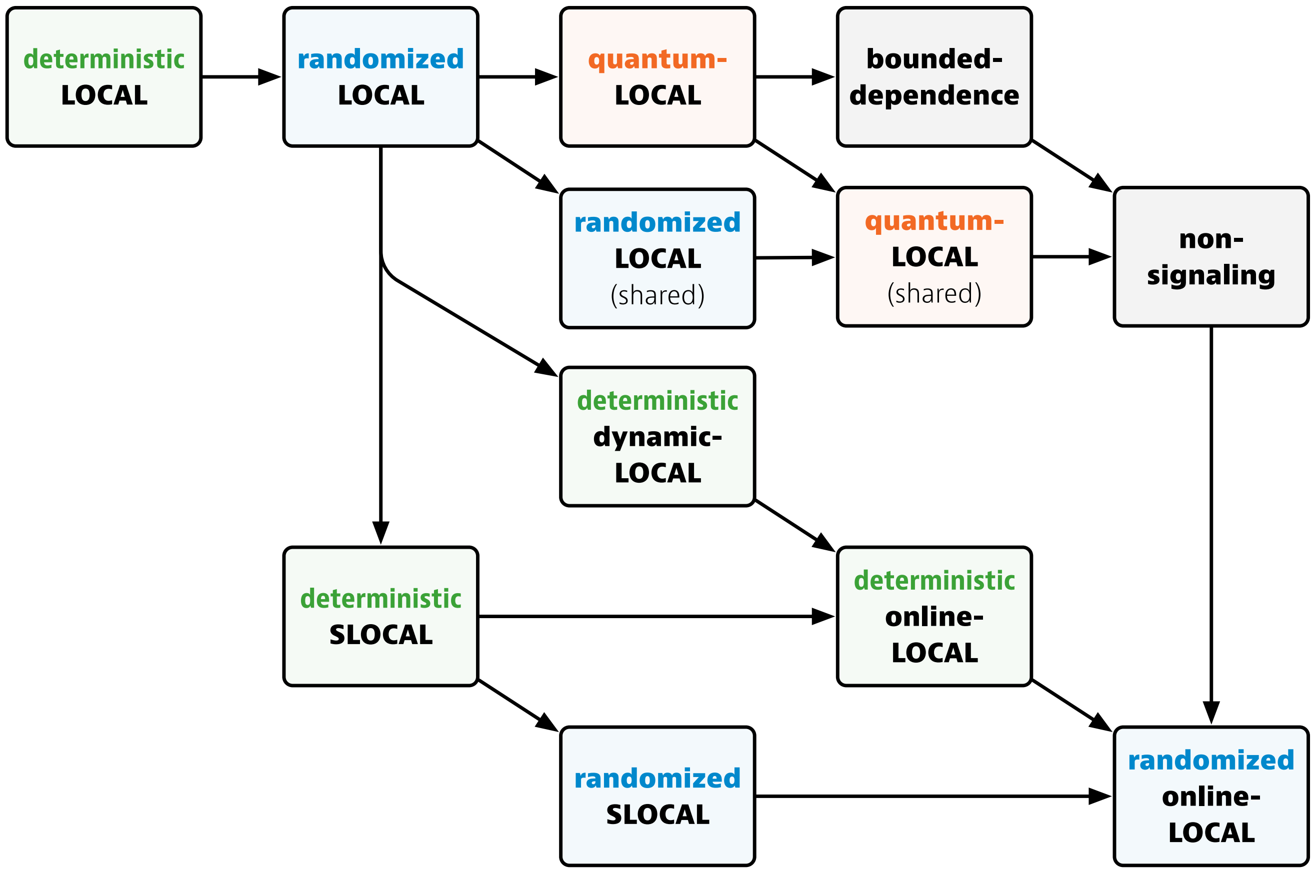

Jukka Suomela
Aalto University

Joint work with e.g.

Amirreza Akbari
Alkida Balliu
Sebastian Brandt
Xavier Coiteux-Roy
Francesco d'Amore
Anubhav Dhar
Massimo Equi
Navid Eslami
Rishikesh Gajjala
Mohsen Ghaffari
Fabian Kuhn
Eli Kujawa
François Le Gall
Henrik Lievonen
Augusto Modanese
Darya Melnyk
Mikail Muftuoglu
Dennis Olivetti
Shreyas Pai
Mikaël Rabie
Marc-Olivier Renou
Václav Rozhoň
Joona Särkijärvi
Gustav Schmid
Jan Studený
Lucas Tendick
Jara Uitto
Isadora Veeren

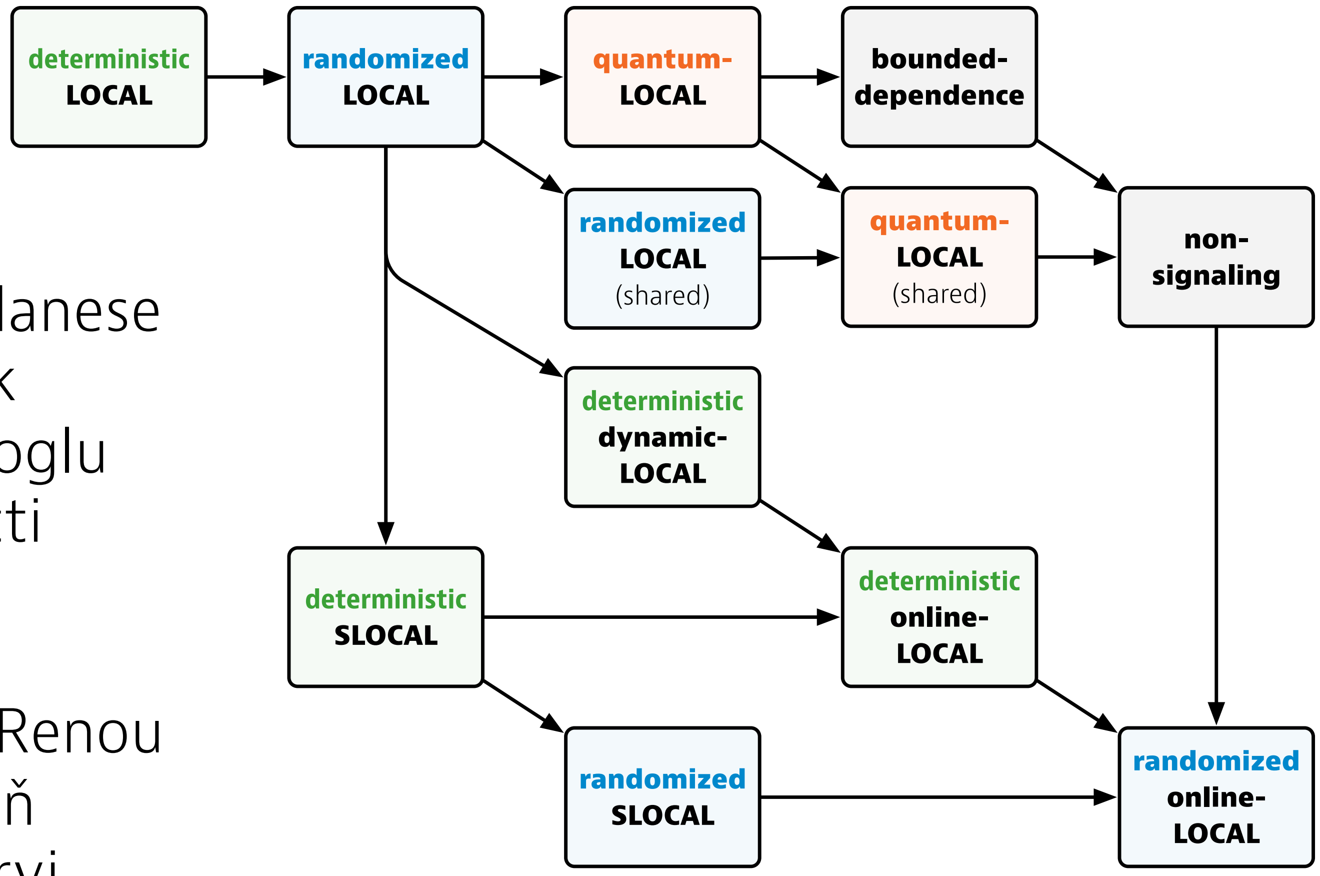

+ many others

Goals

1. Informal overview of all these models
2. Why is this relevant and interesting?
3. What is known, what is open?

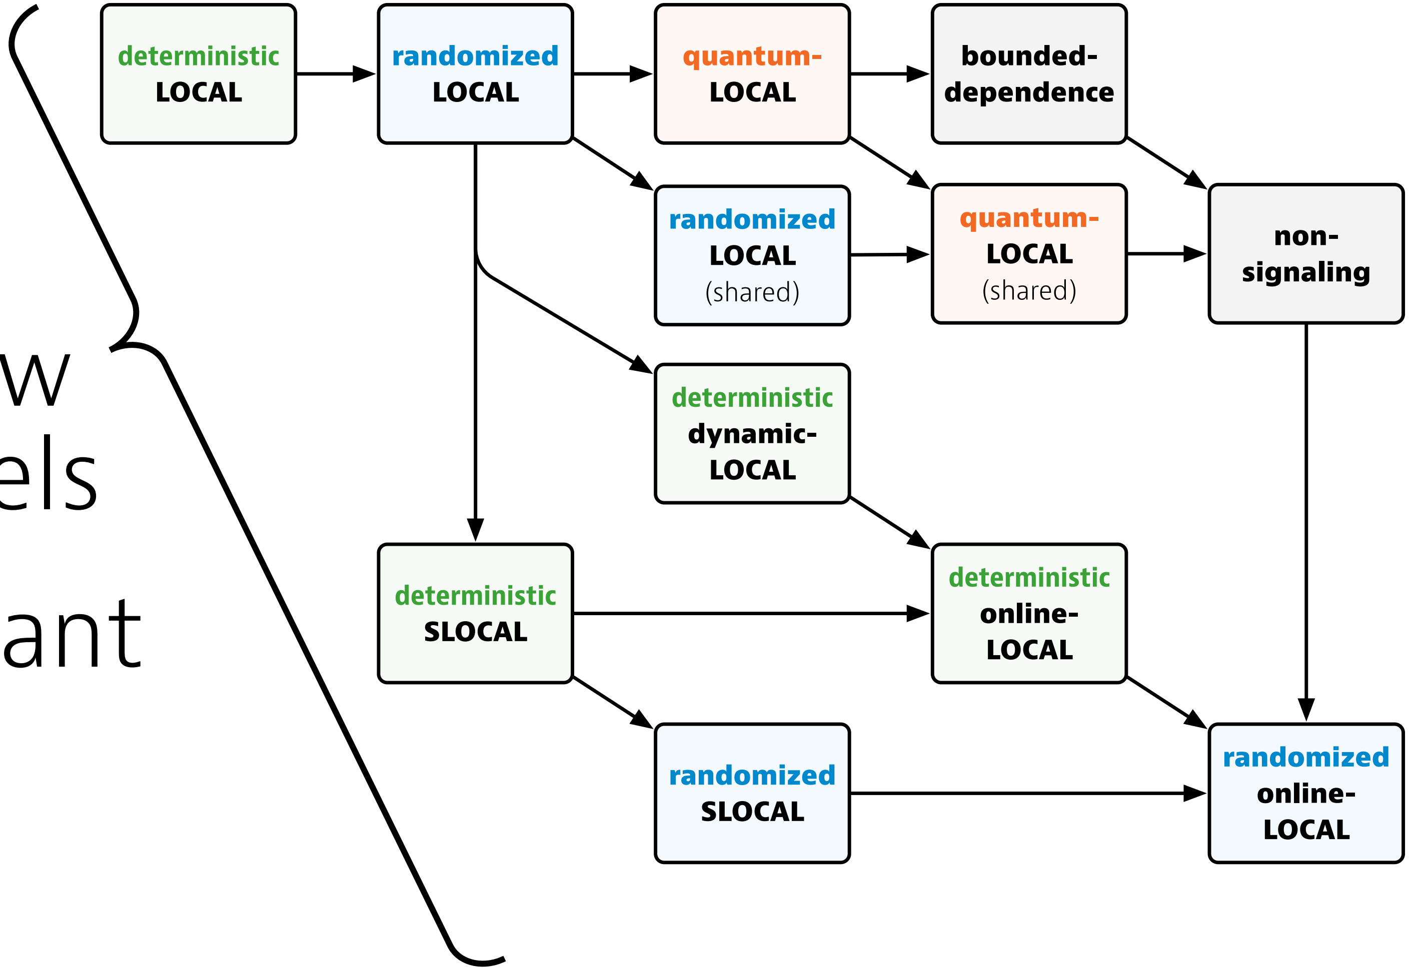

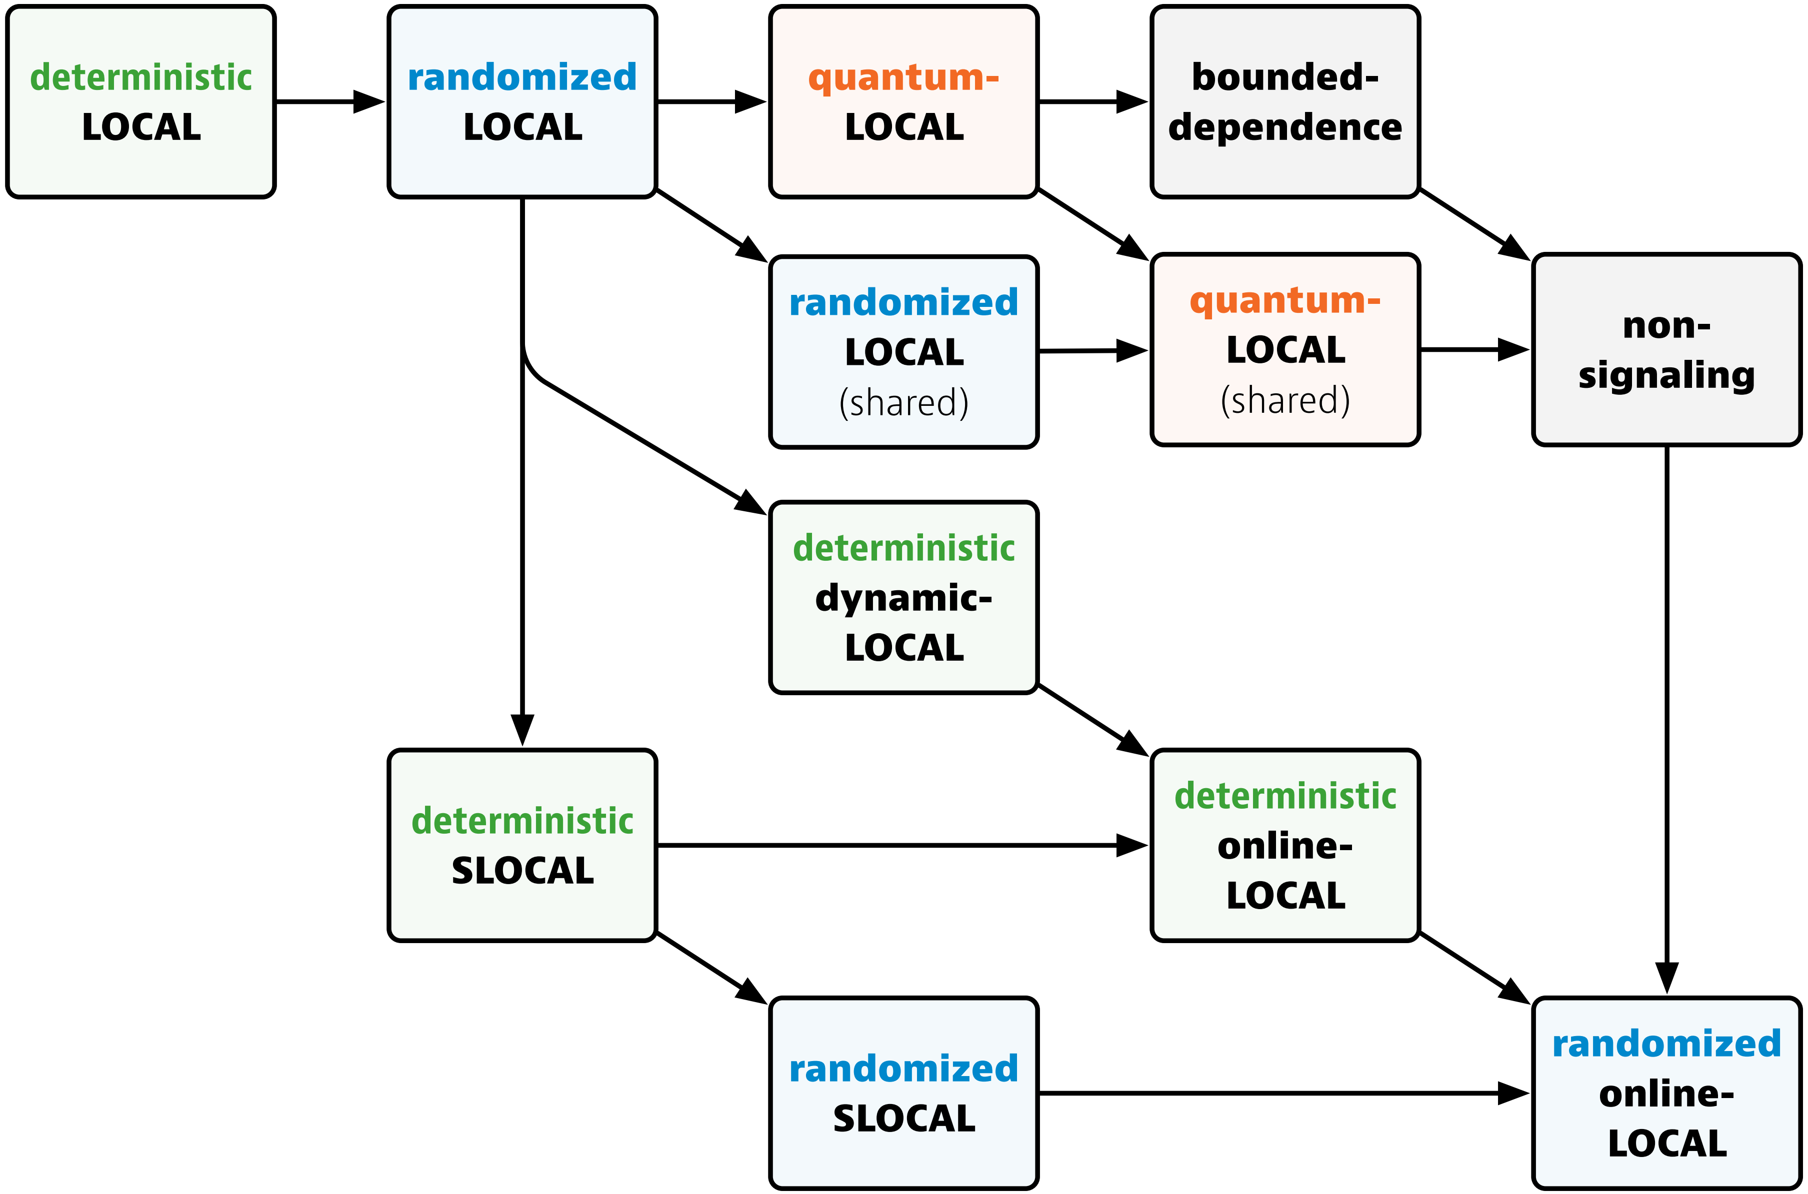

deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

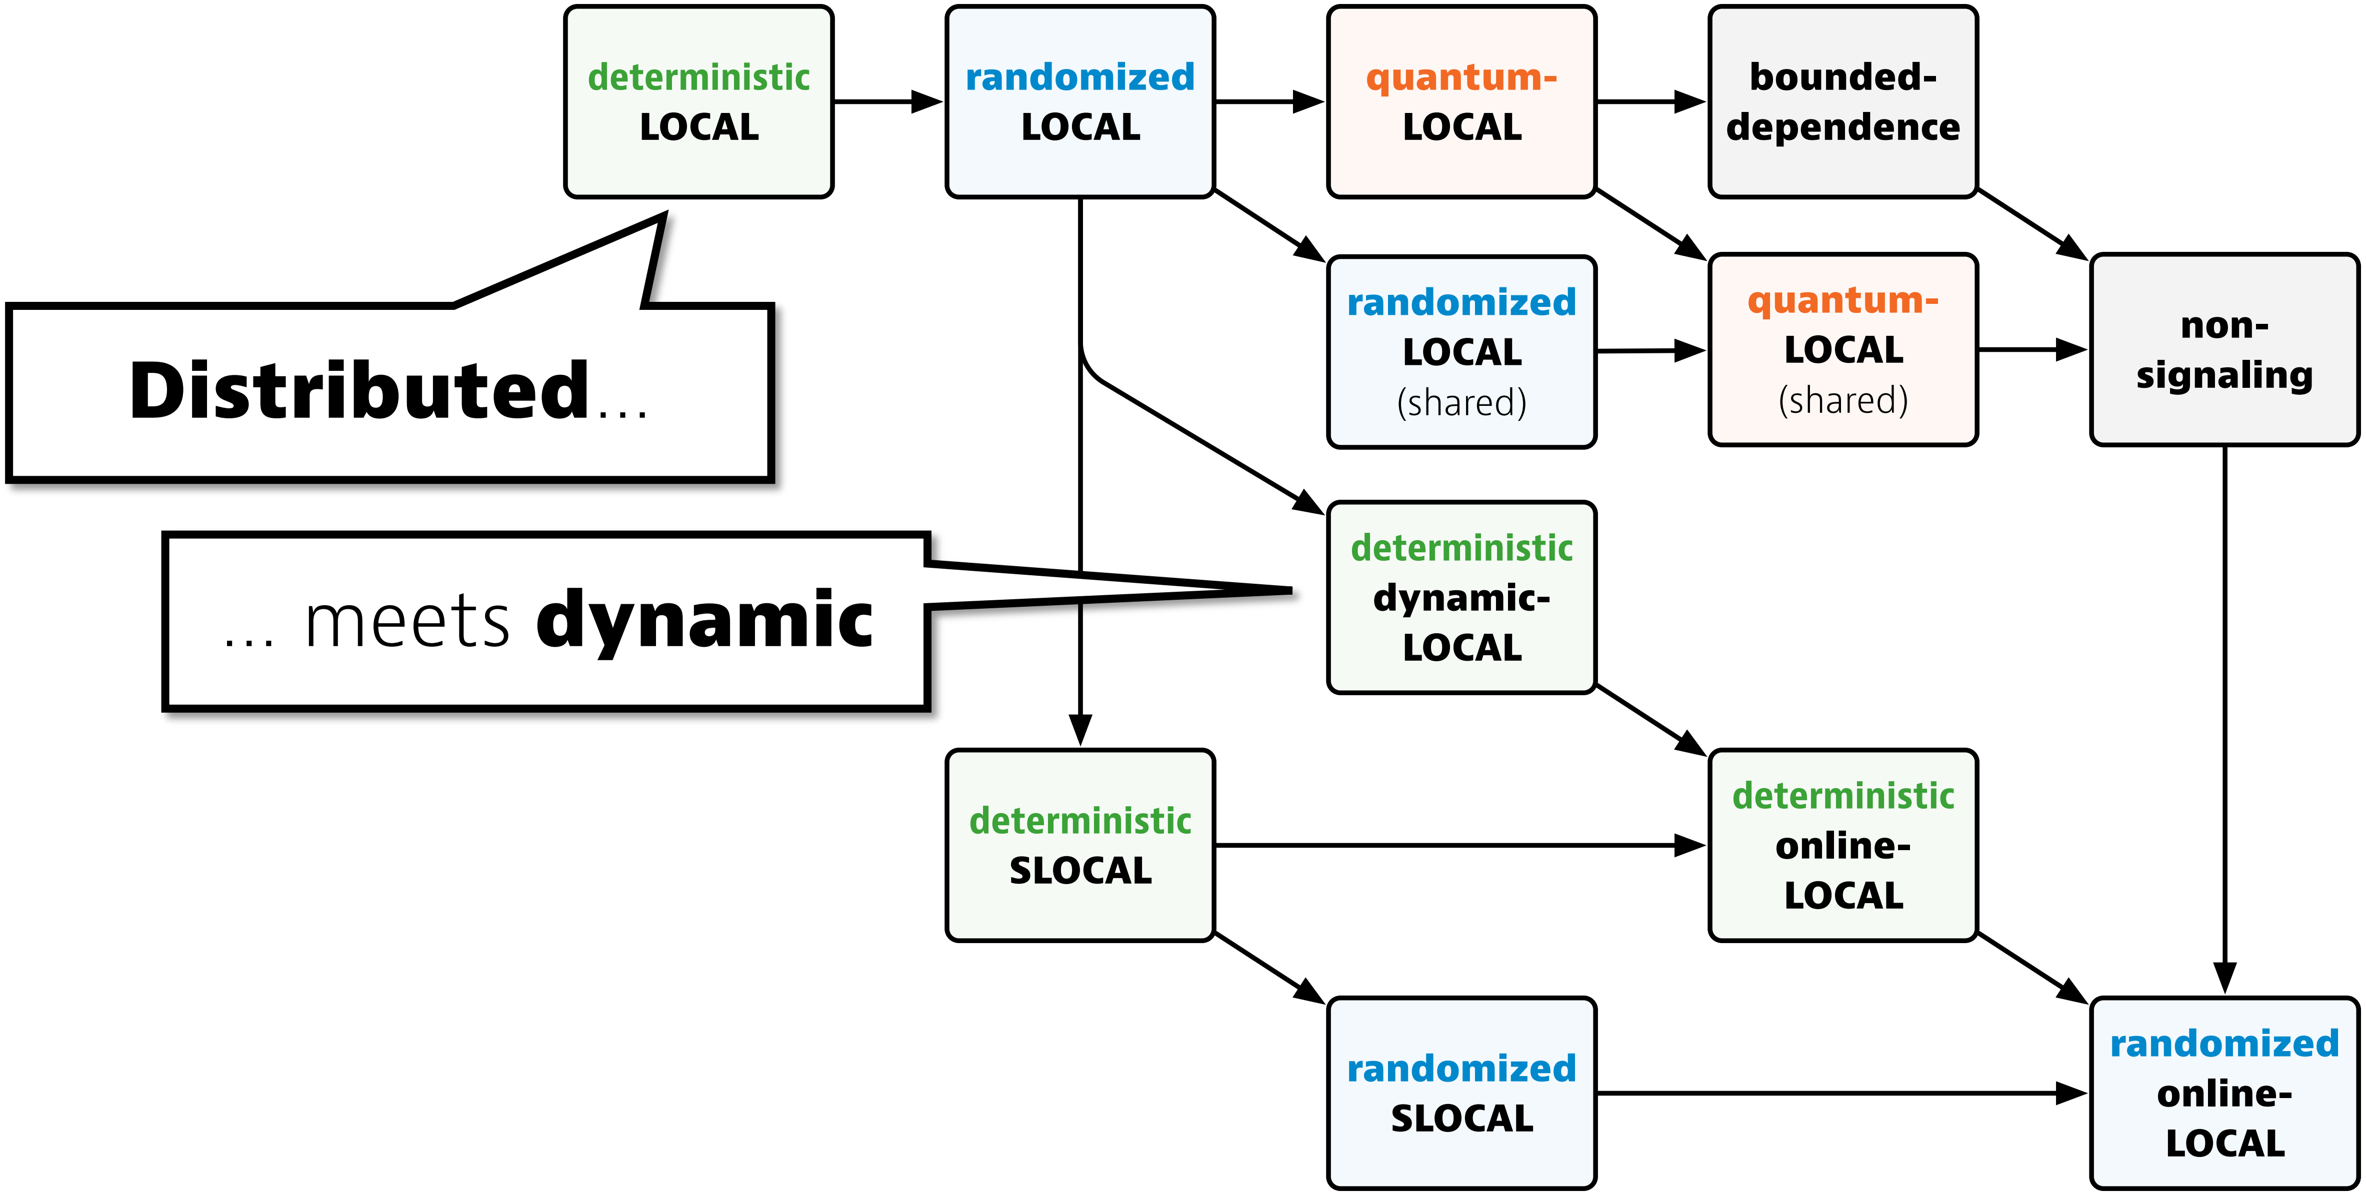

deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
Distributed…
… meets dynamic
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

Context:
locally checkable labelings (LCL)

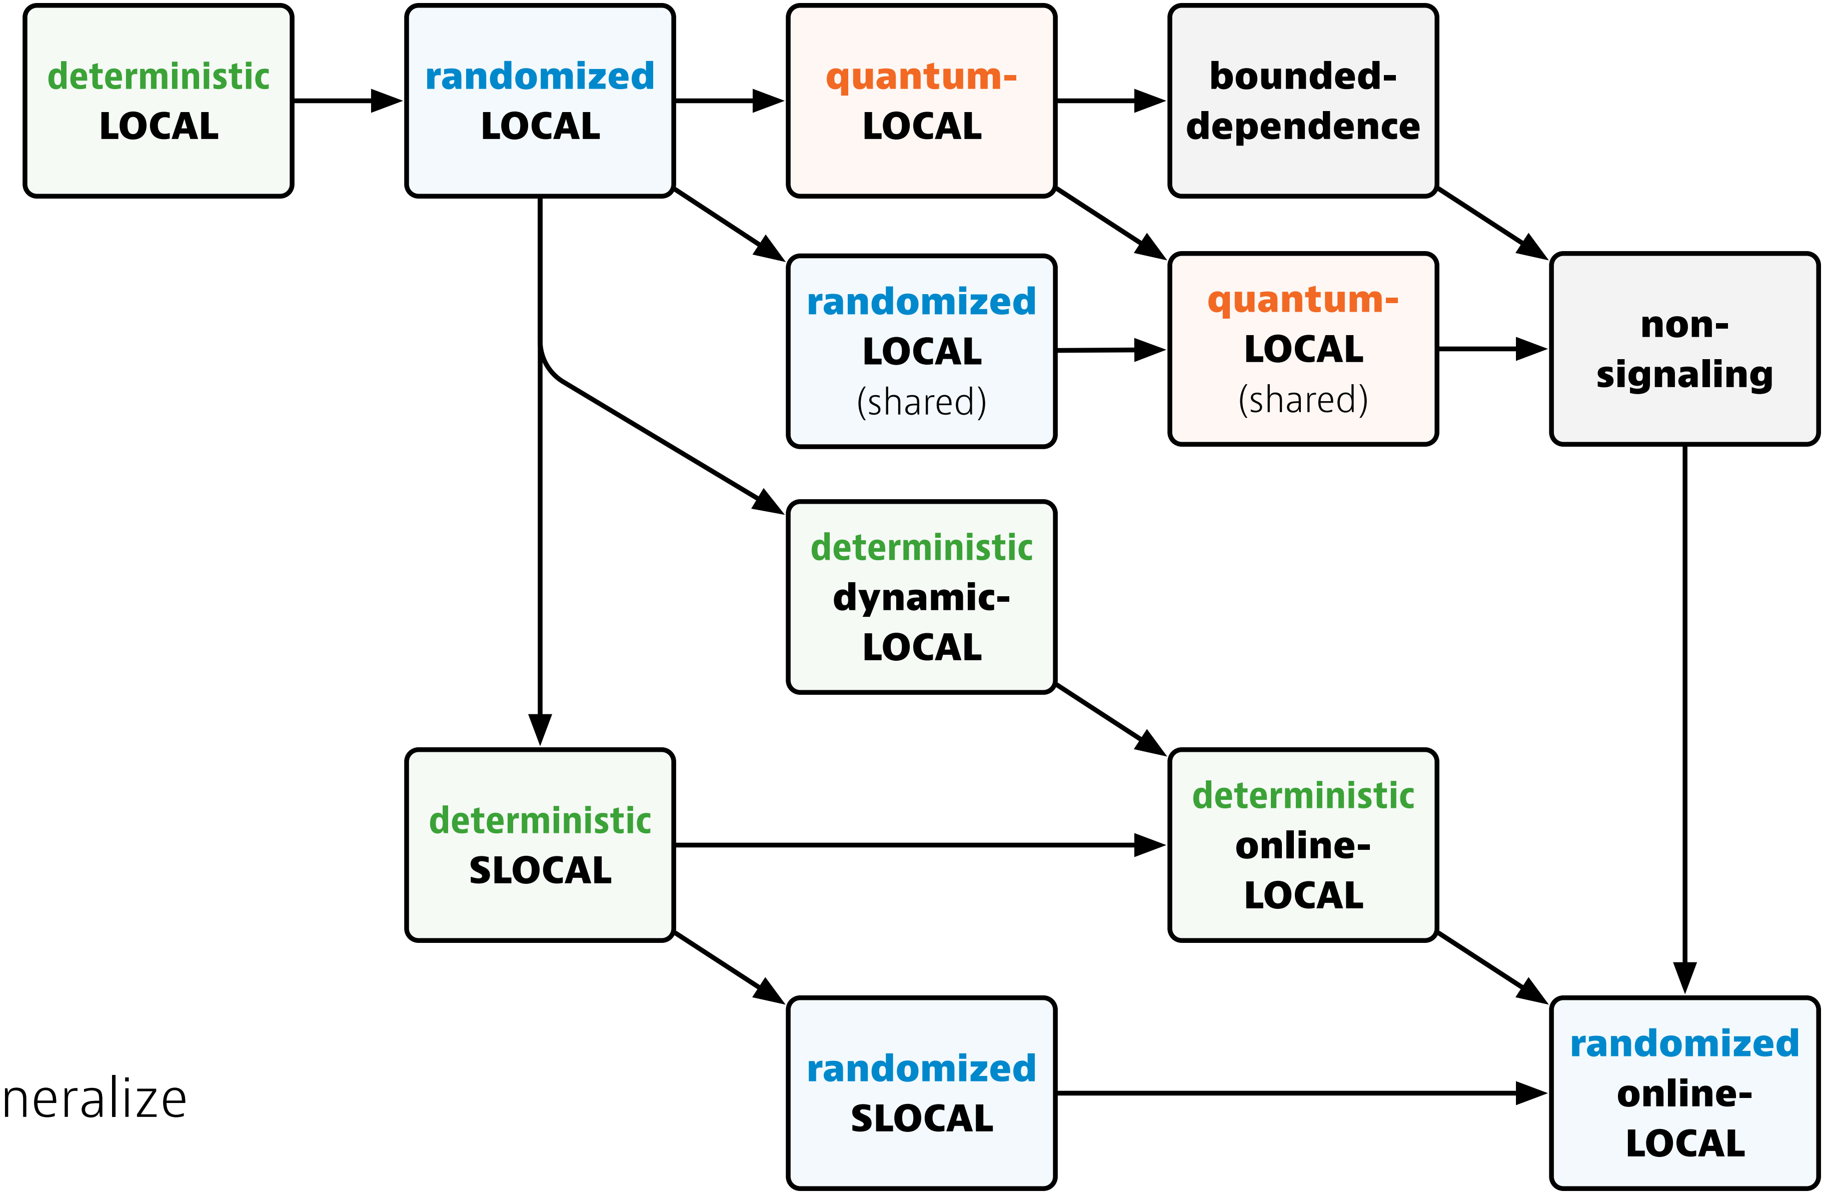

but many things generalize far beyond LCLs

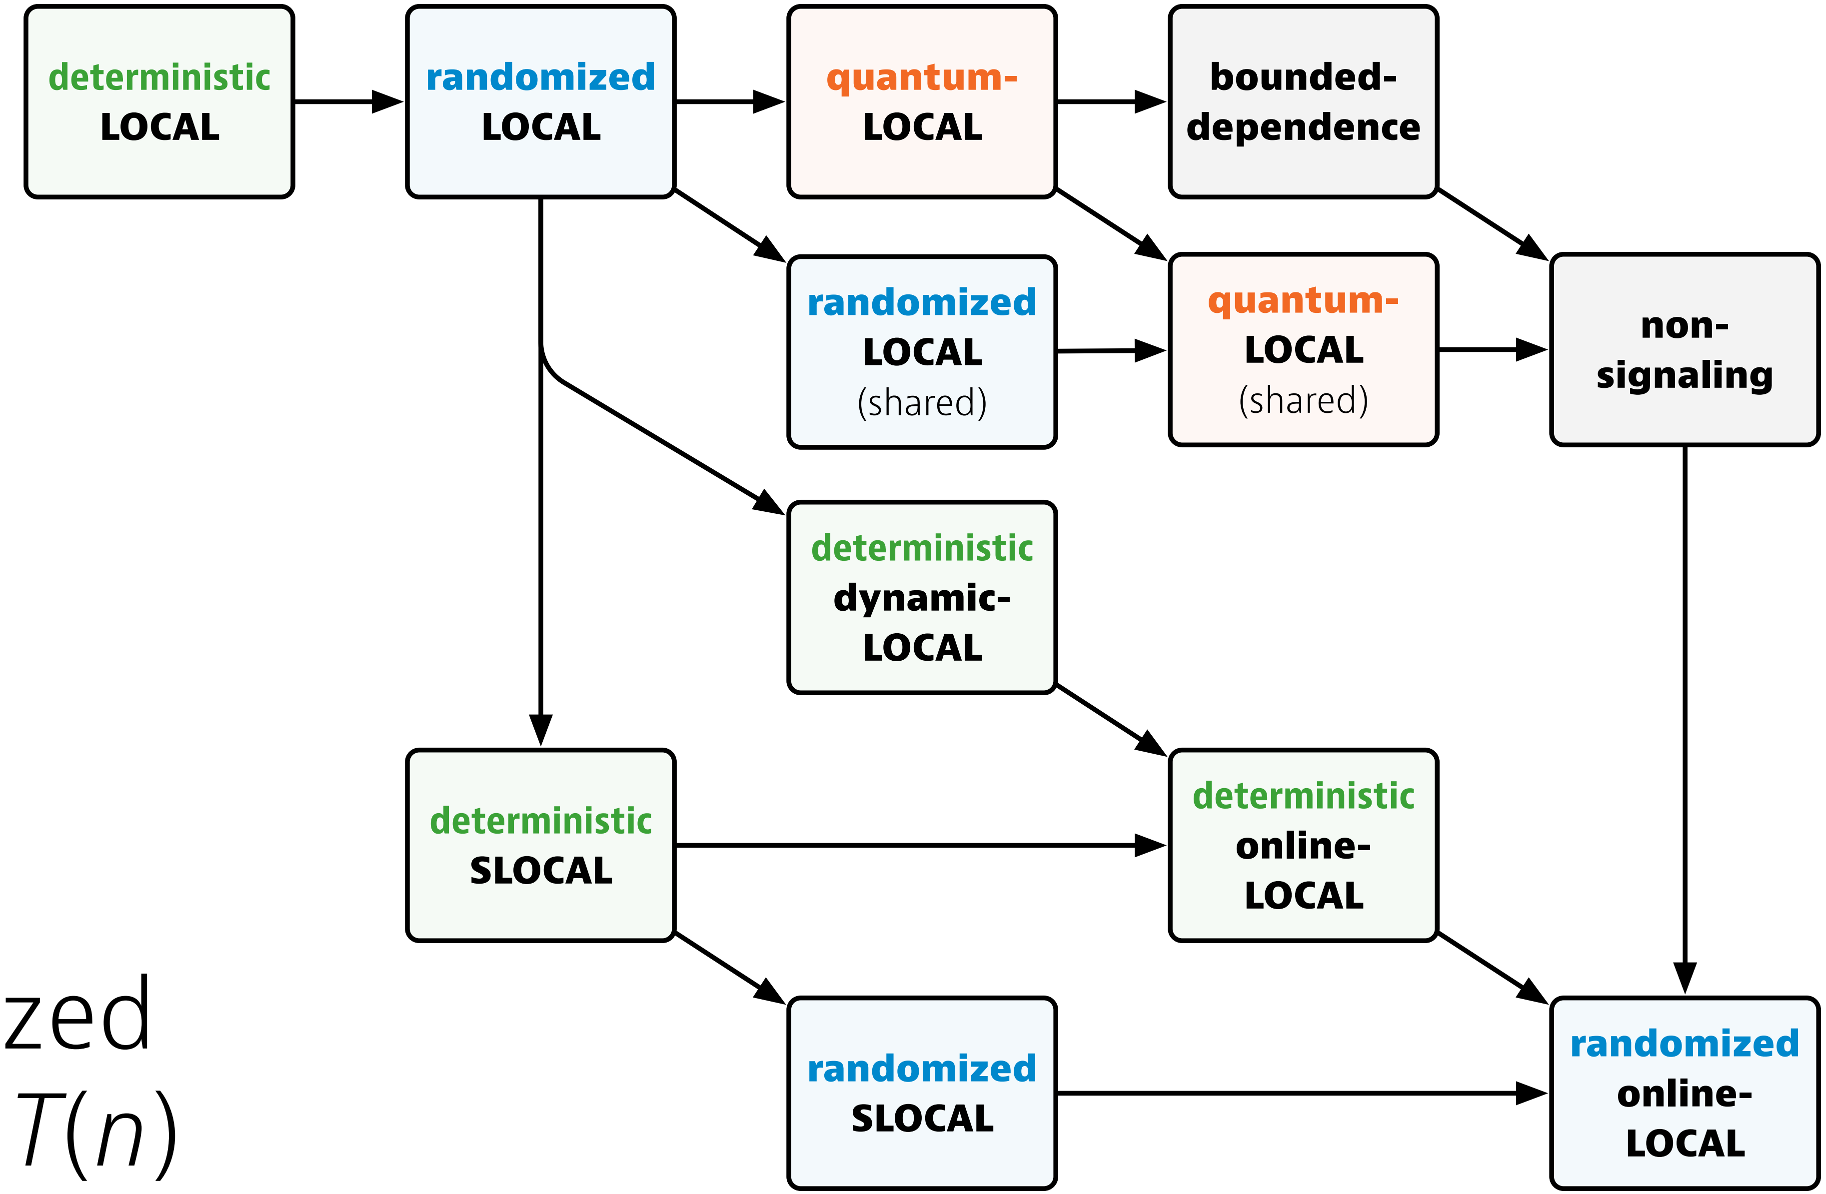

Each box:
model of computing parameterized by locality $T(n)$

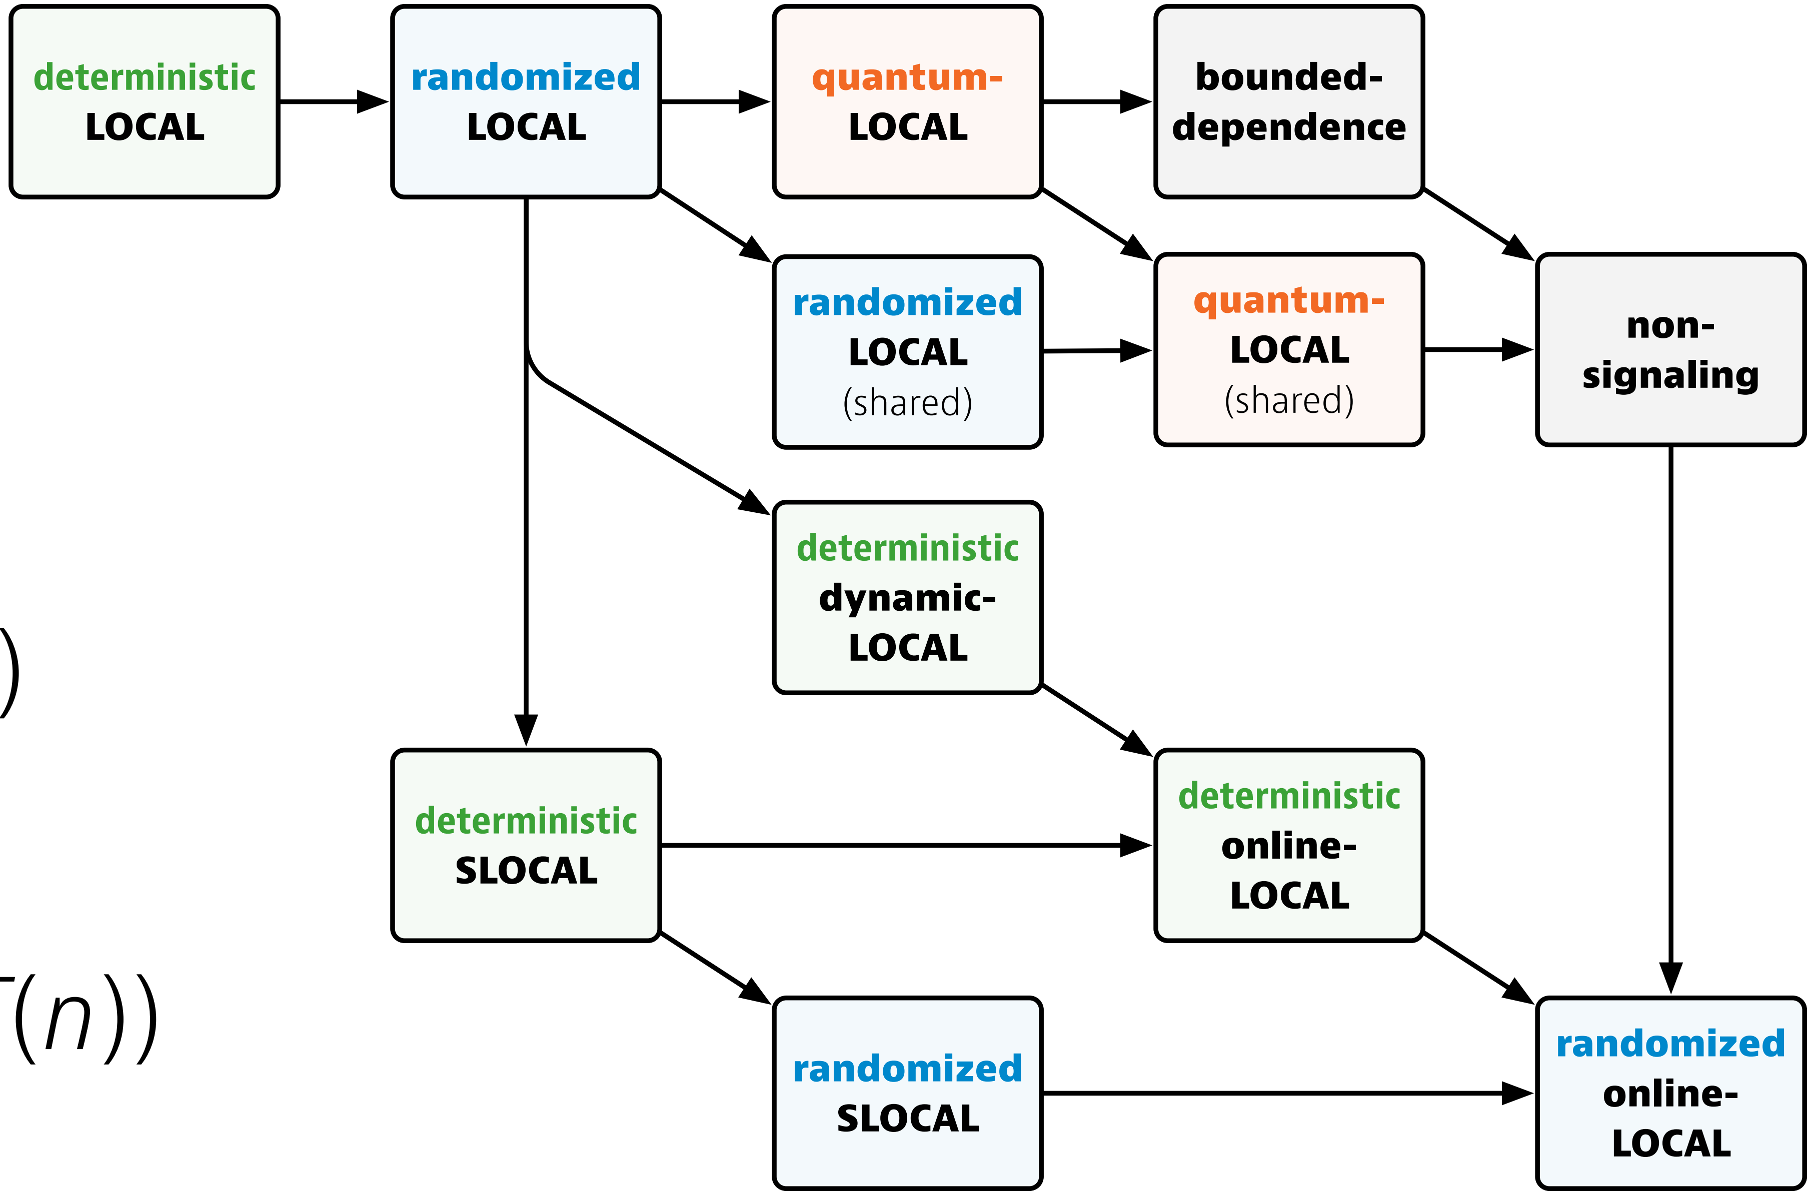

A* → *B:
locality $T(n)$
in model A
implies
locality $O(T(n))$
in model B

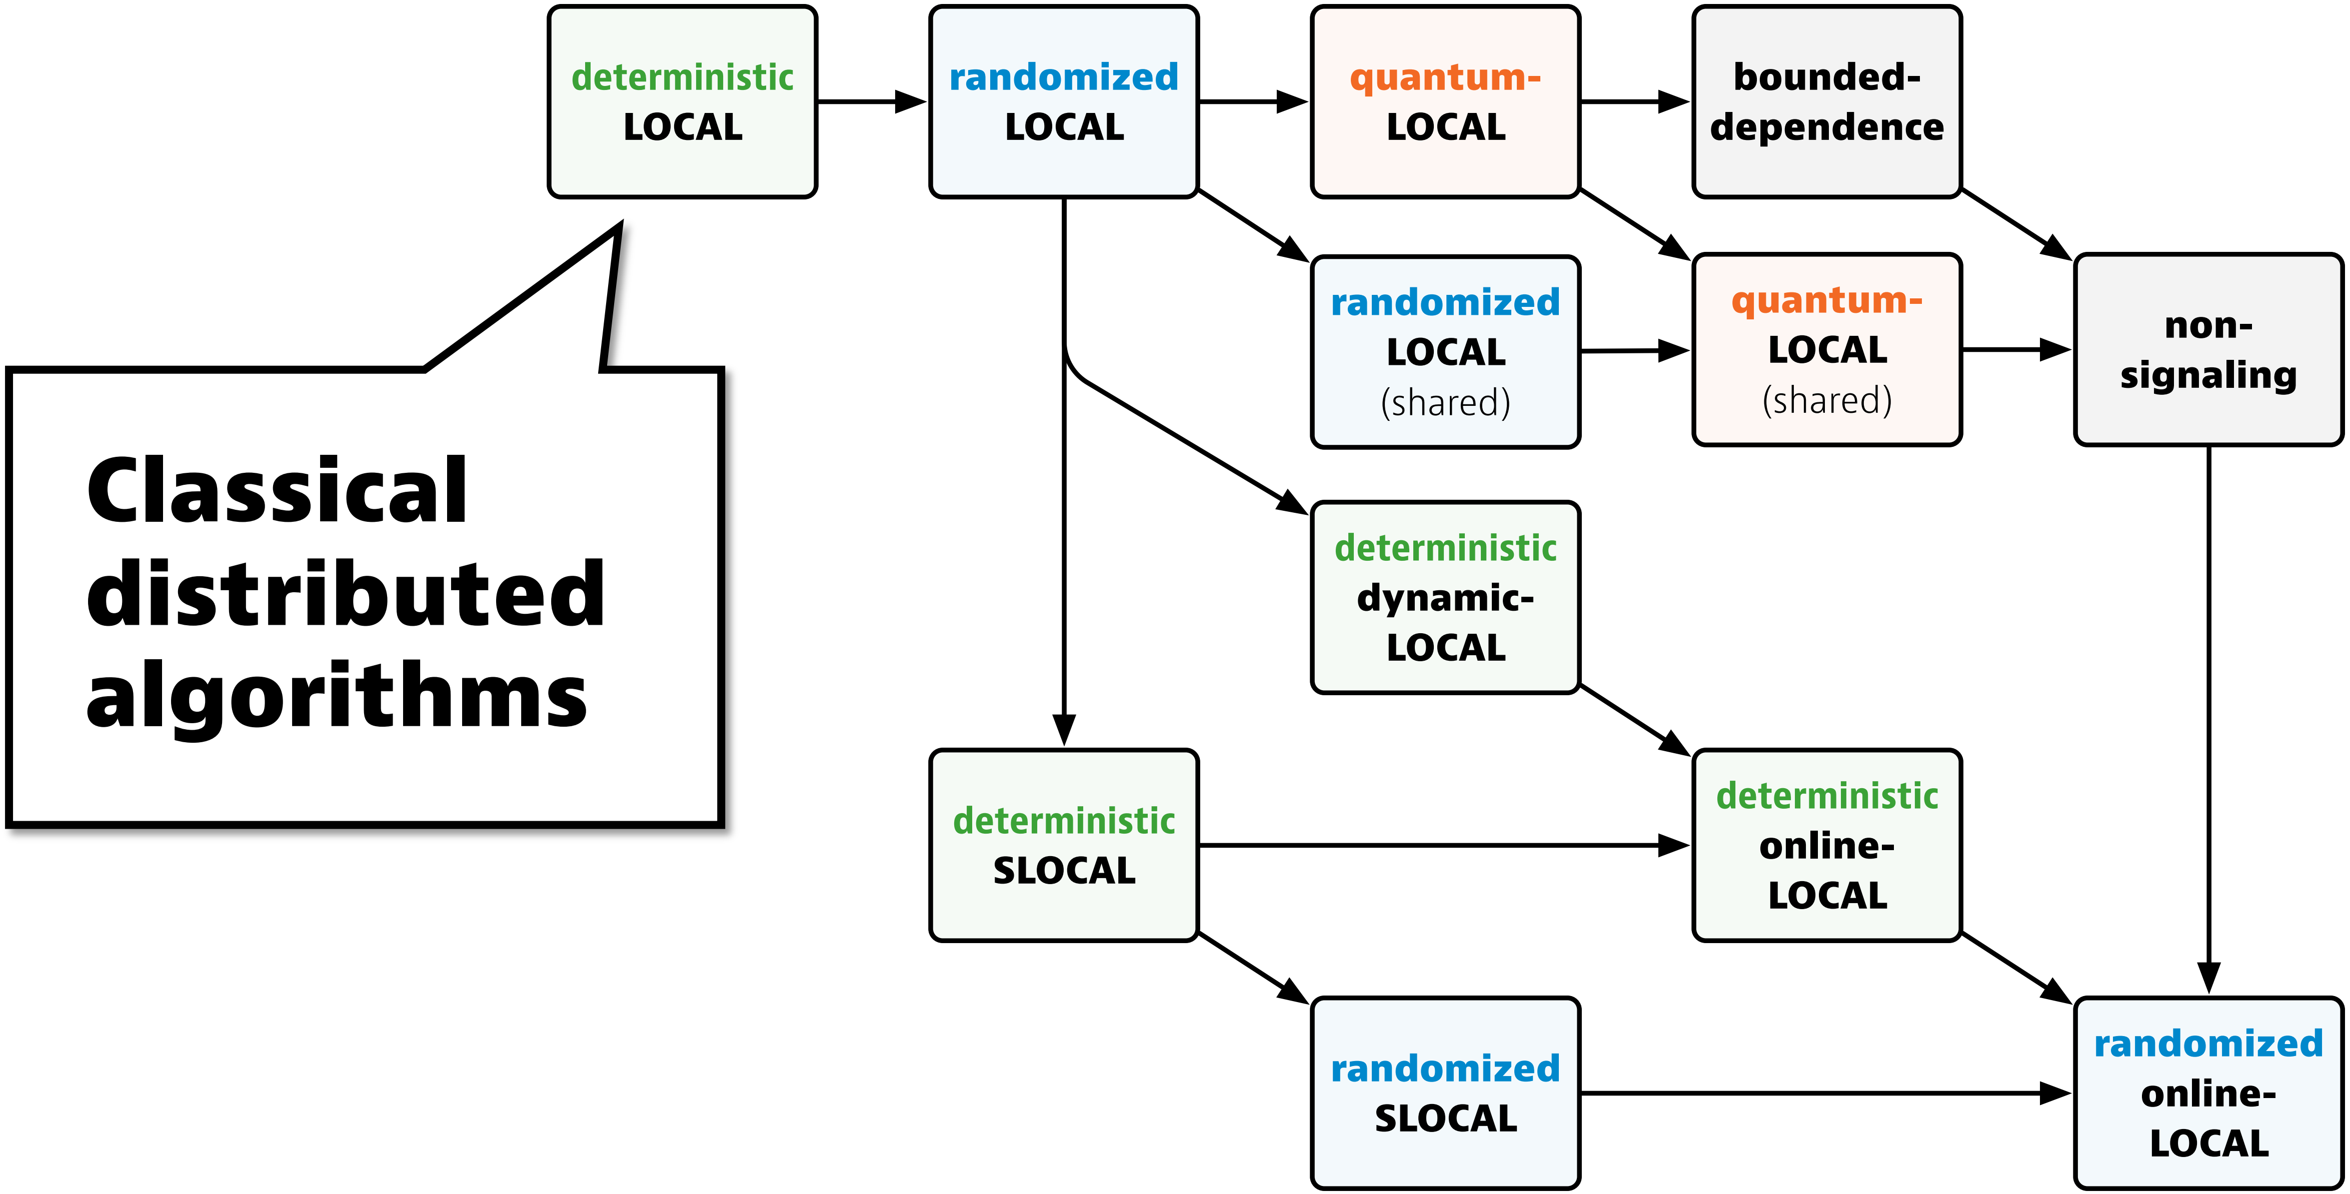
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
Classical distributed algorithms

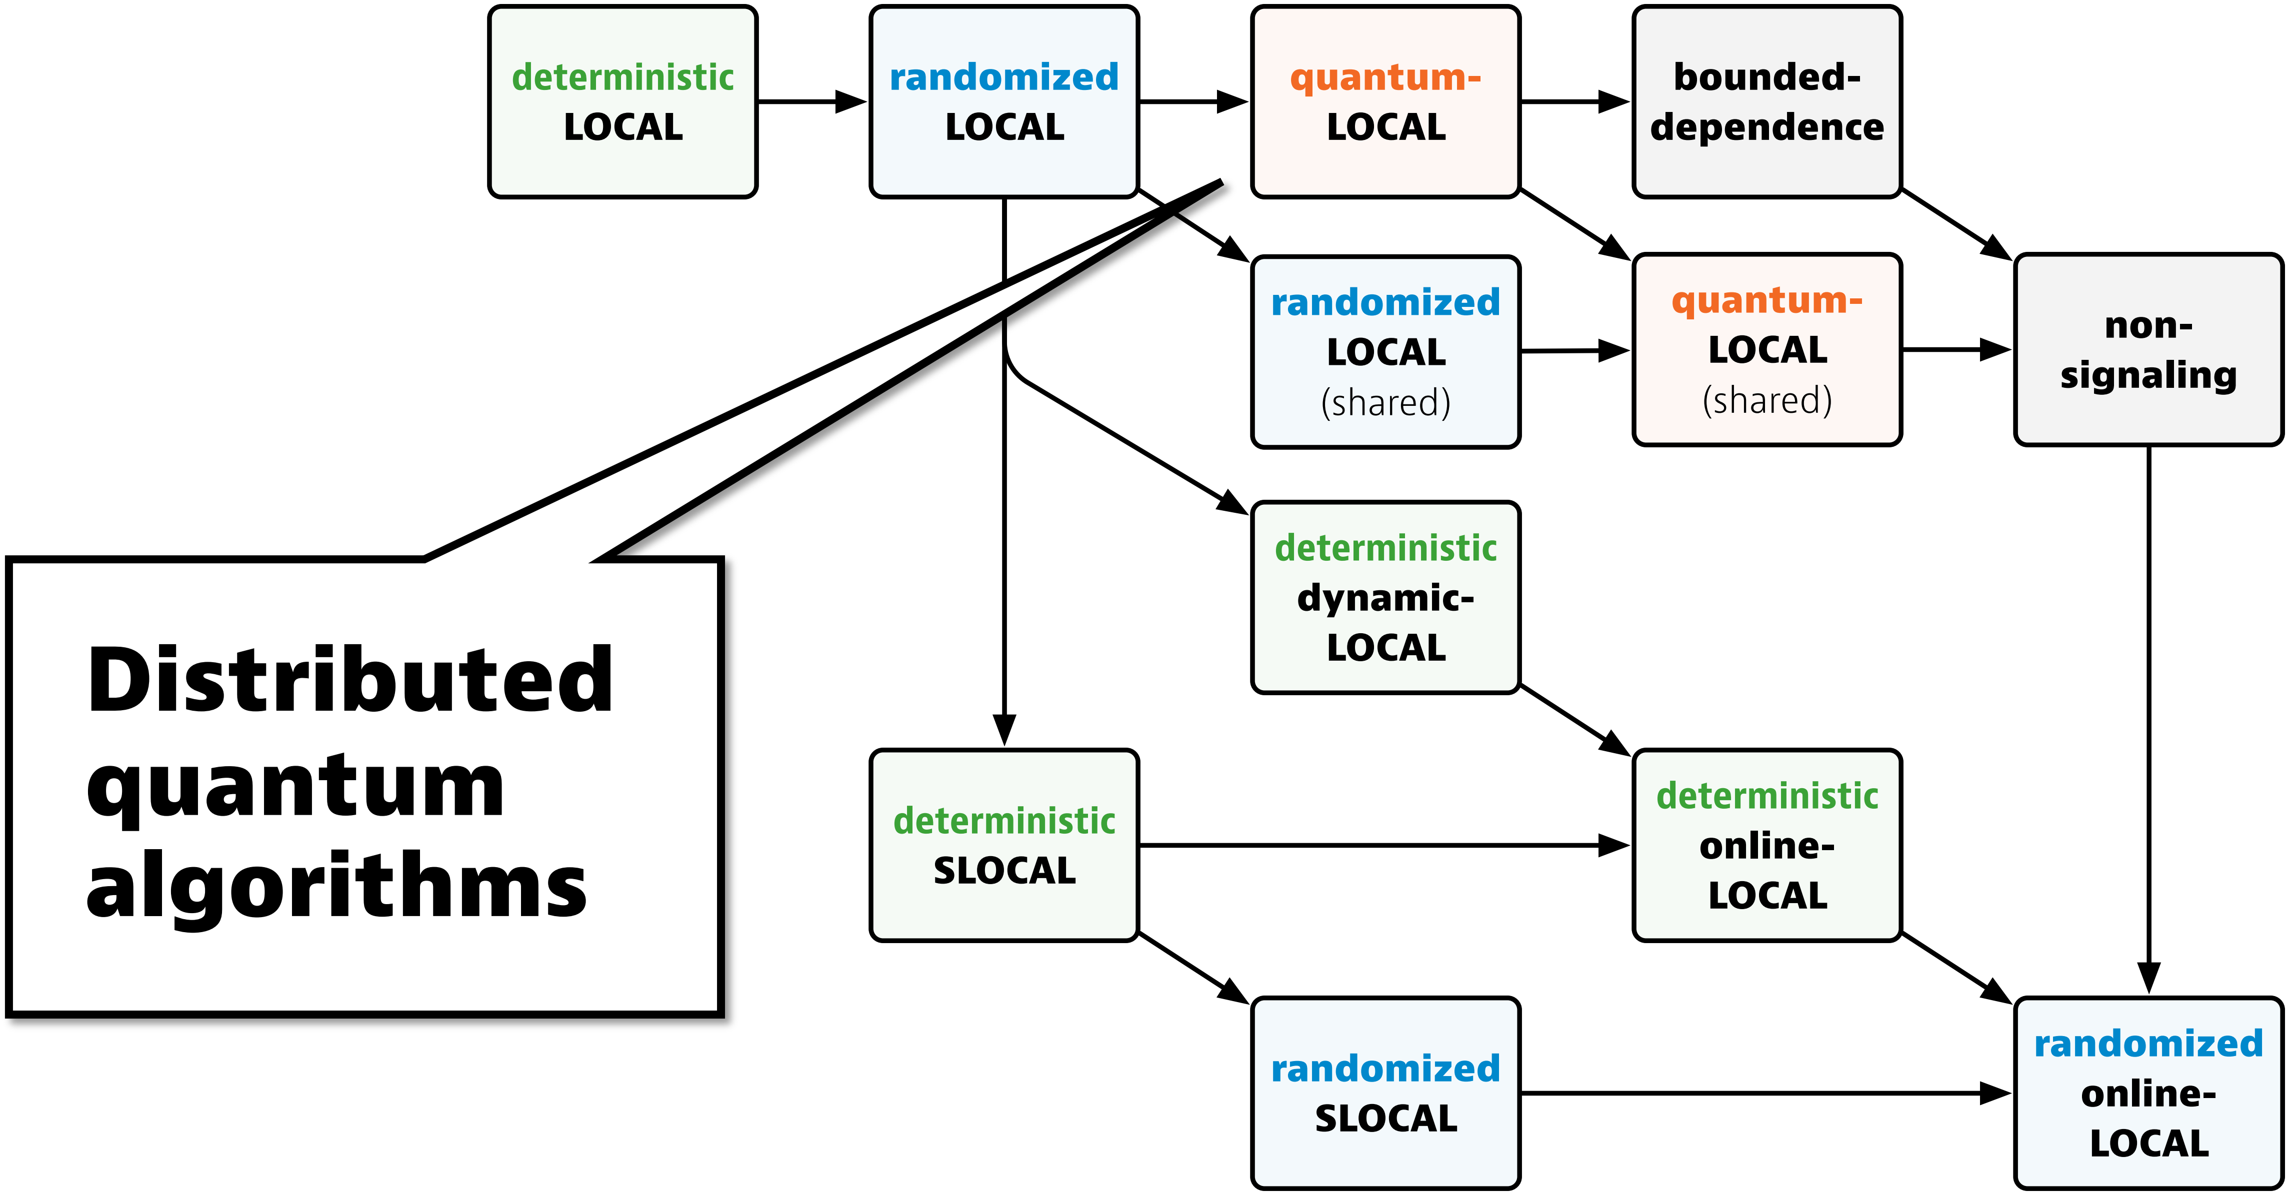
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
Distributed quantum algorithms

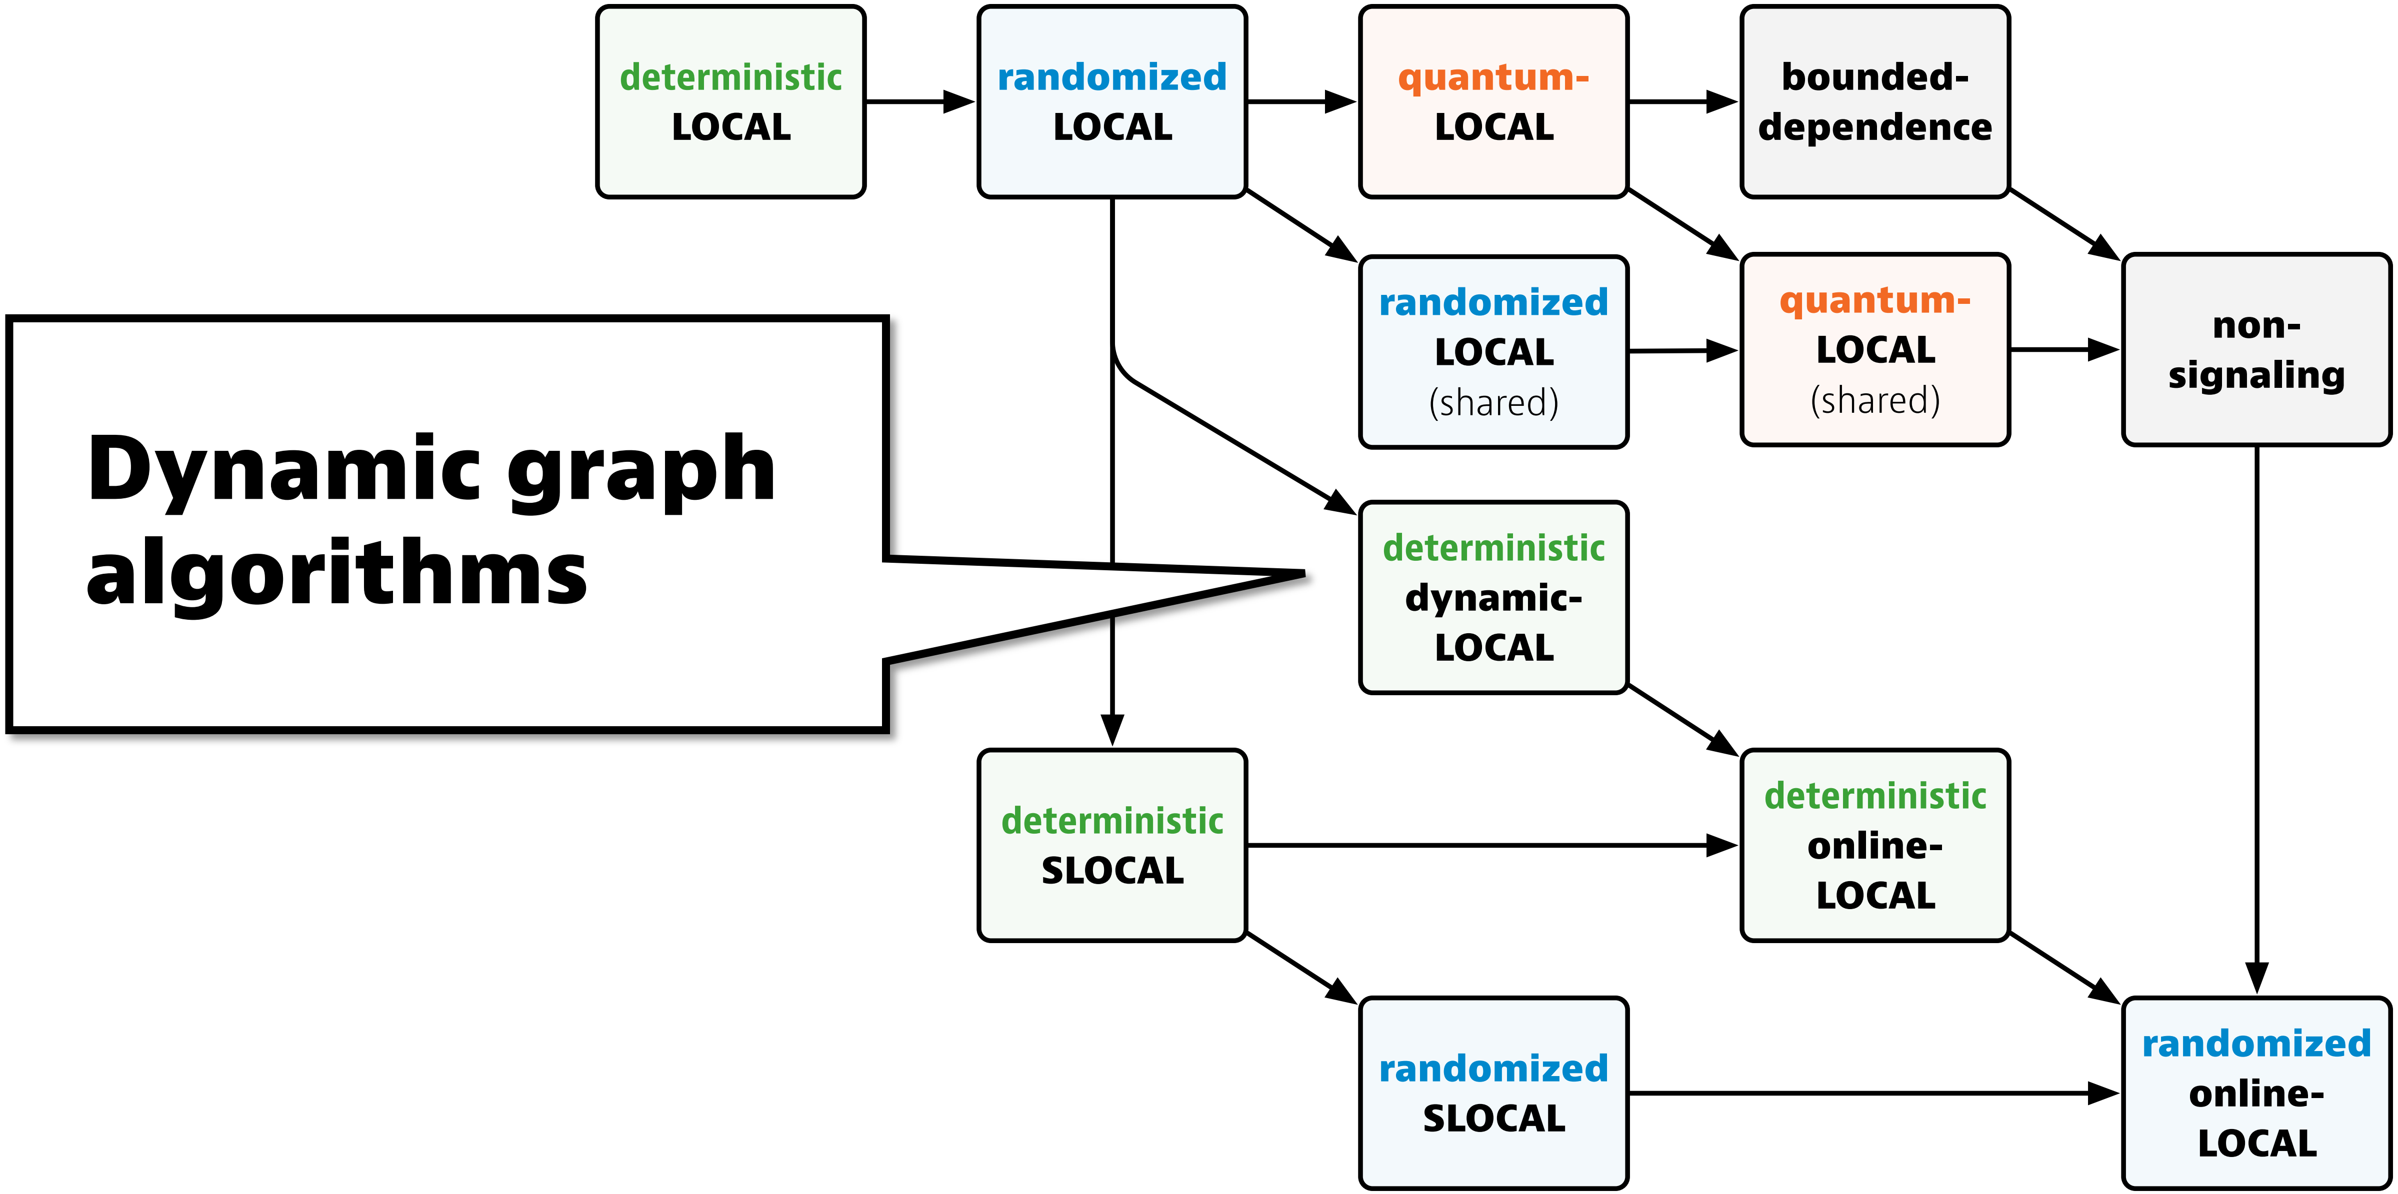
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
Dynamic graph algorithms
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

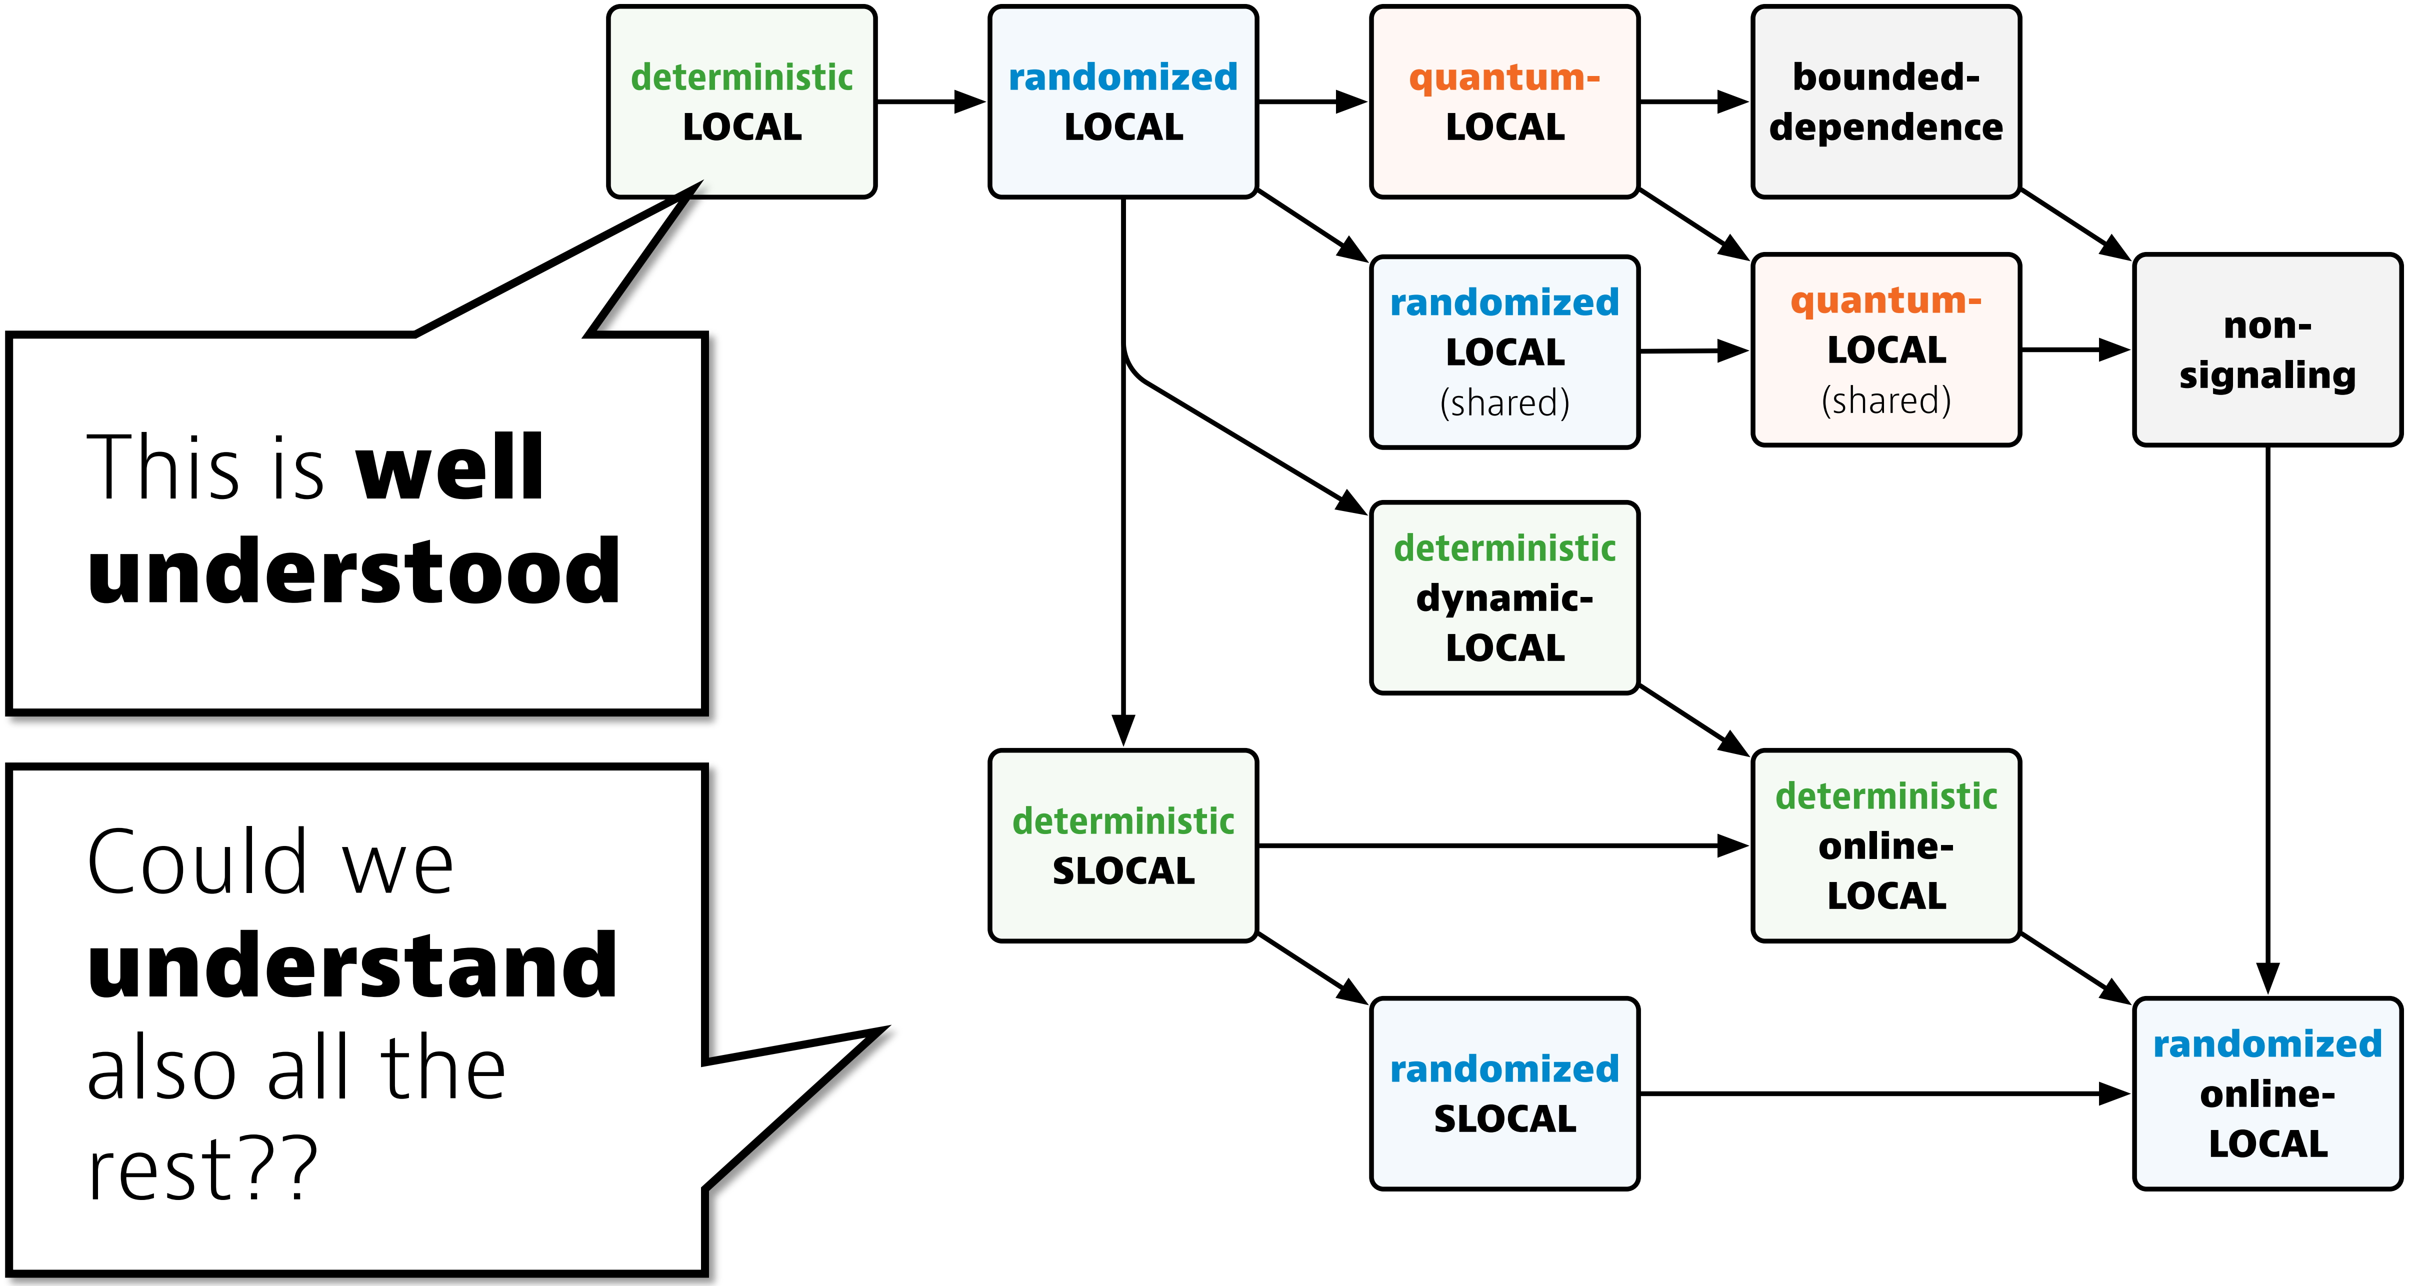

deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
This is well understood
Could we understand also all the rest??

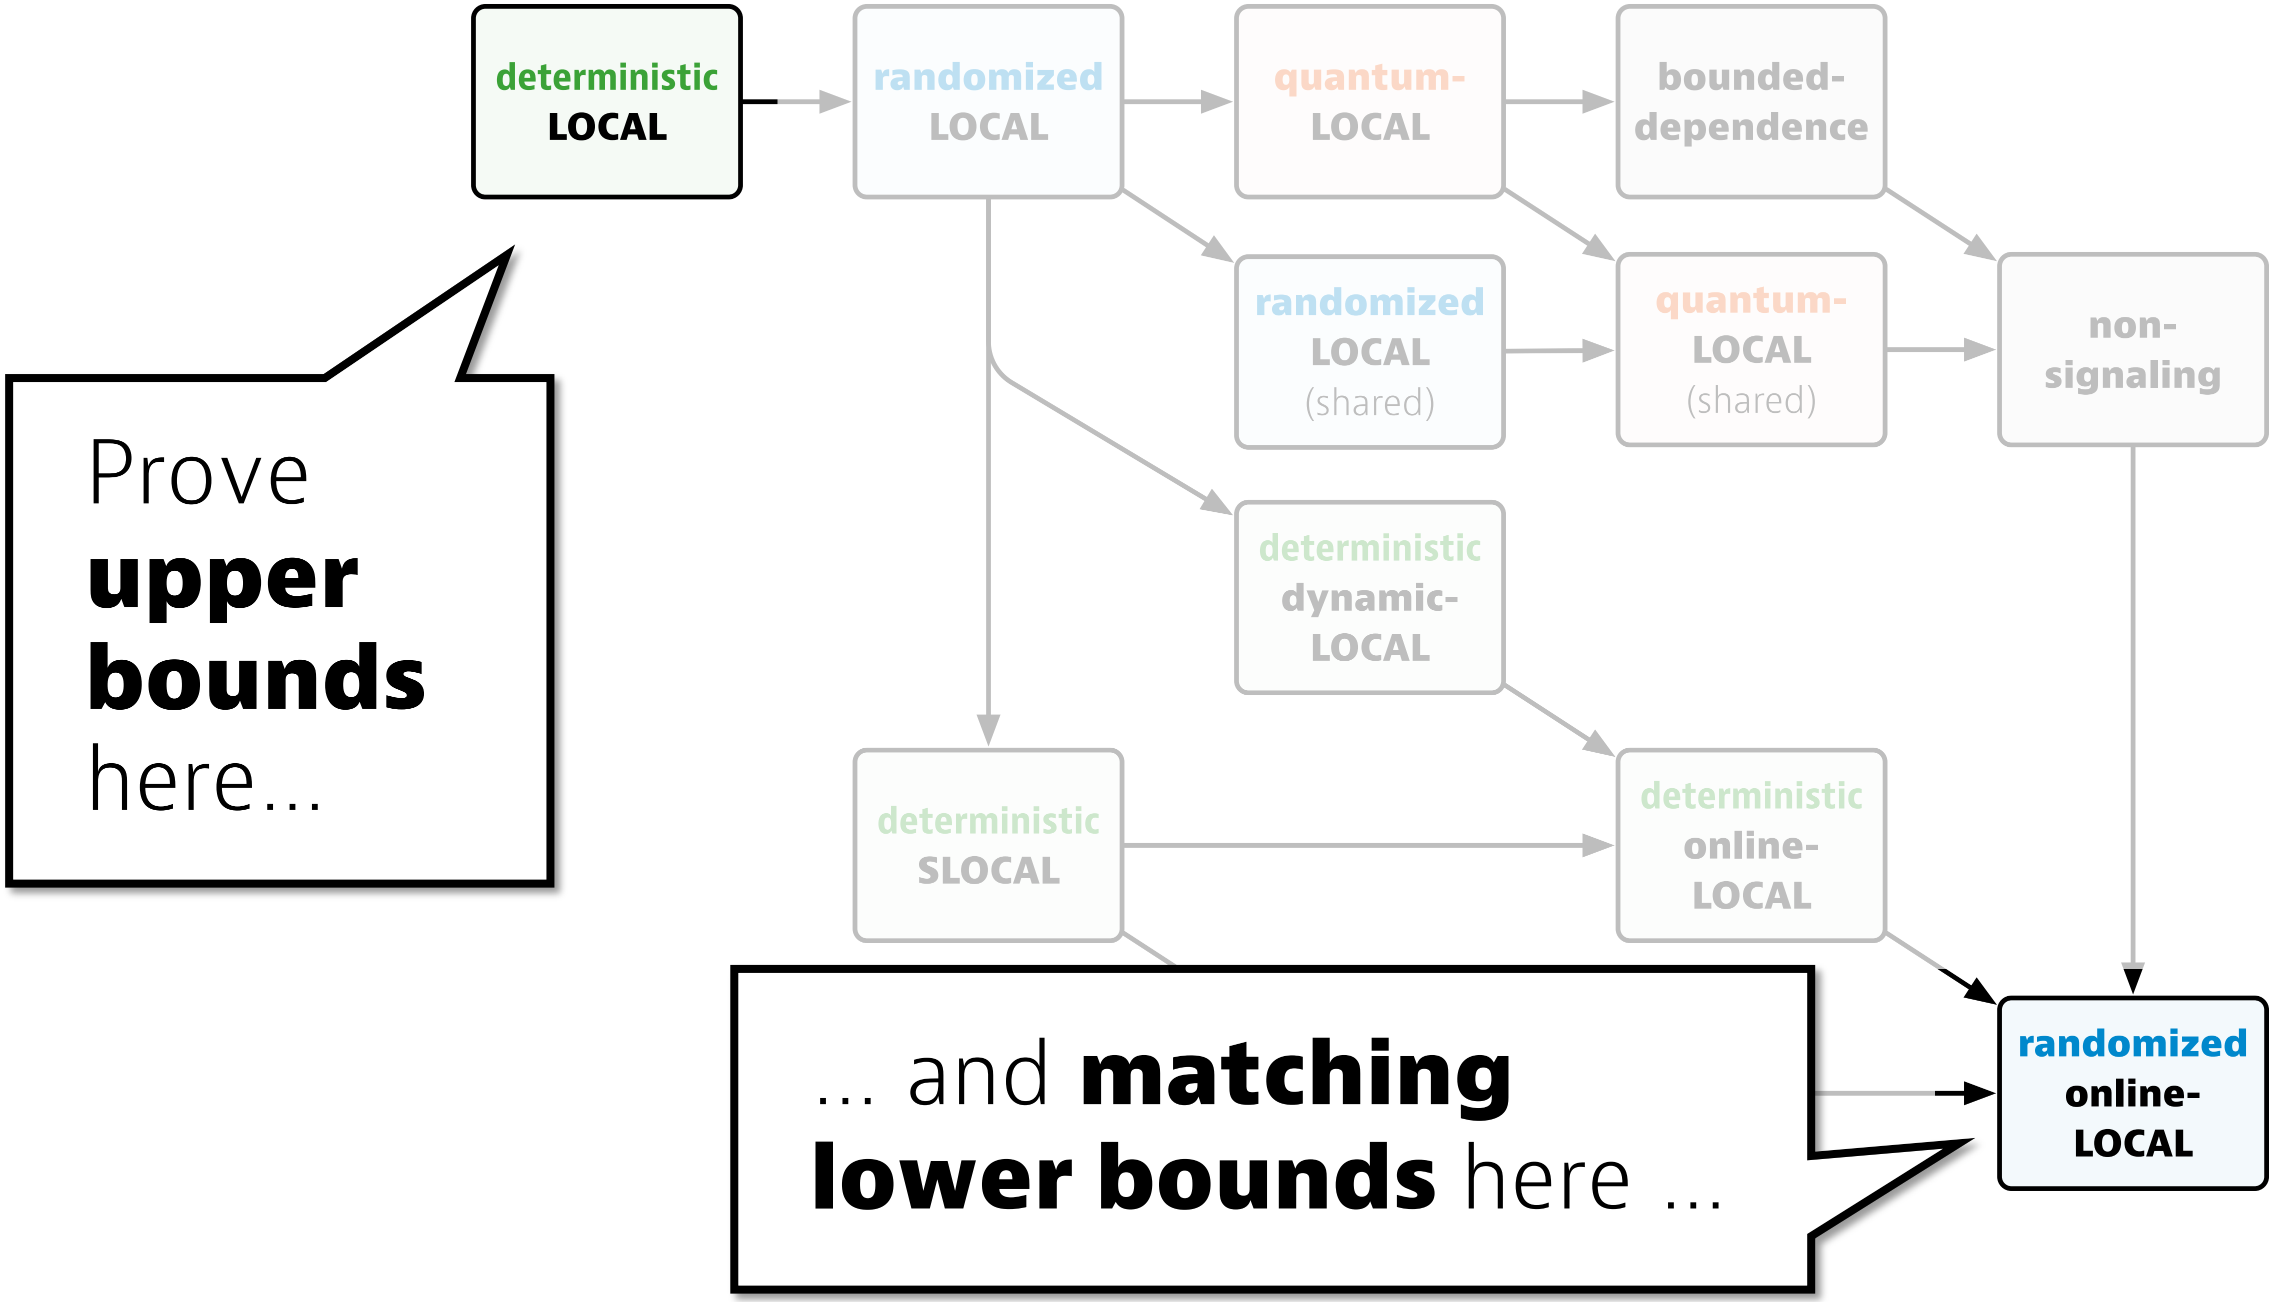
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized online-LOCAL
Prove upper bounds here…
… and matching lower bounds here …

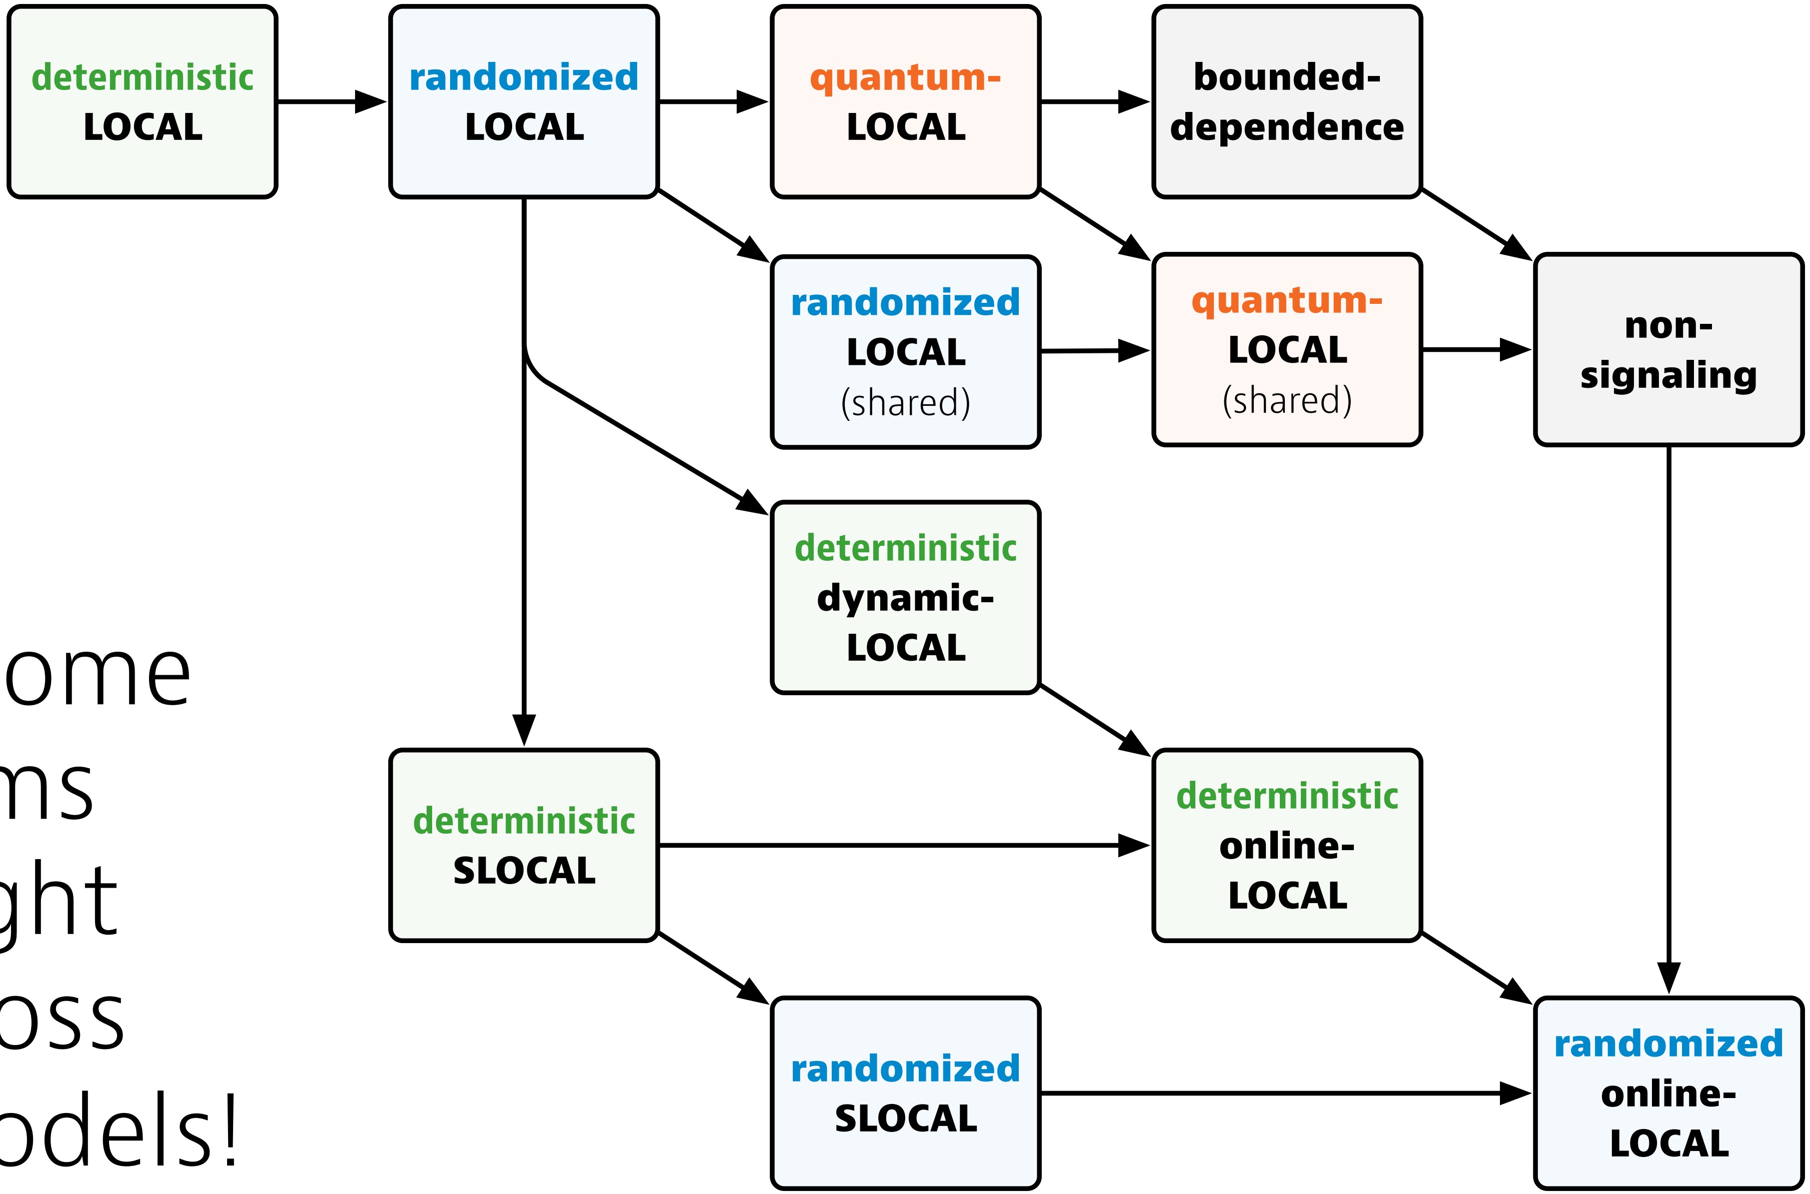

… and for some LCL problems we have tight bounds across all these models!

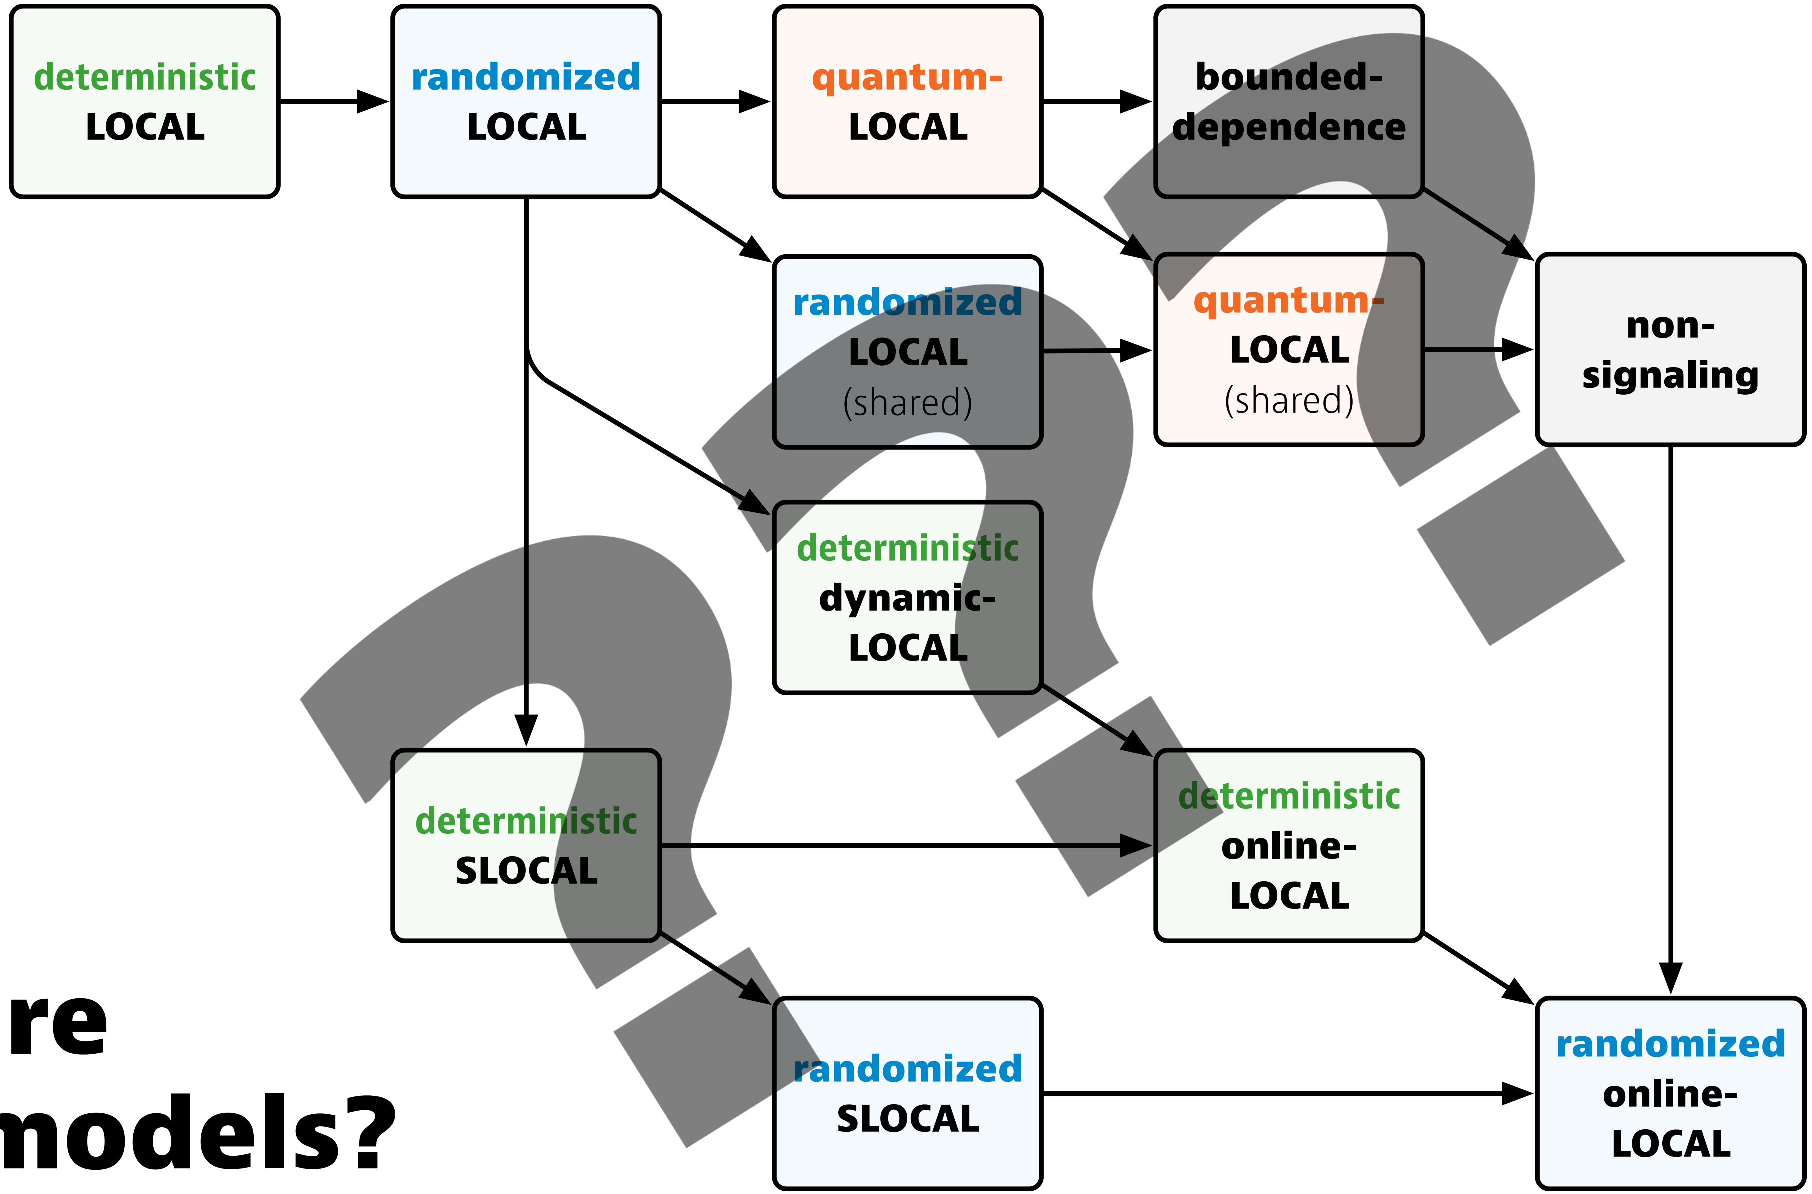

So what are all these models?

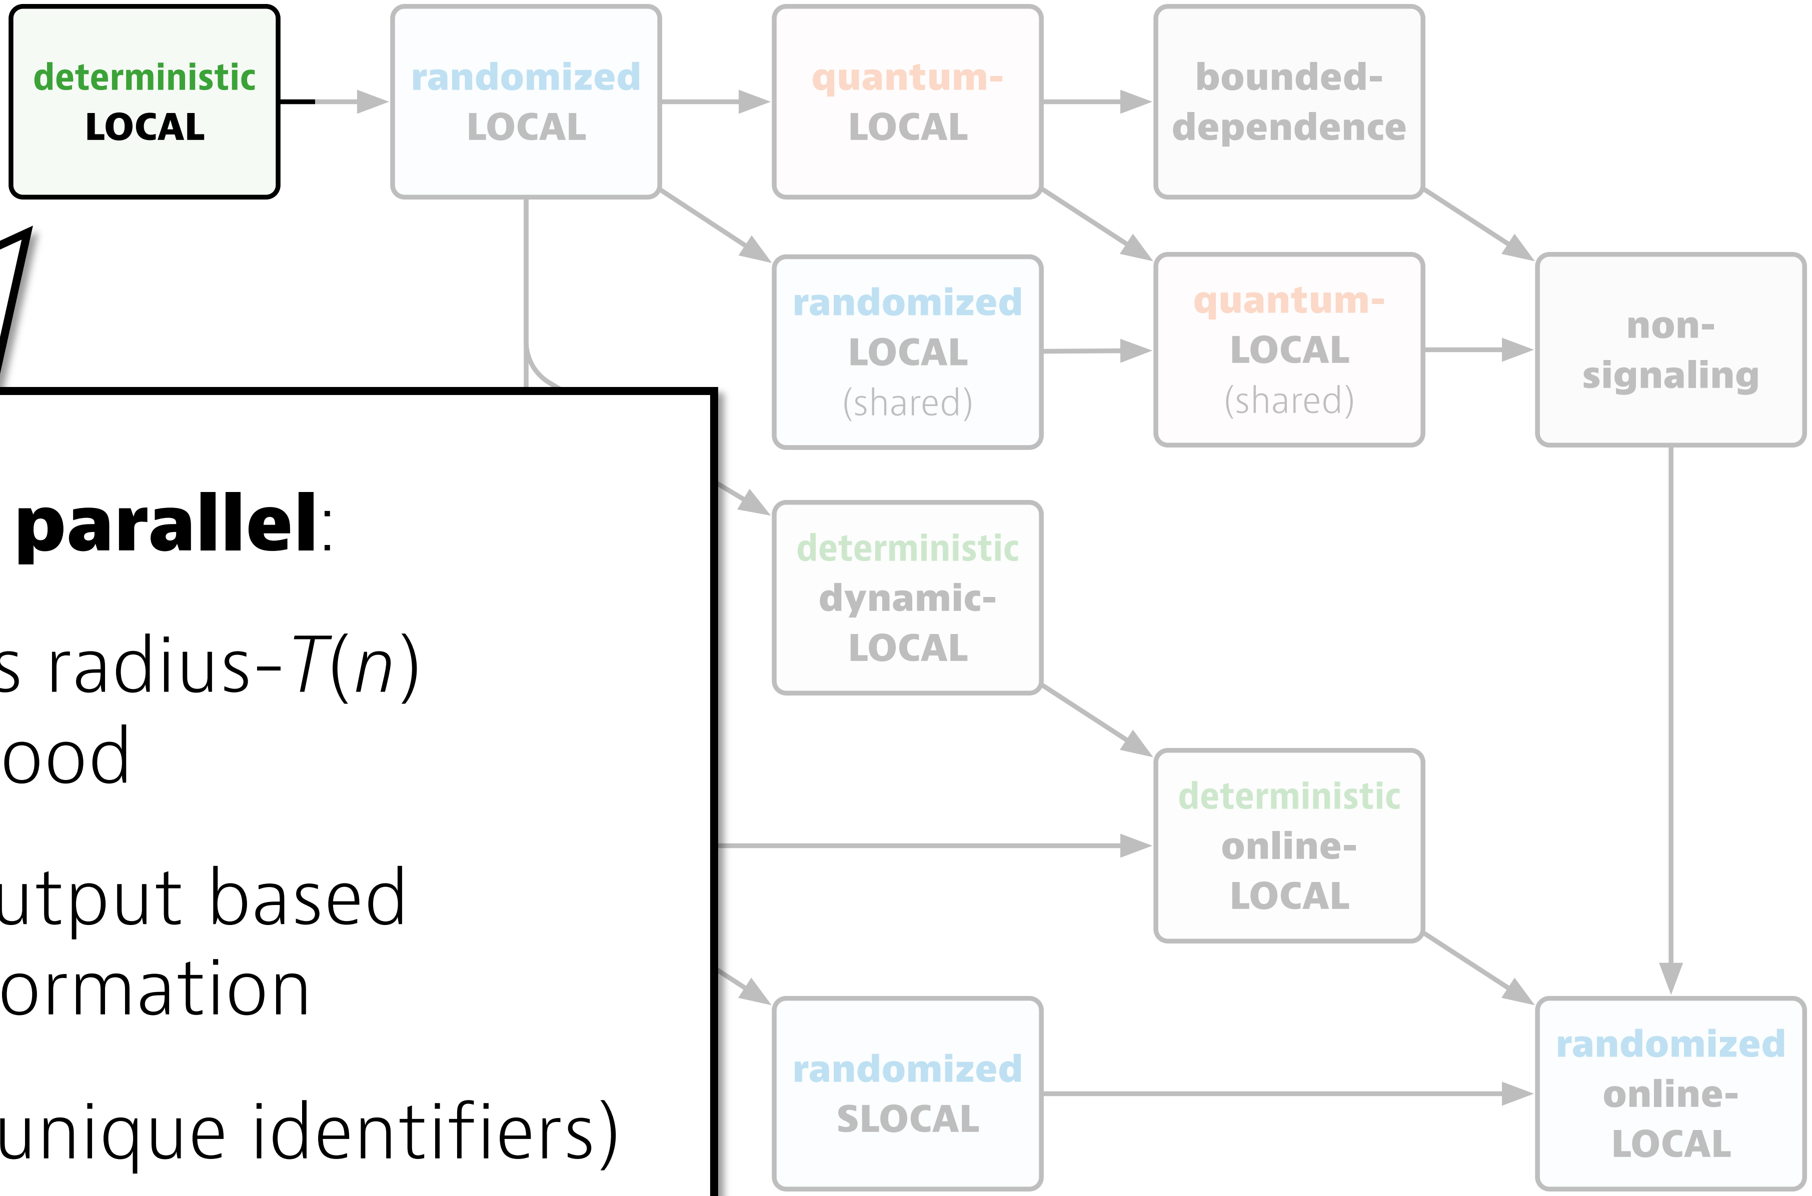

Each node in **parallel**:

- looks at its radius-$T(n)$ neighborhood
- picks its output based on this information

(nodes have unique identifiers)

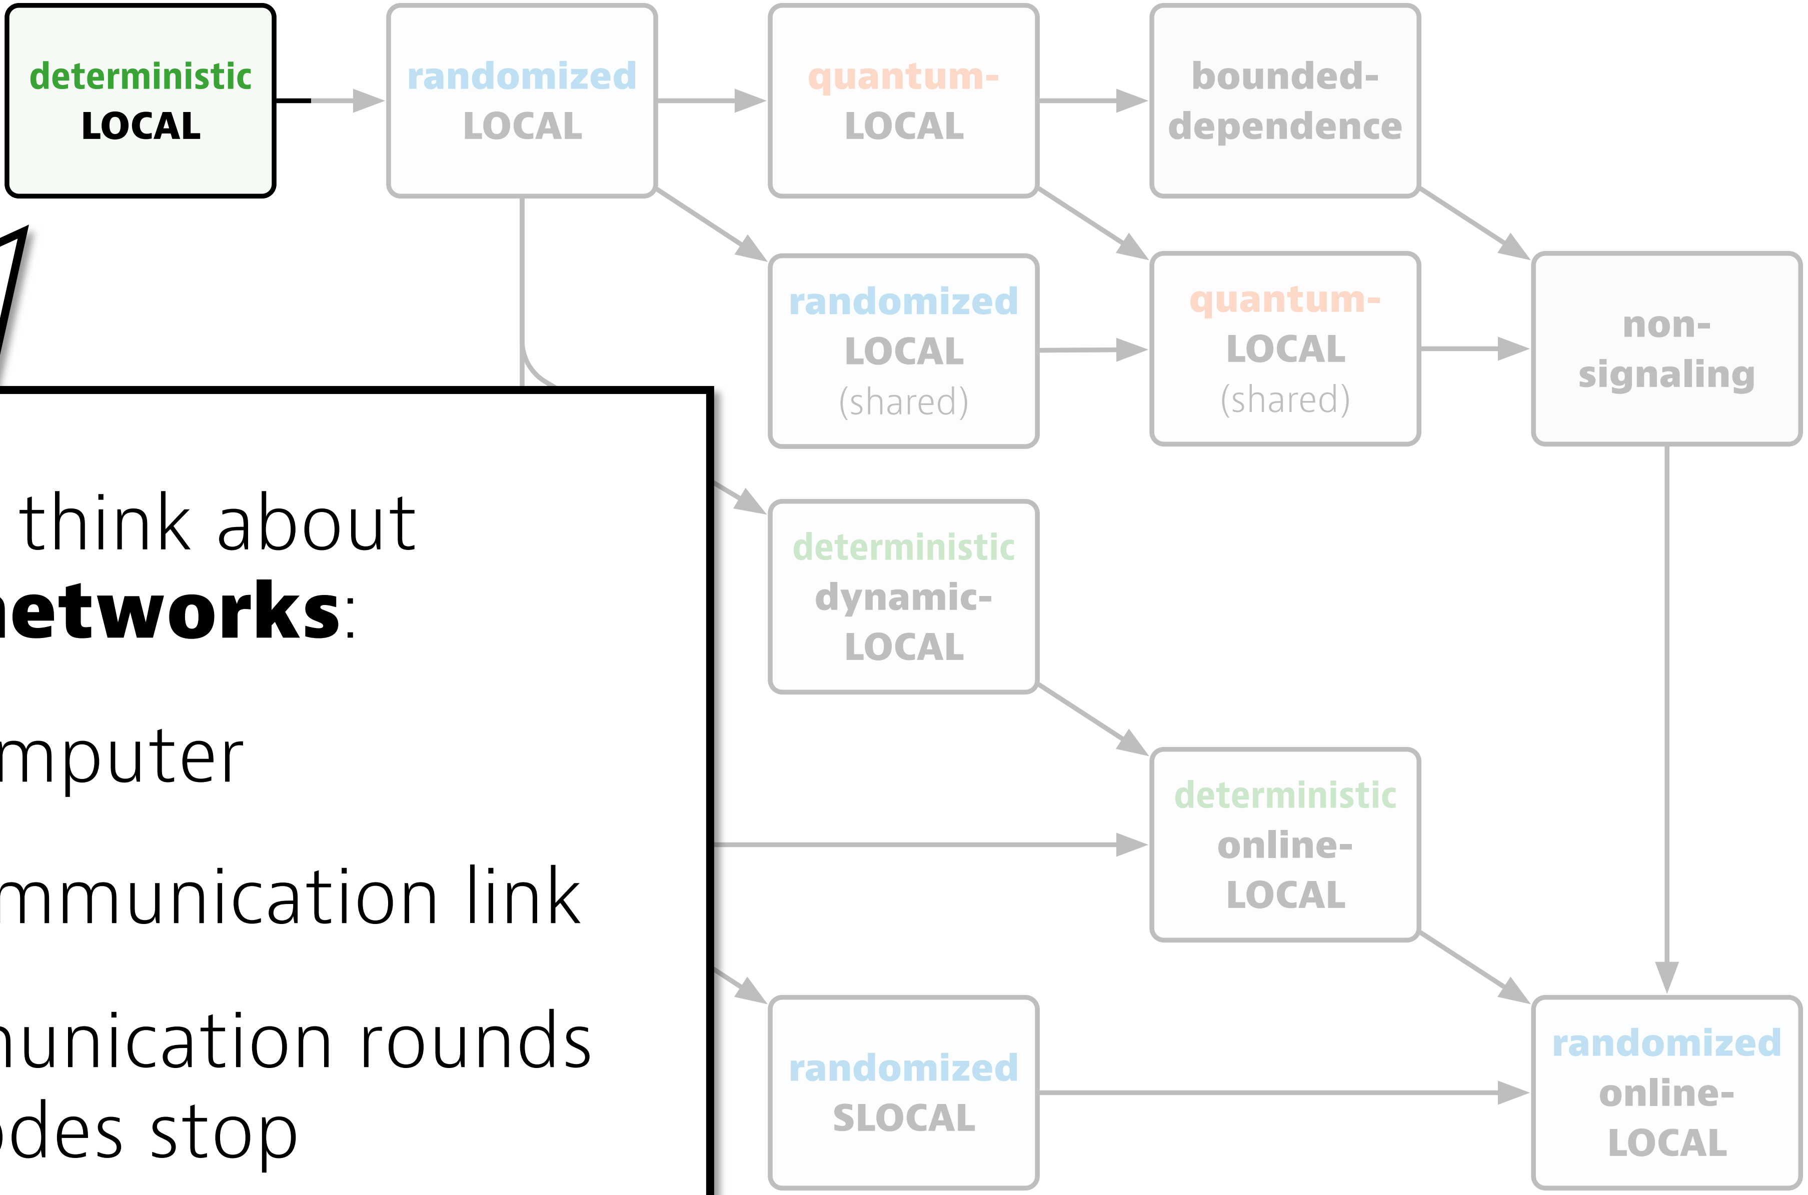

… or you can think about **computer networks**:

- node = computer
- edge = communication link
- $T(n)$ communication rounds until all nodes stop

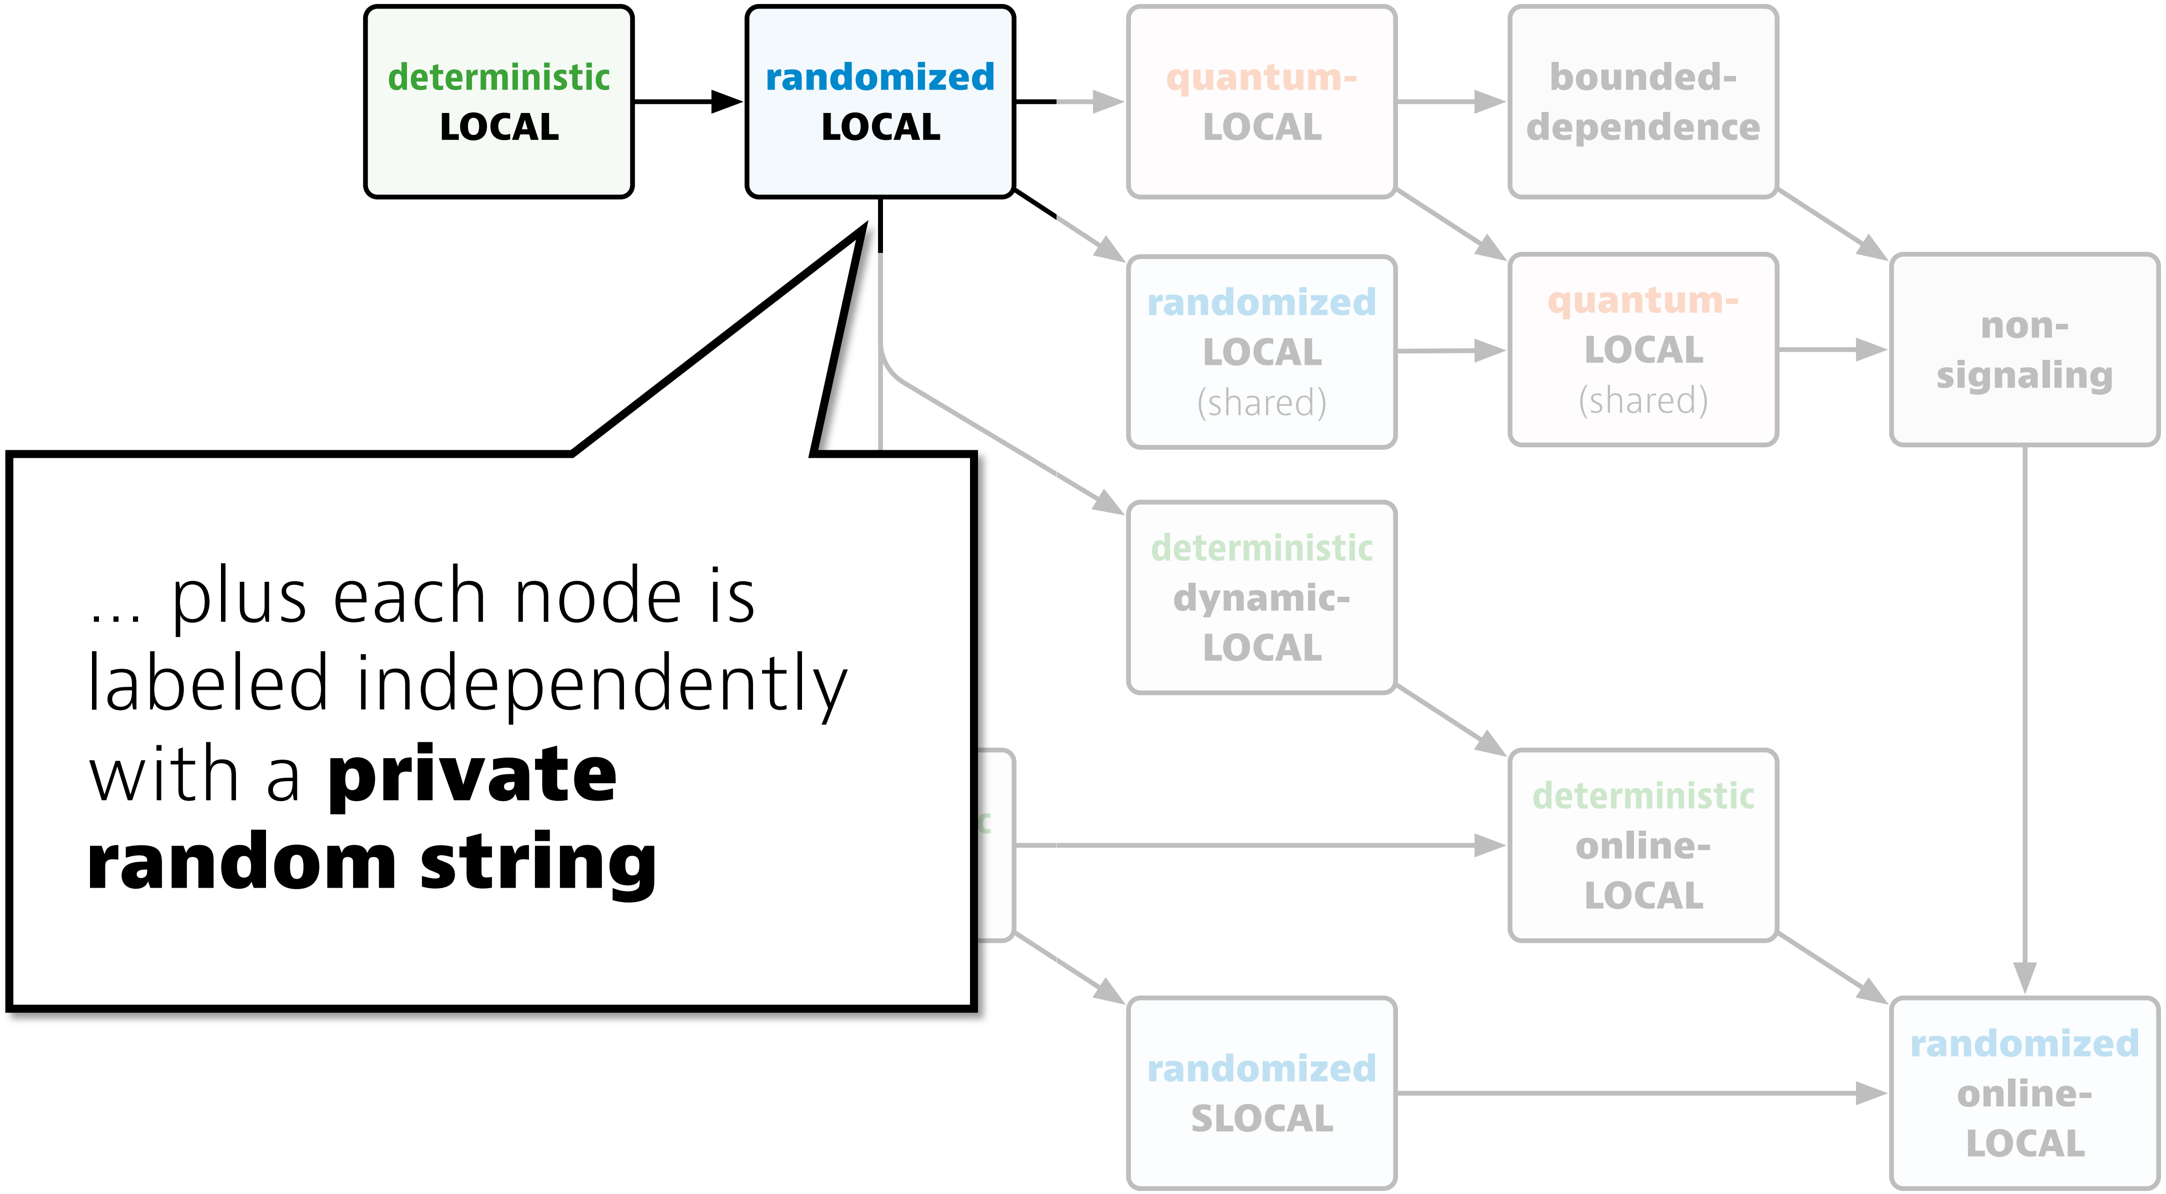
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
… plus each node is labeled independently with a **private random string**

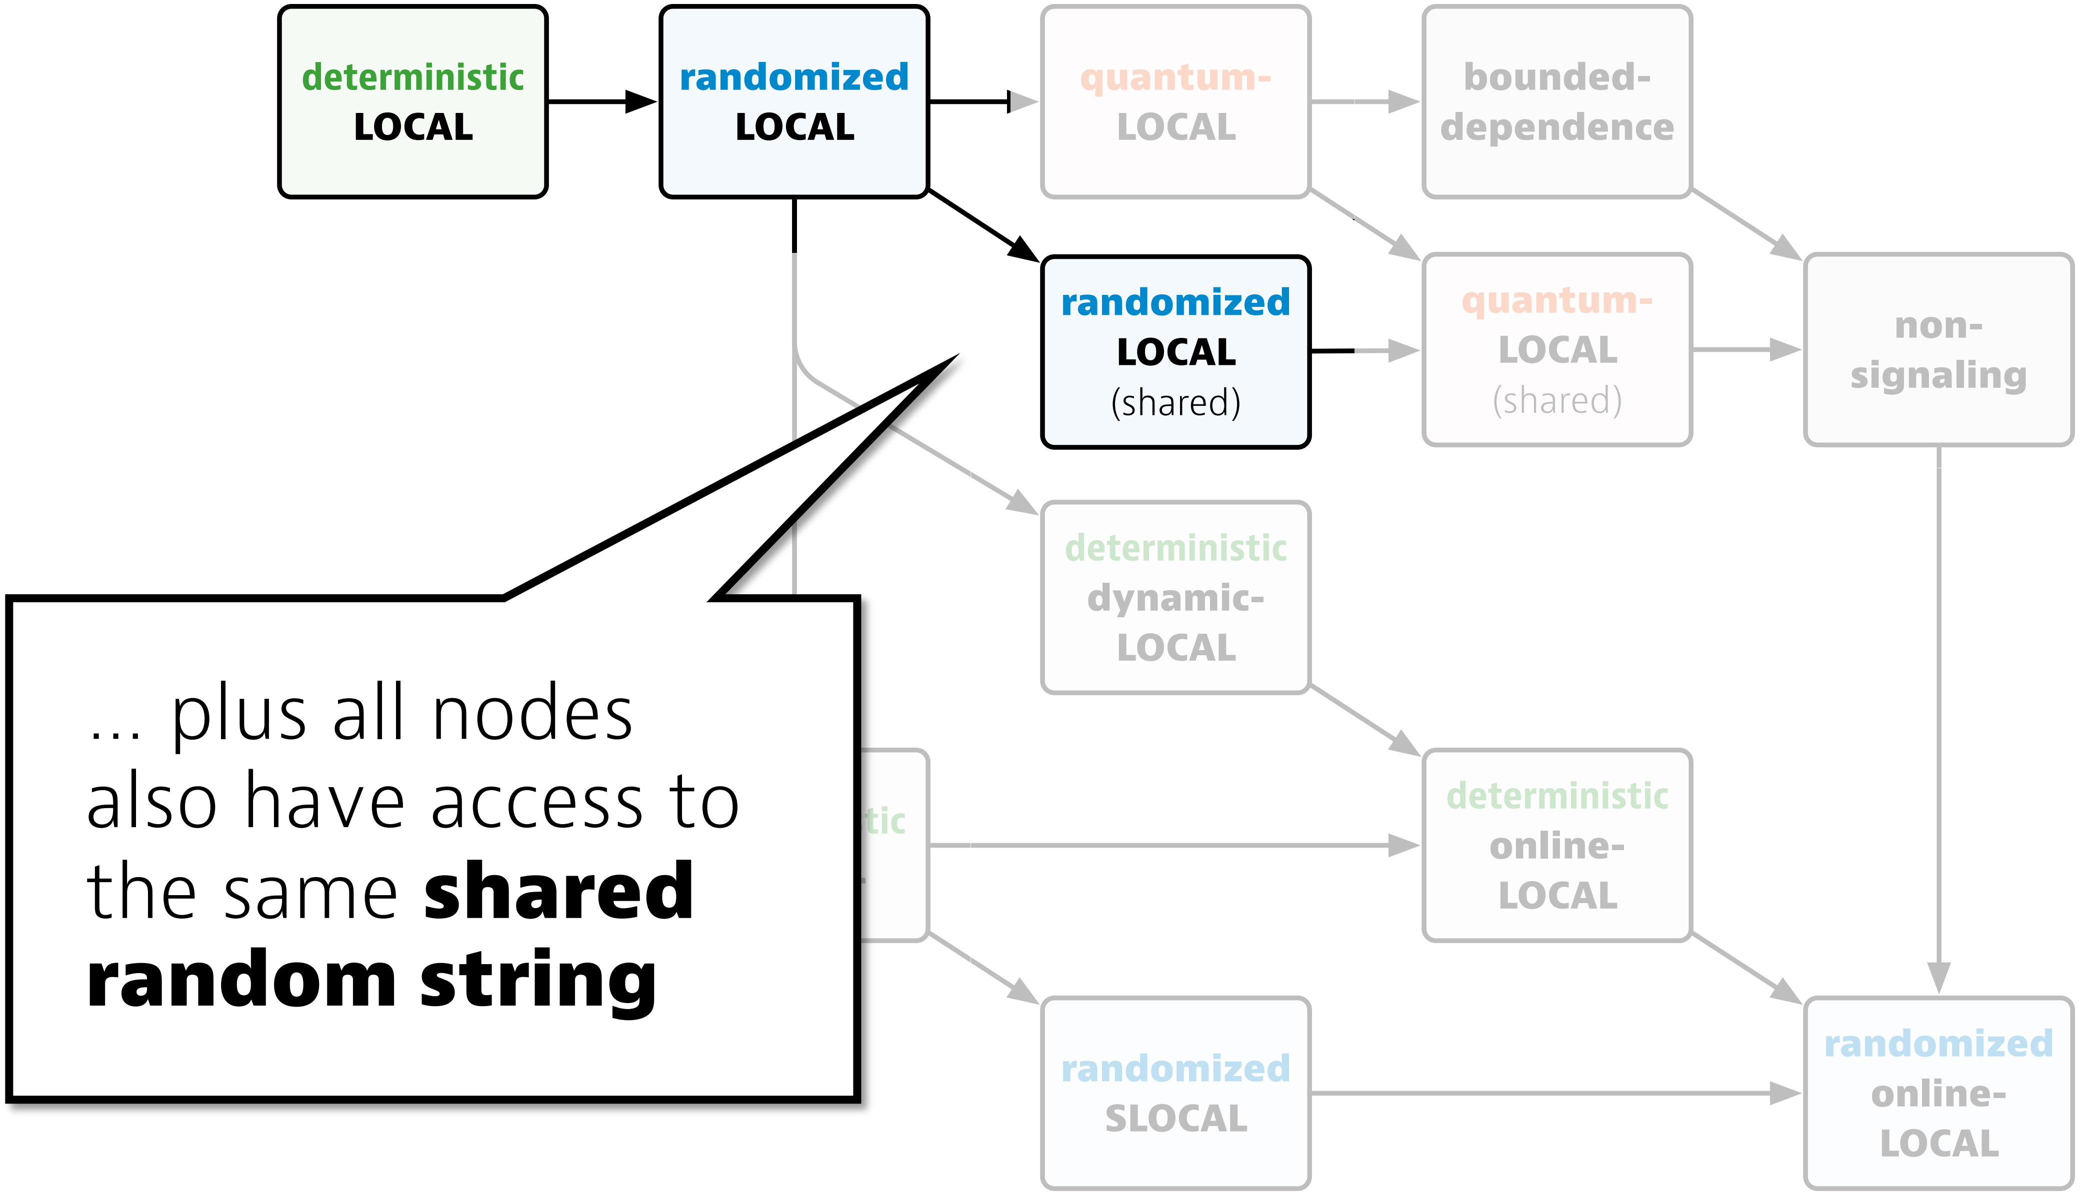
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
… plus all nodes also have access to the same **shared random string**

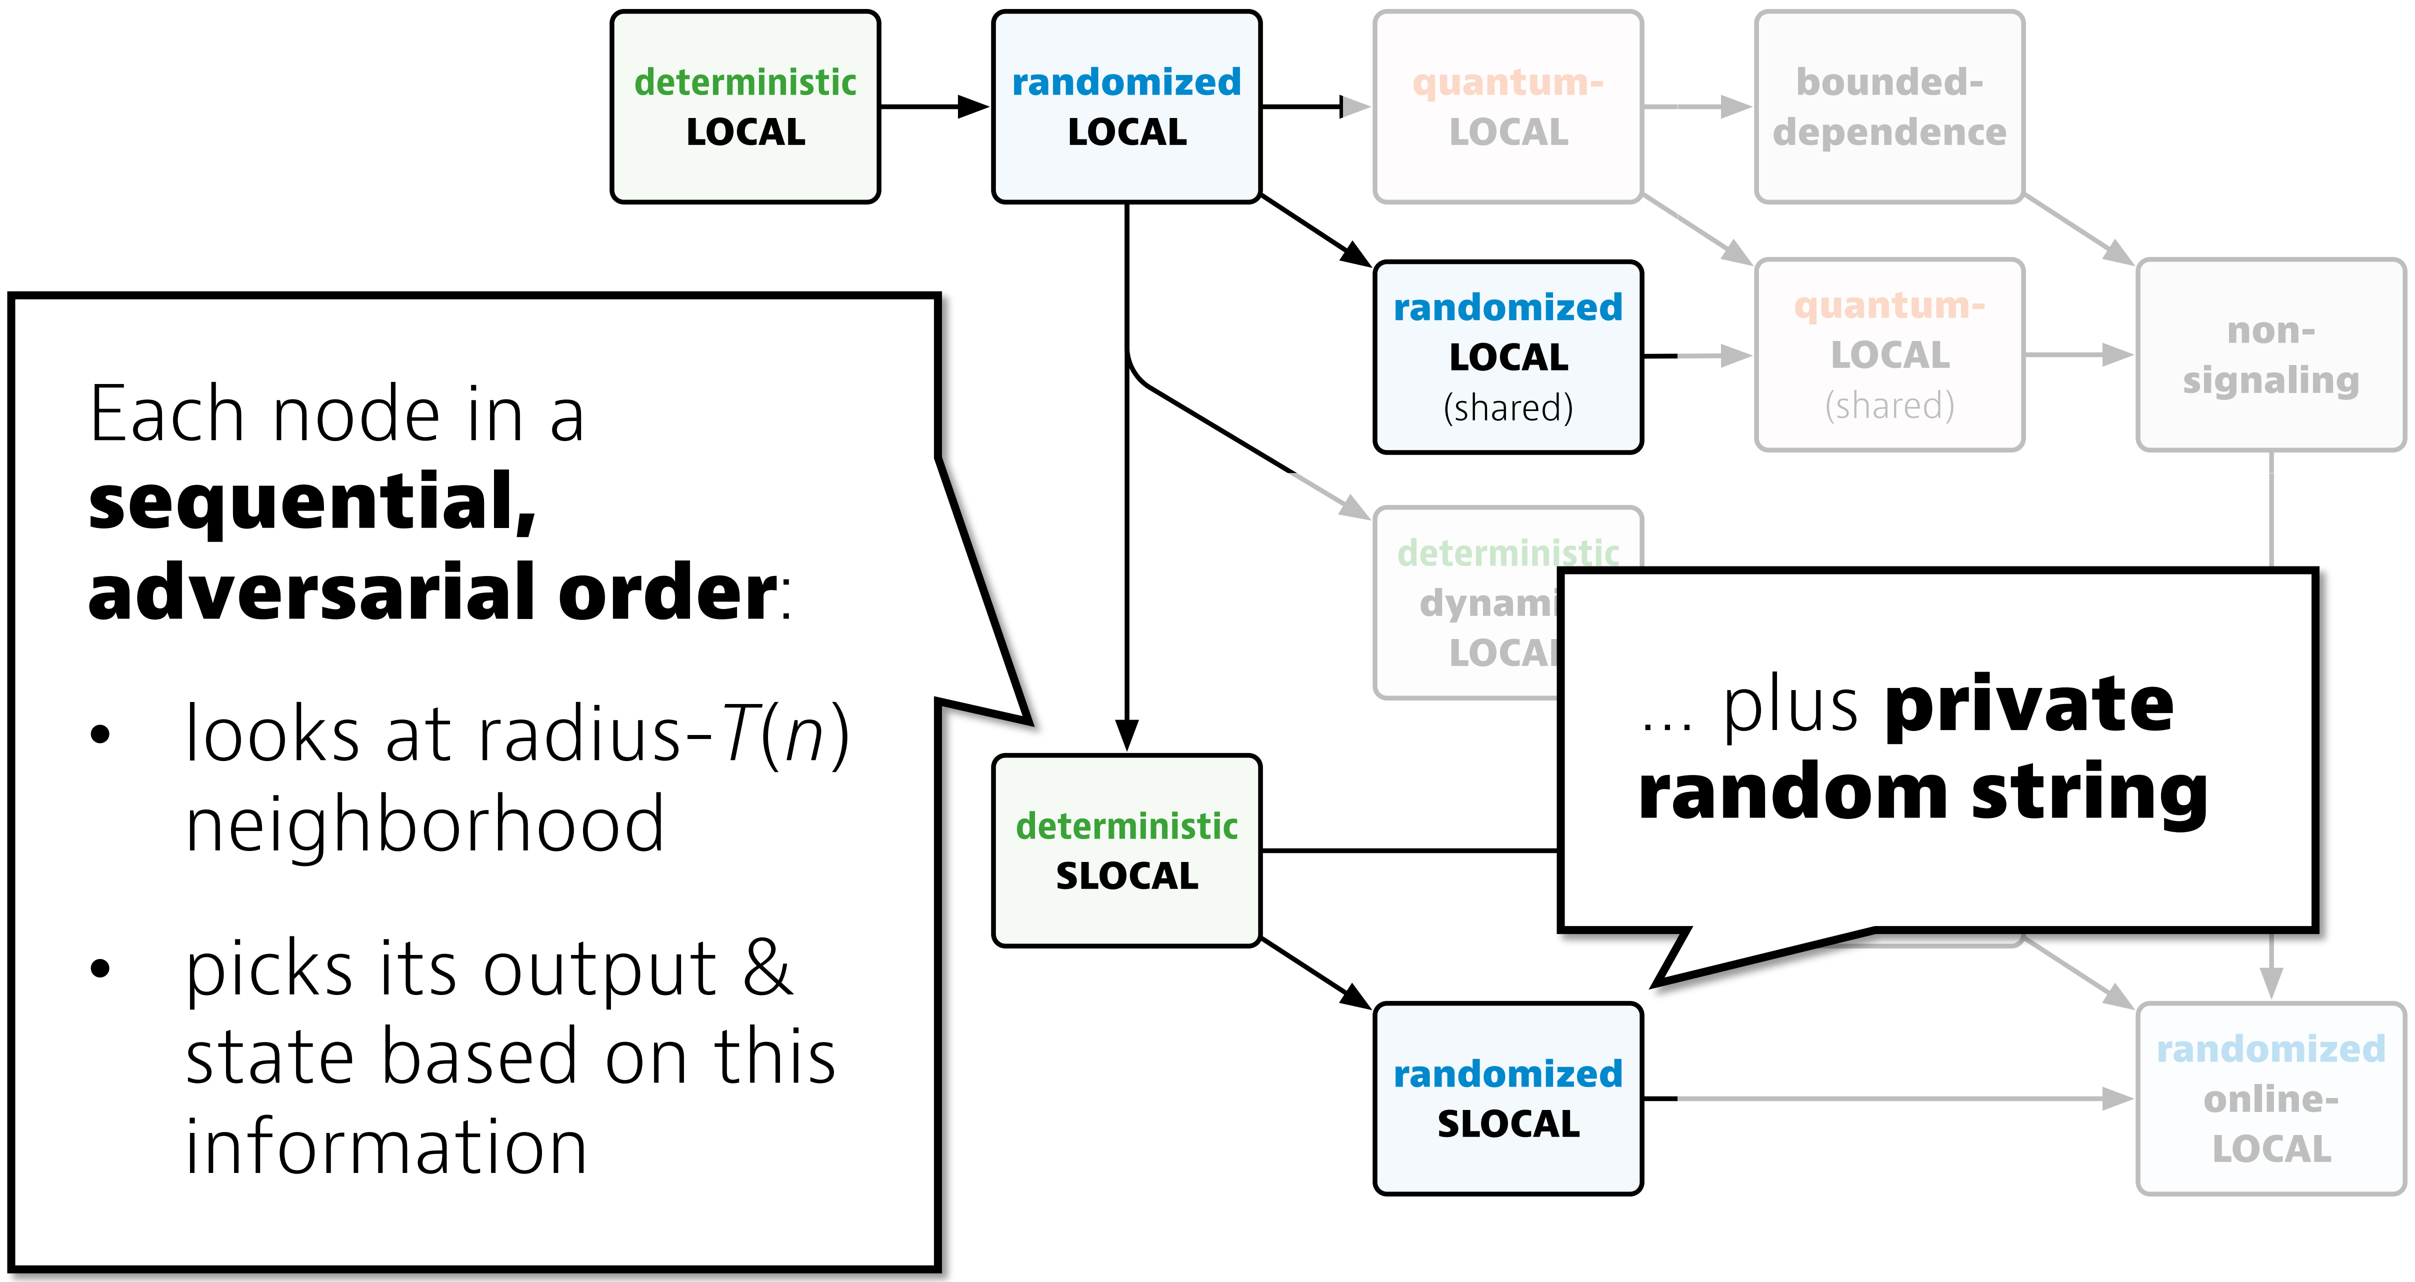
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynami LOCAL
deterministic SLOCAL
randomized SLOCAL
randomized online-LOCAL
Each node in a **sequential, adversarial order**:
• looks at radius-$T(n)$ neighborhood
• picks its output & state based on this information
… plus **private random string**

Sequential ordering gives enough power to **derandomize** local algorithms!

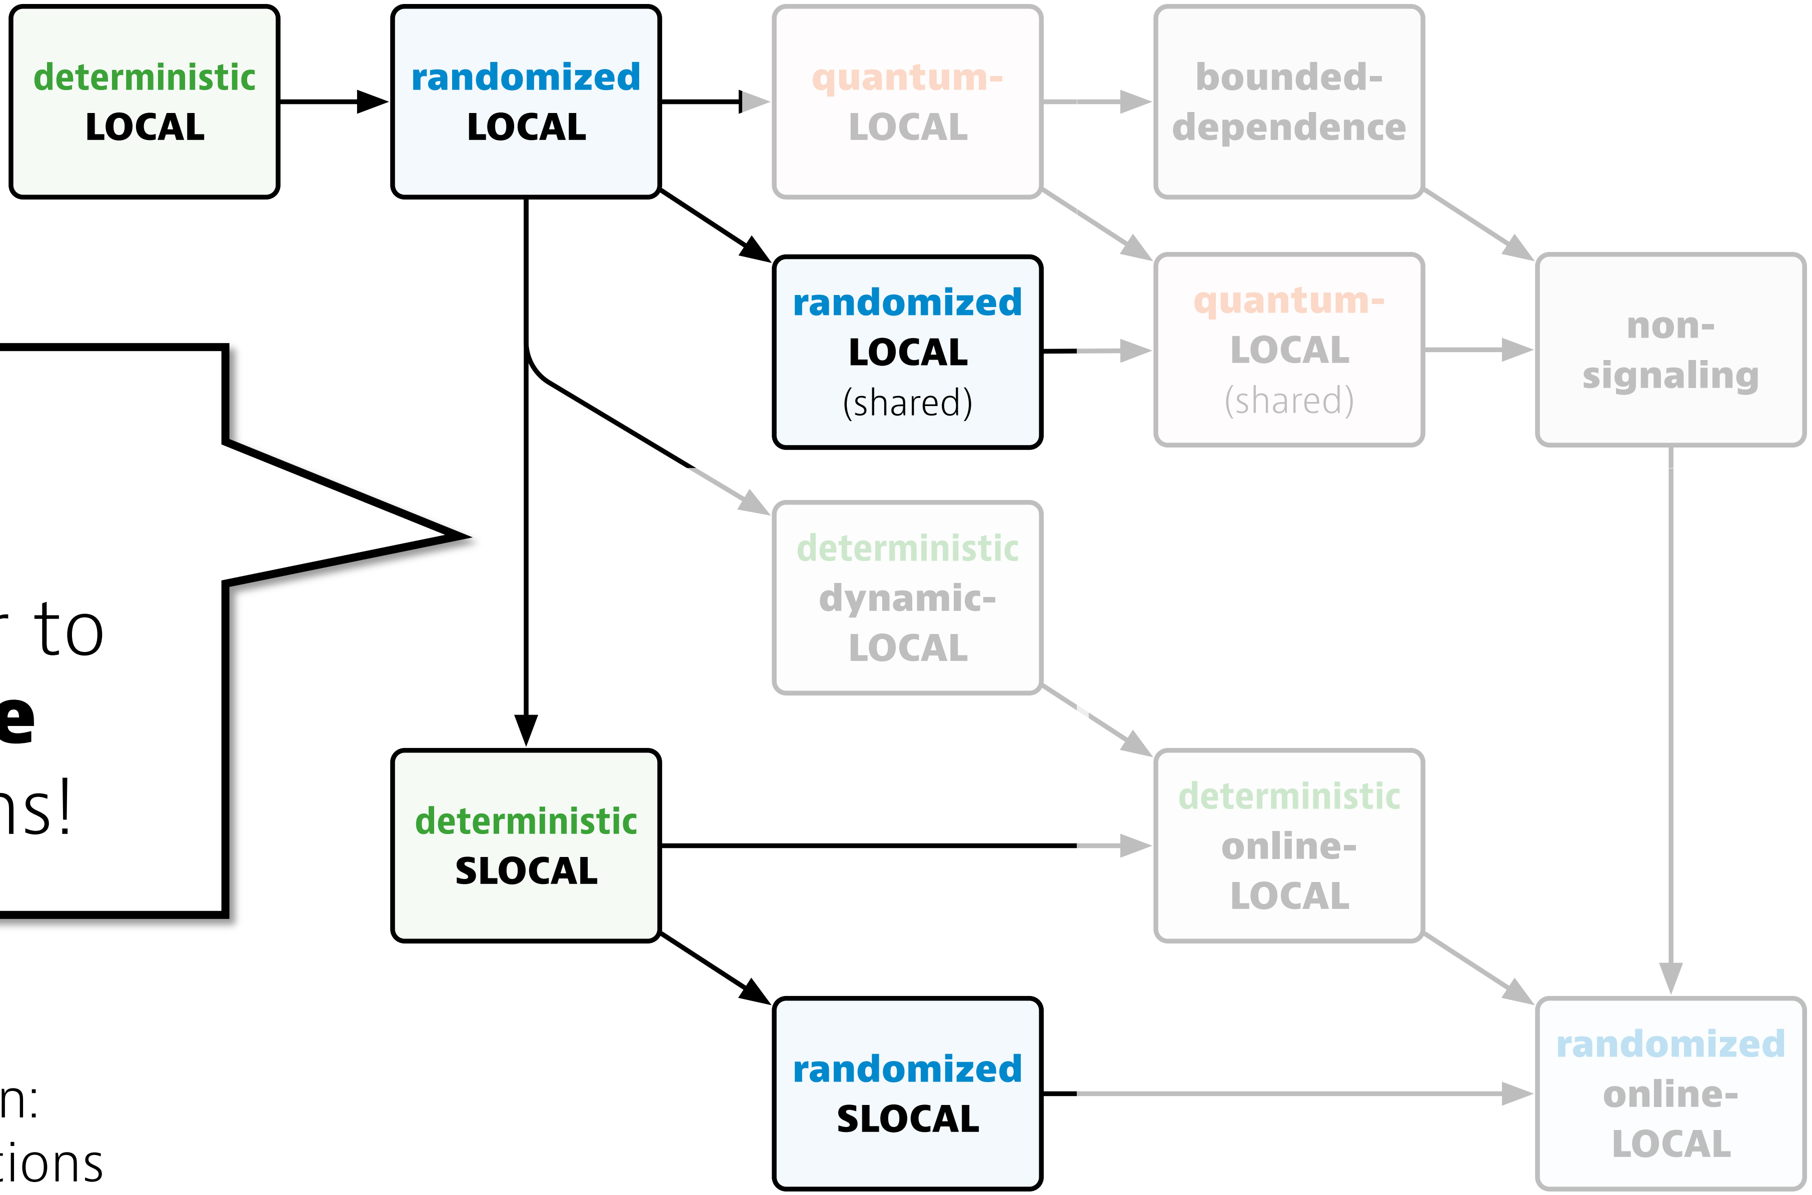

Ghaffari, Harris, Kuhn:
conditional expectations

Adversary adds nodes and edges one by one

- We can **see everything**
- Have to maintain valid solution

We can **change** our output only within distance $T(n)$ from a point of change

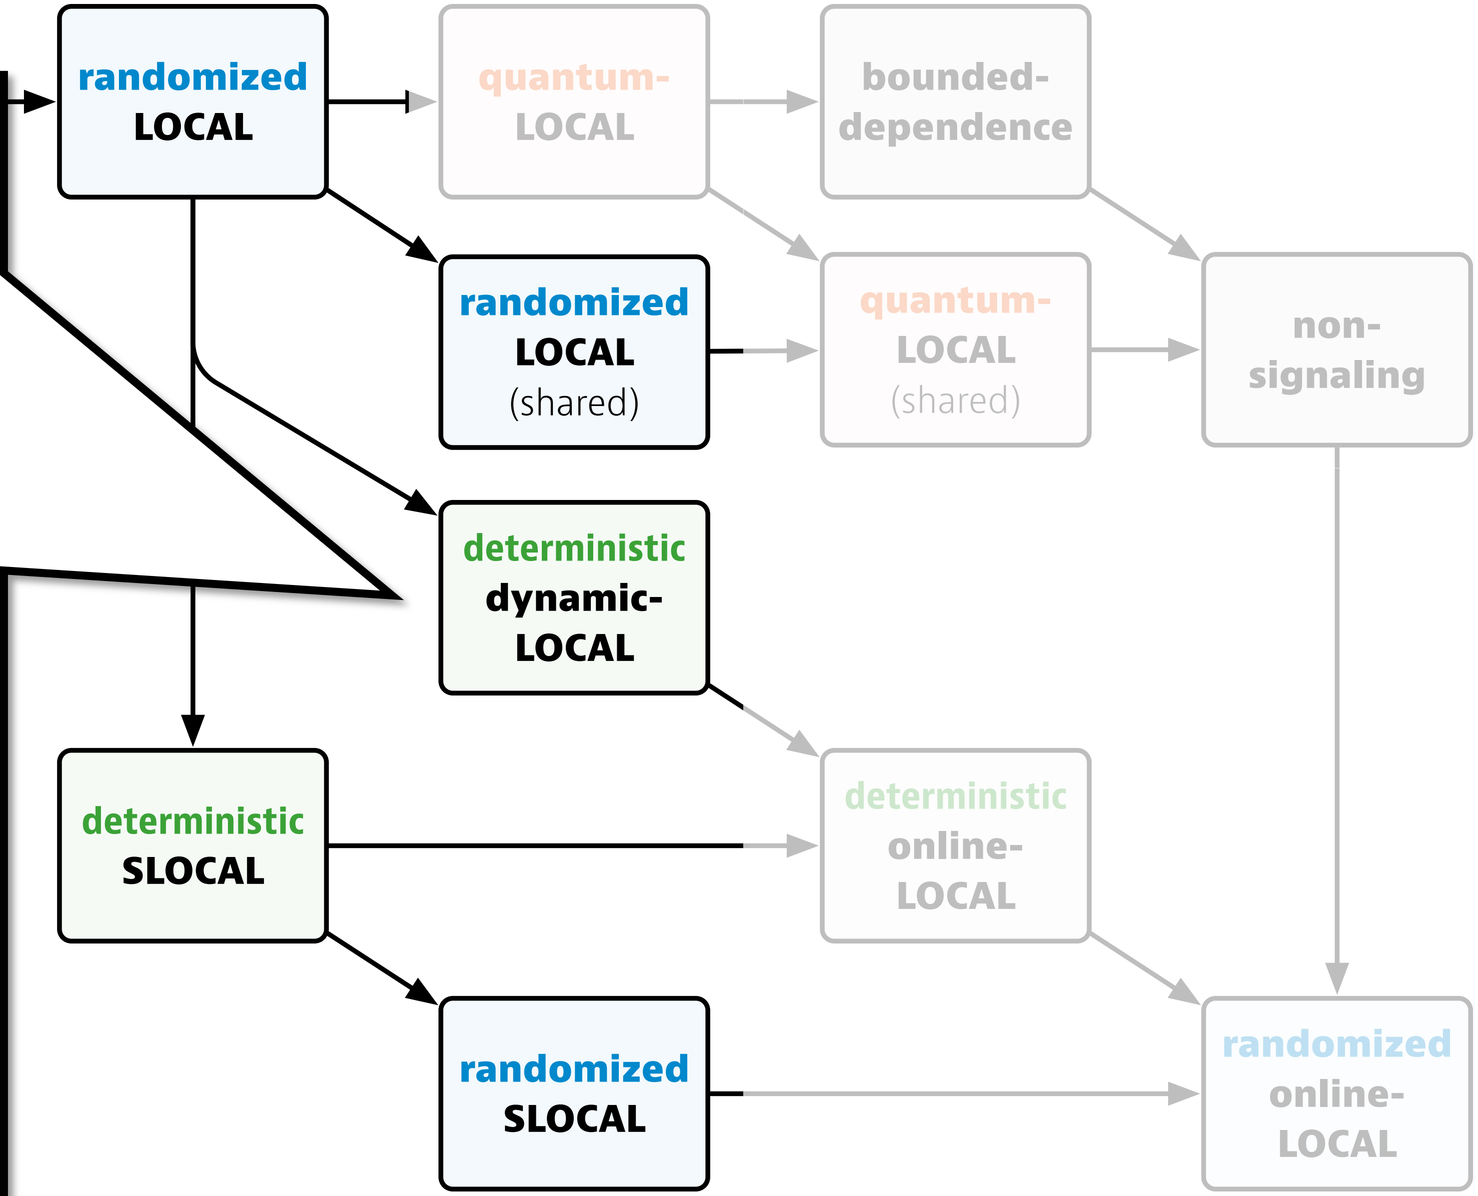

Limited what we can **see**

Limited what we can **change**

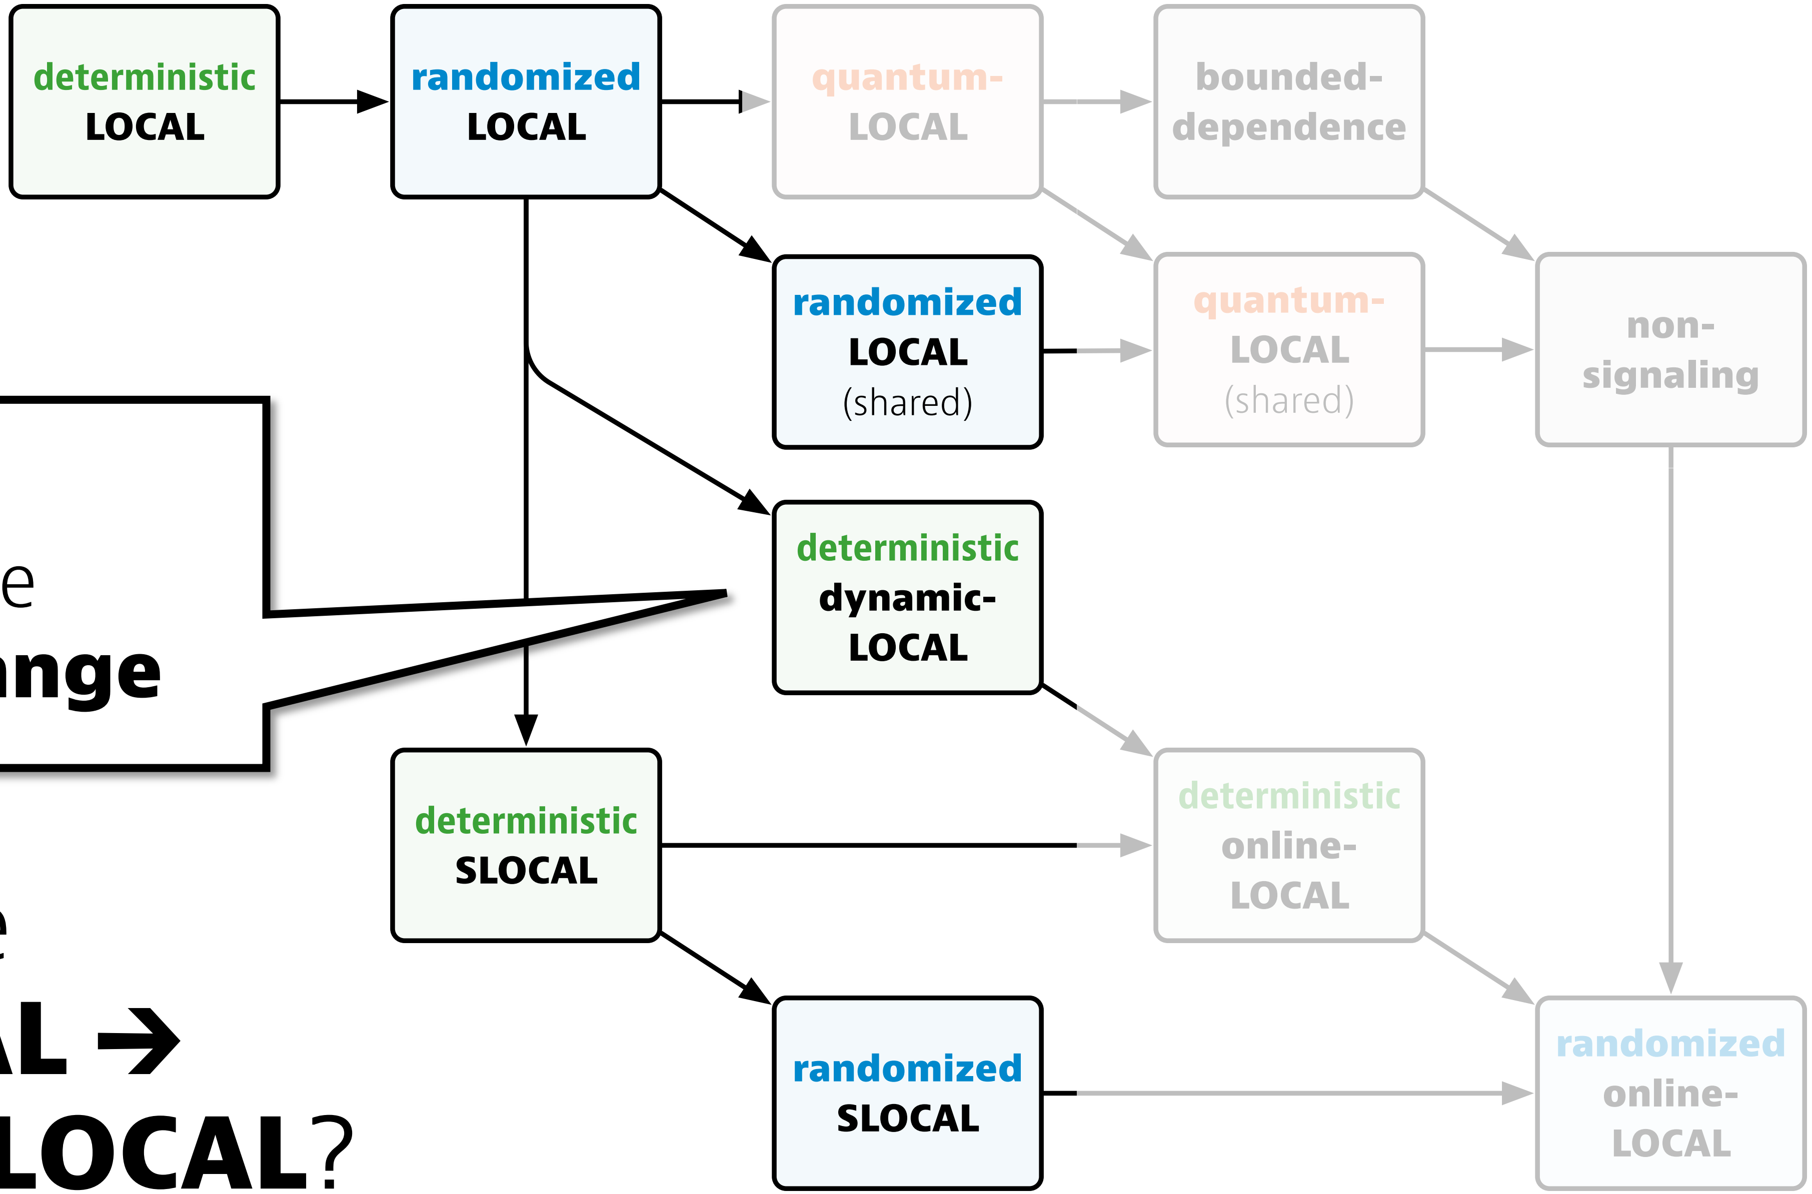

Why do we have **LOCAL → dynamic-LOCAL**?

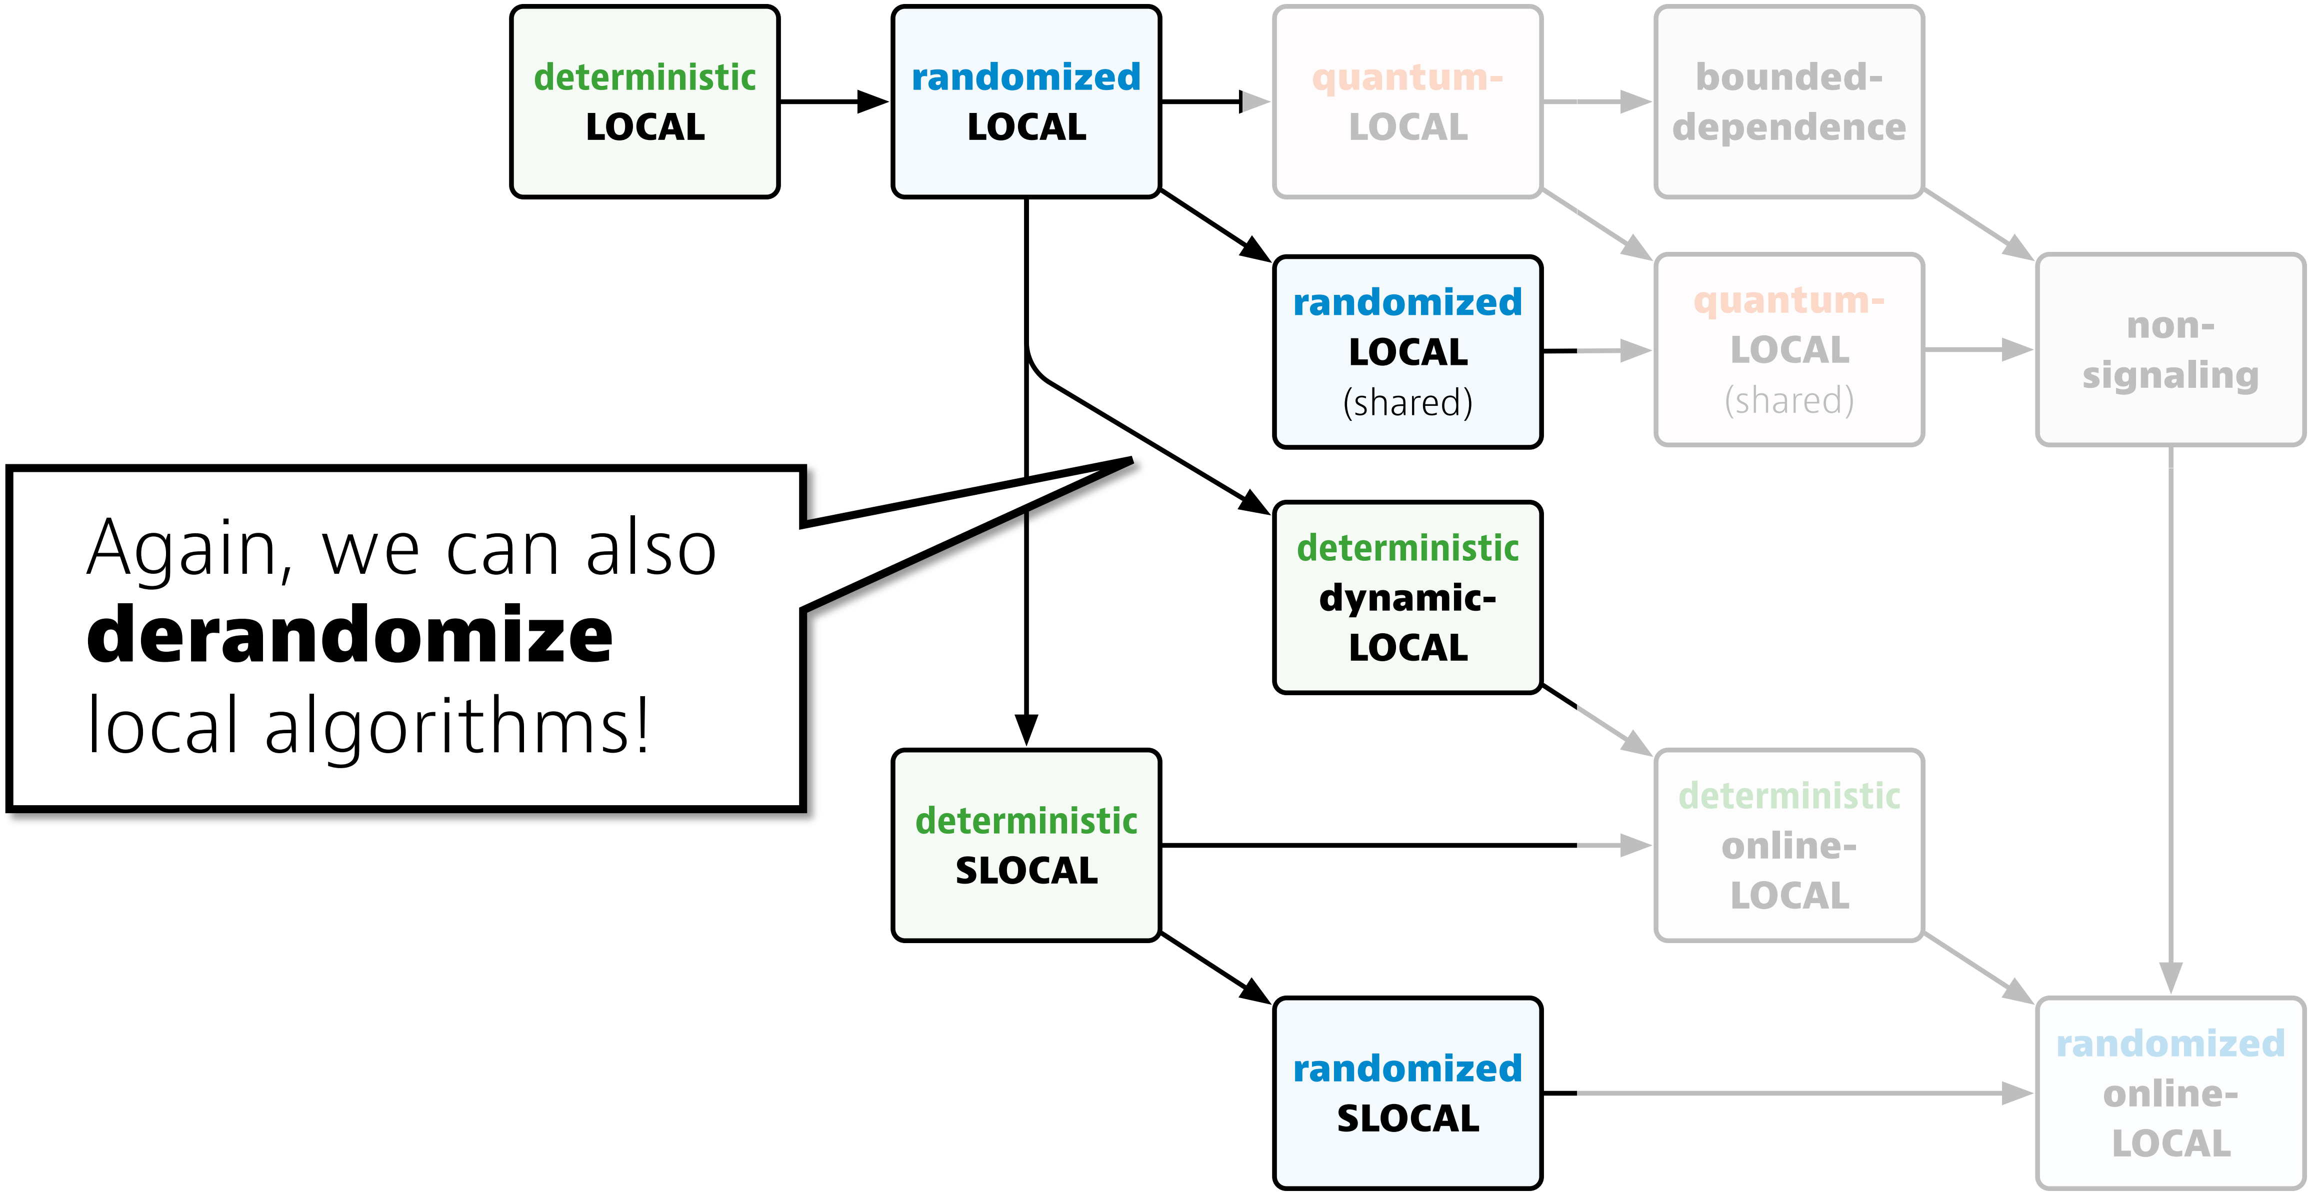

deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
Again, we can also **derandomize** local algorithms!
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

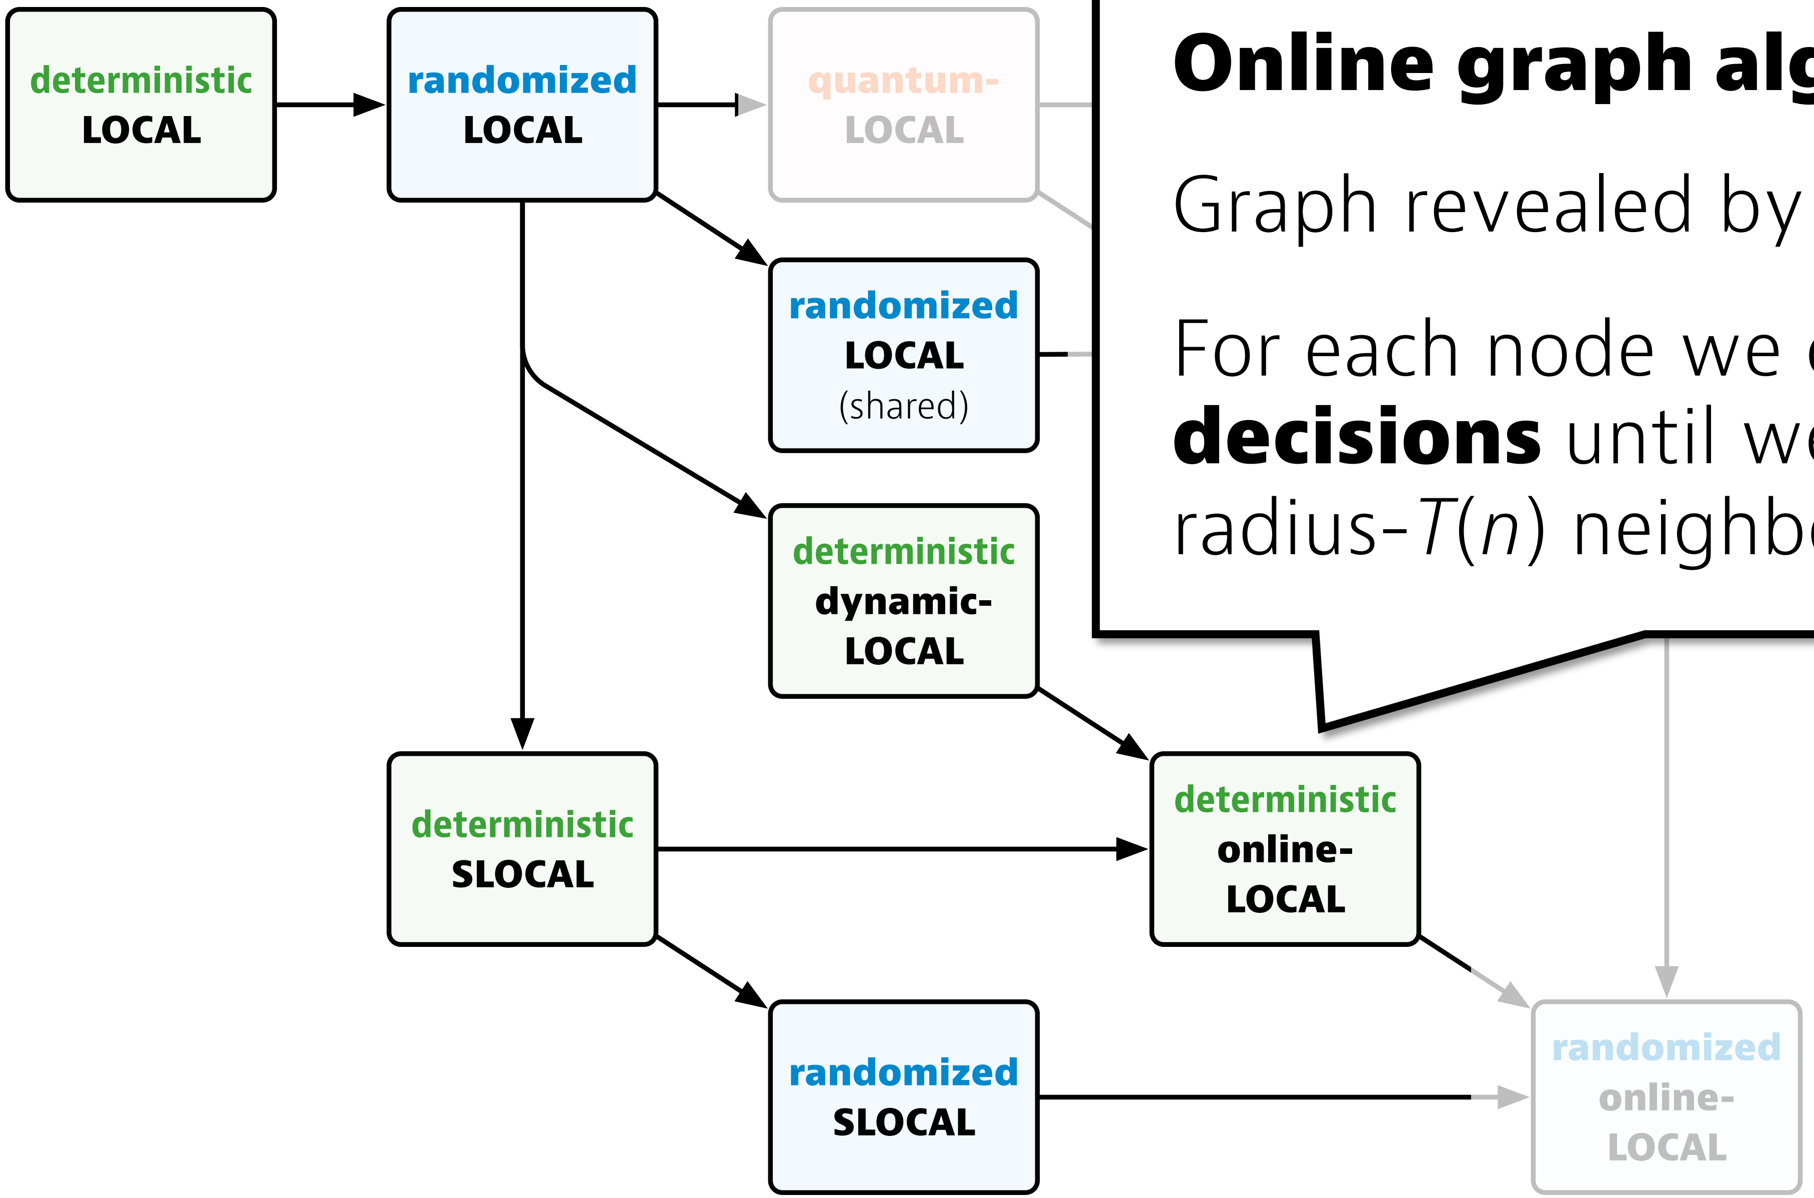

Online graph algorithms

Graph revealed by adversary

For each node we can **delay decisions** until we see its radius-$T(n)$ neighborhood

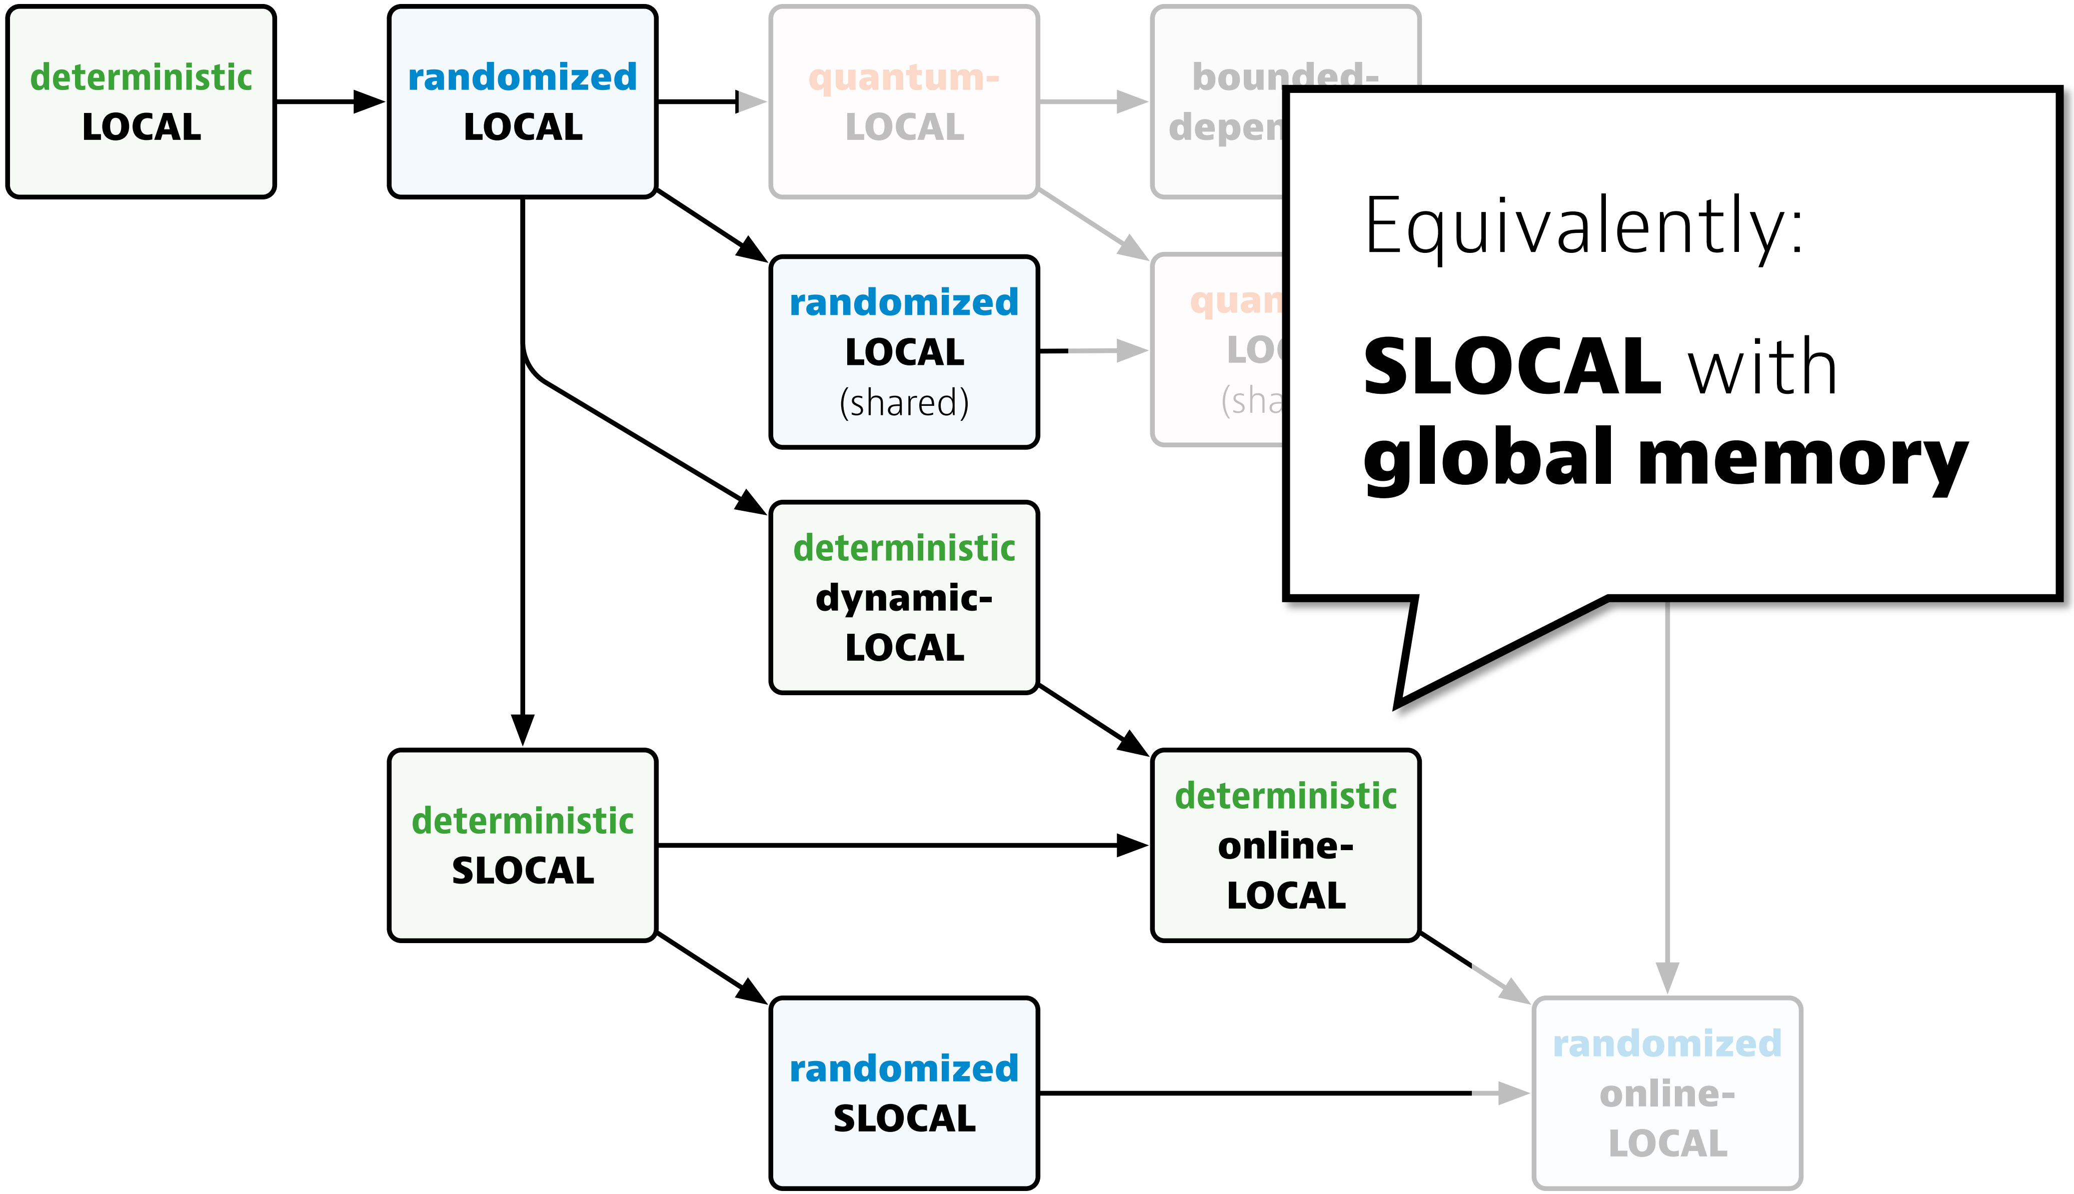
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-depen
randomized LOCAL (shared)
quan LO (sha
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
Equivalently:
SLOCAL with global memory

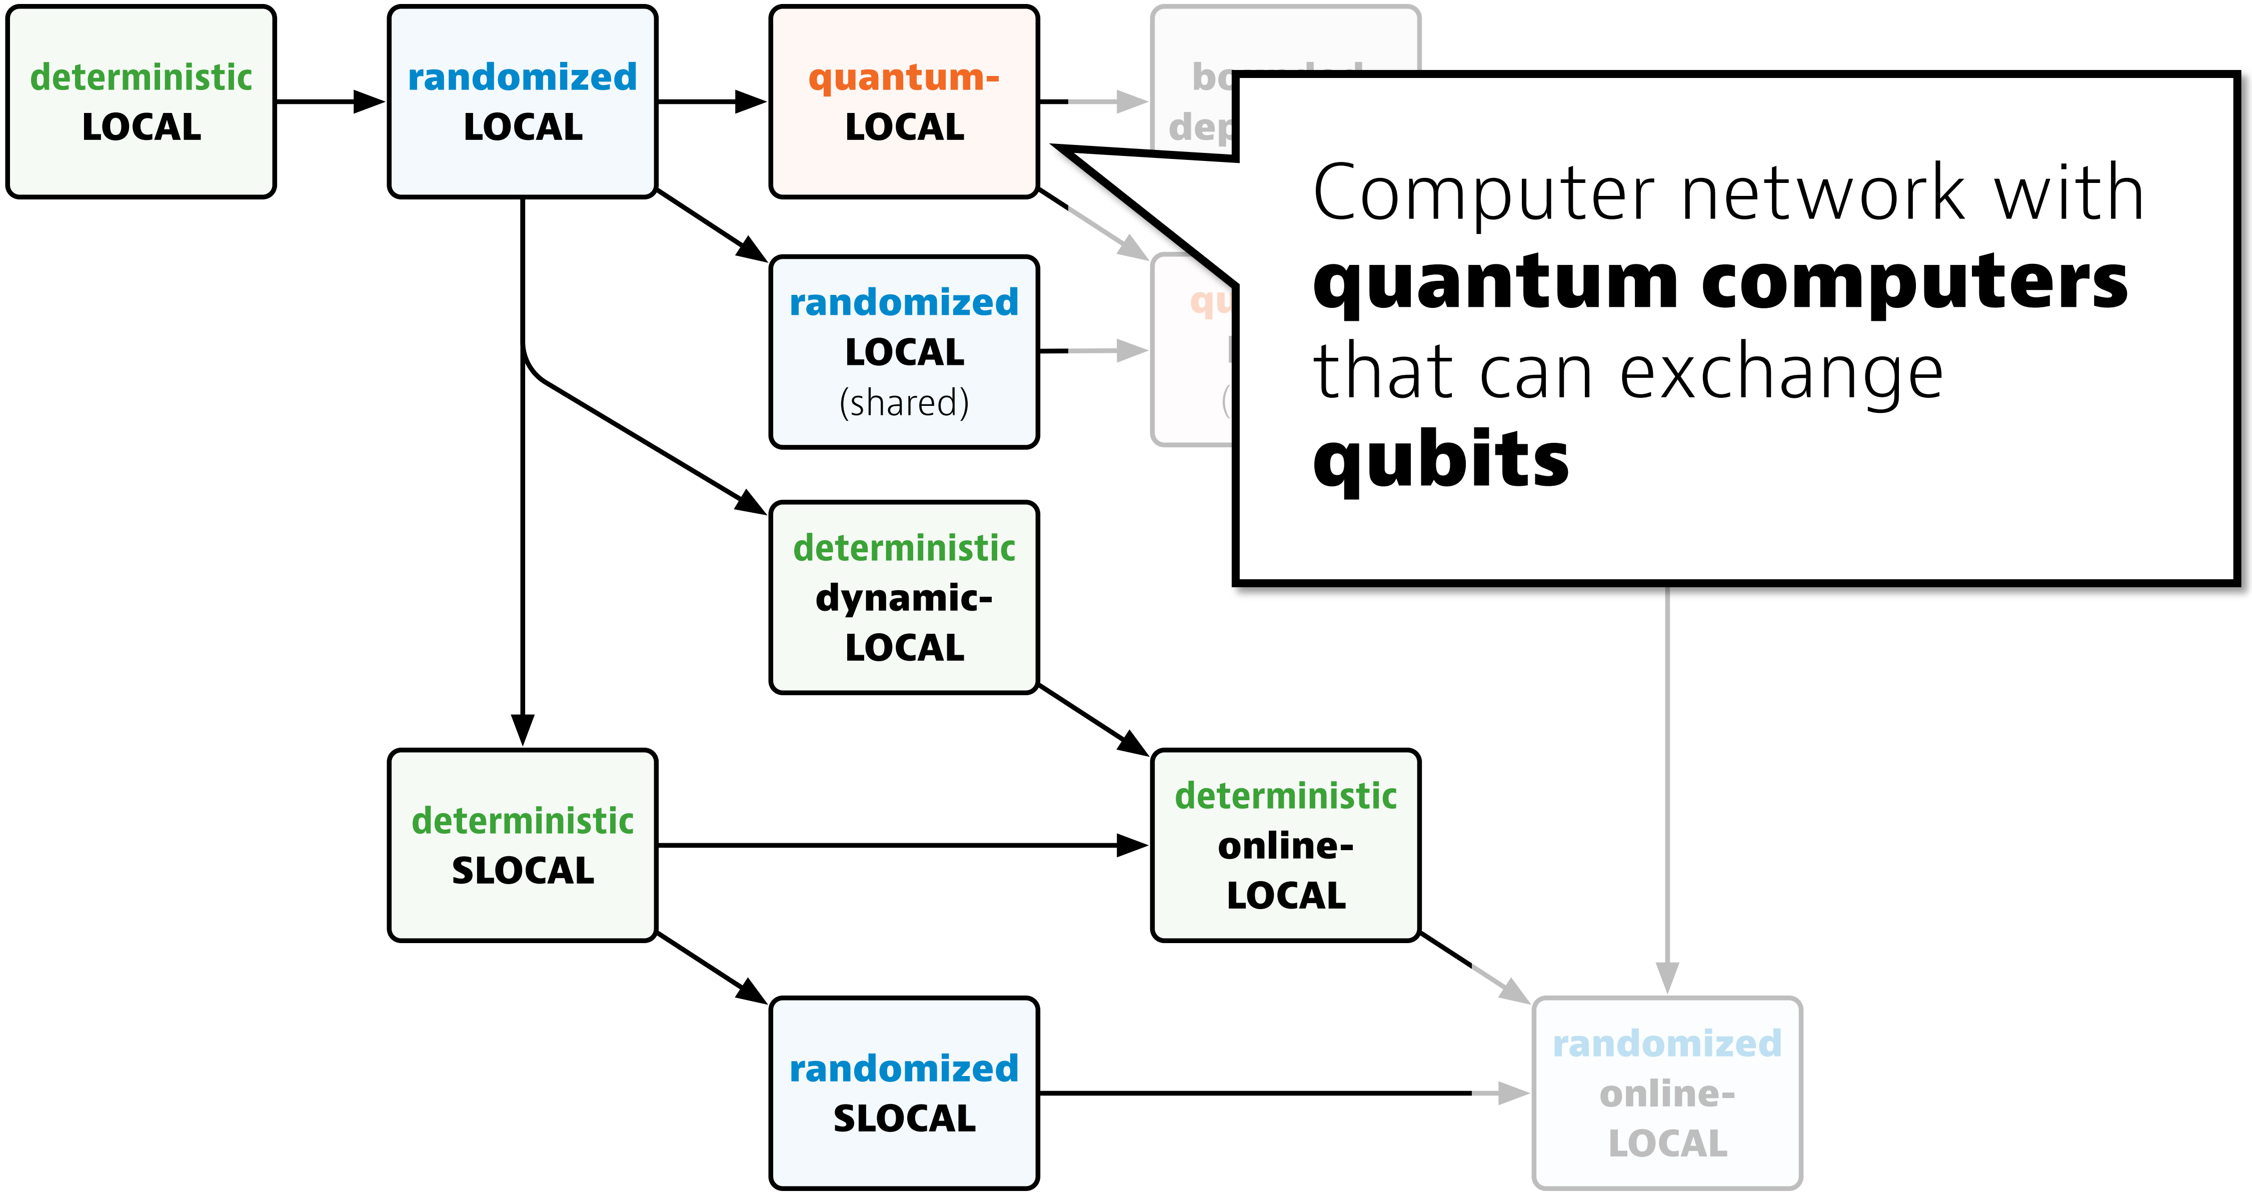
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
randomized LOCAL (shared)
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
Computer network with **quantum computers** that can exchange **qubits**

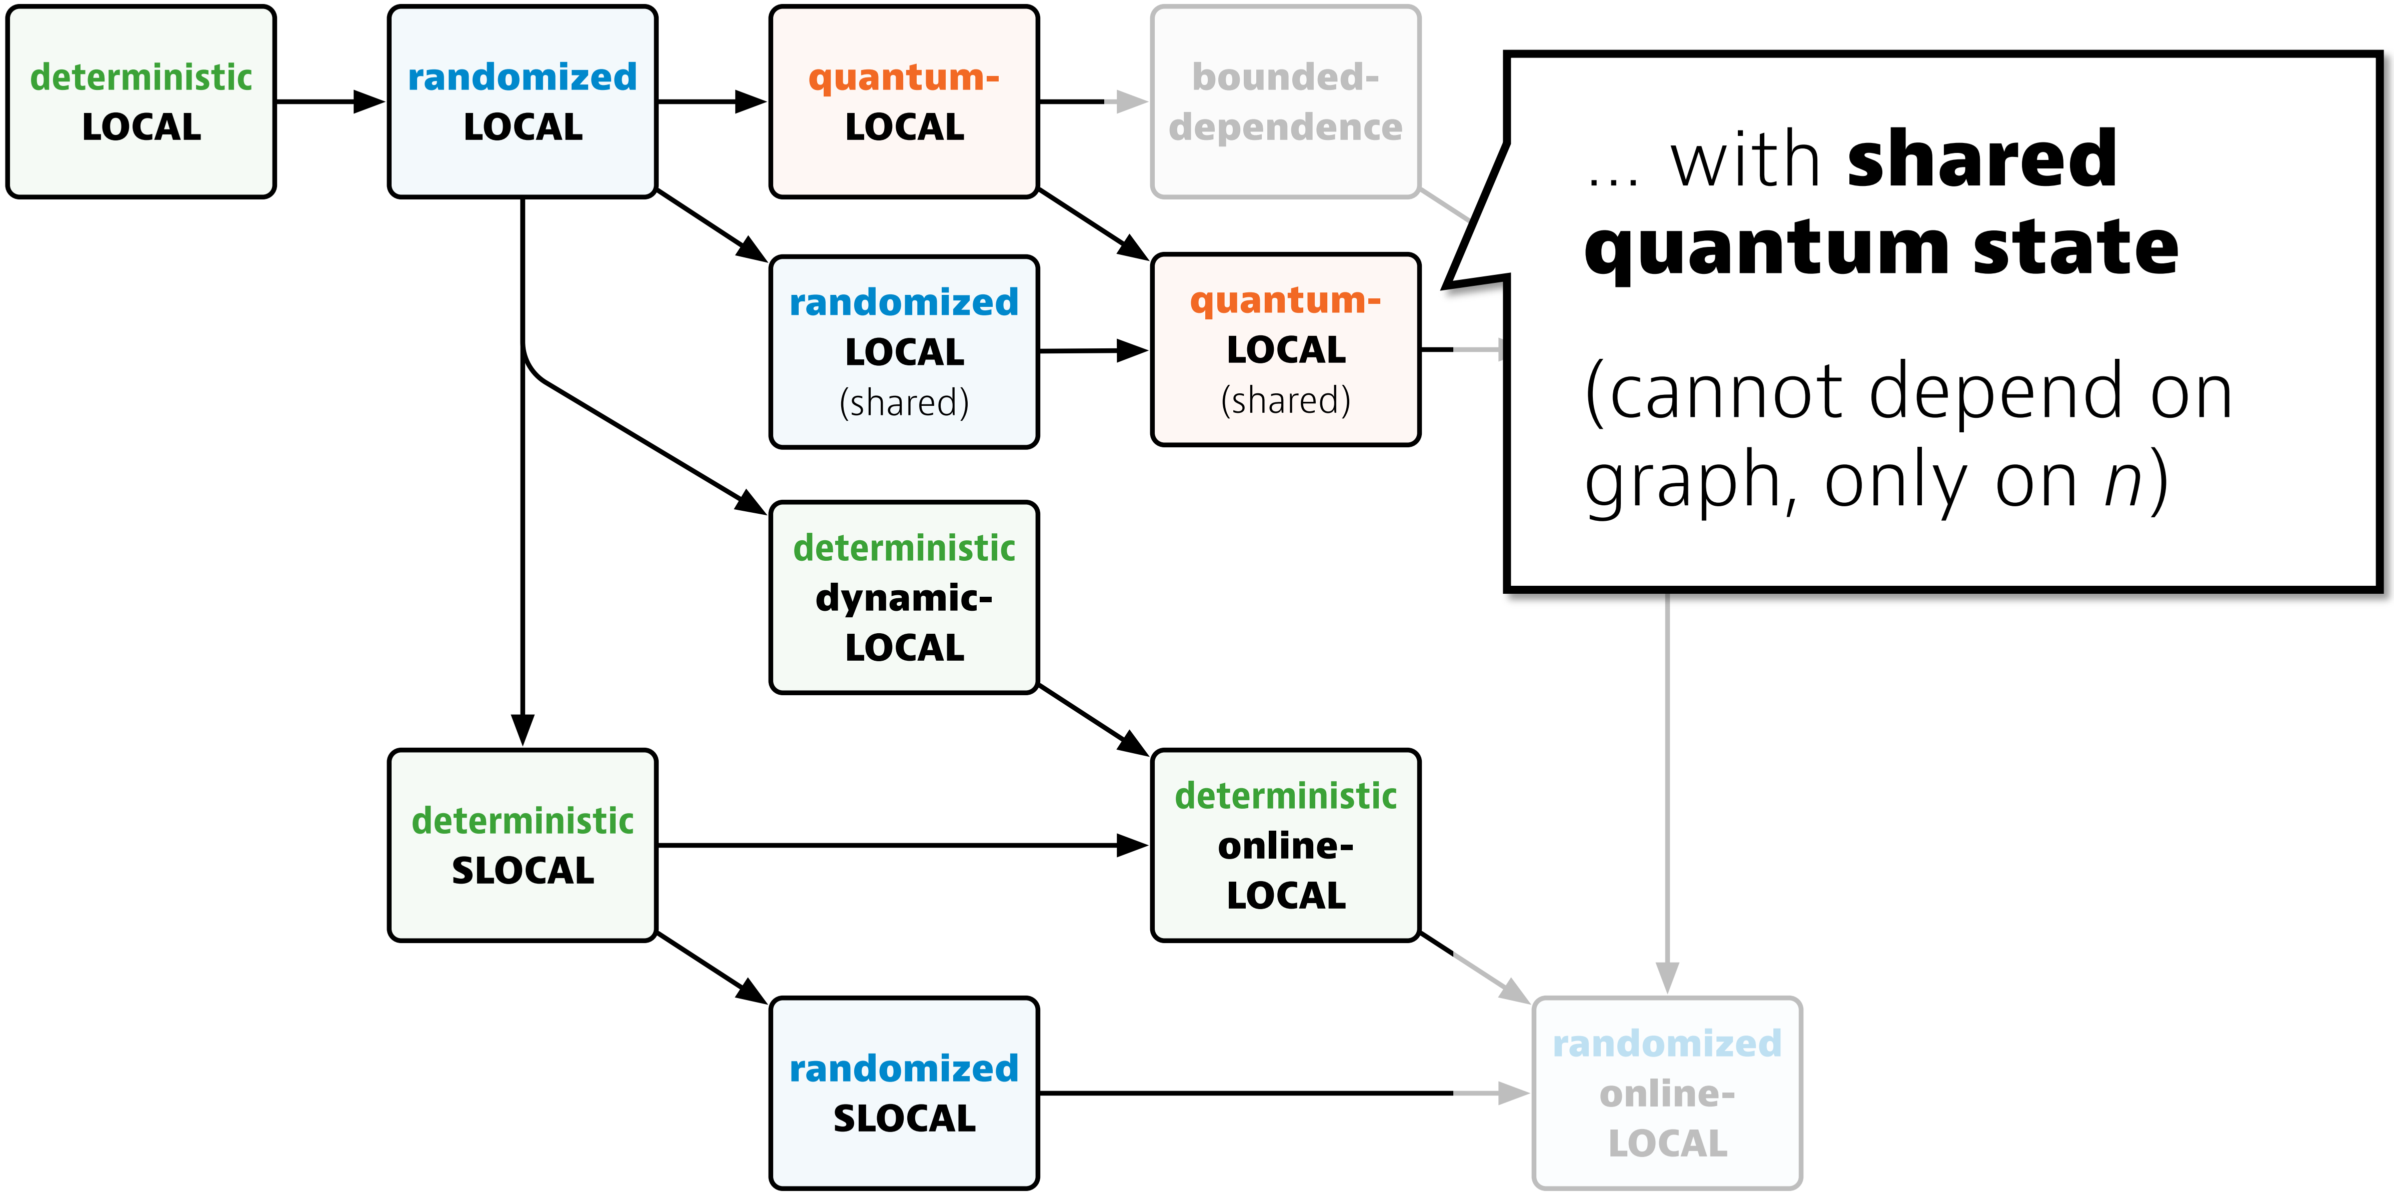
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
… with shared quantum state
(cannot depend on graph, only on n)

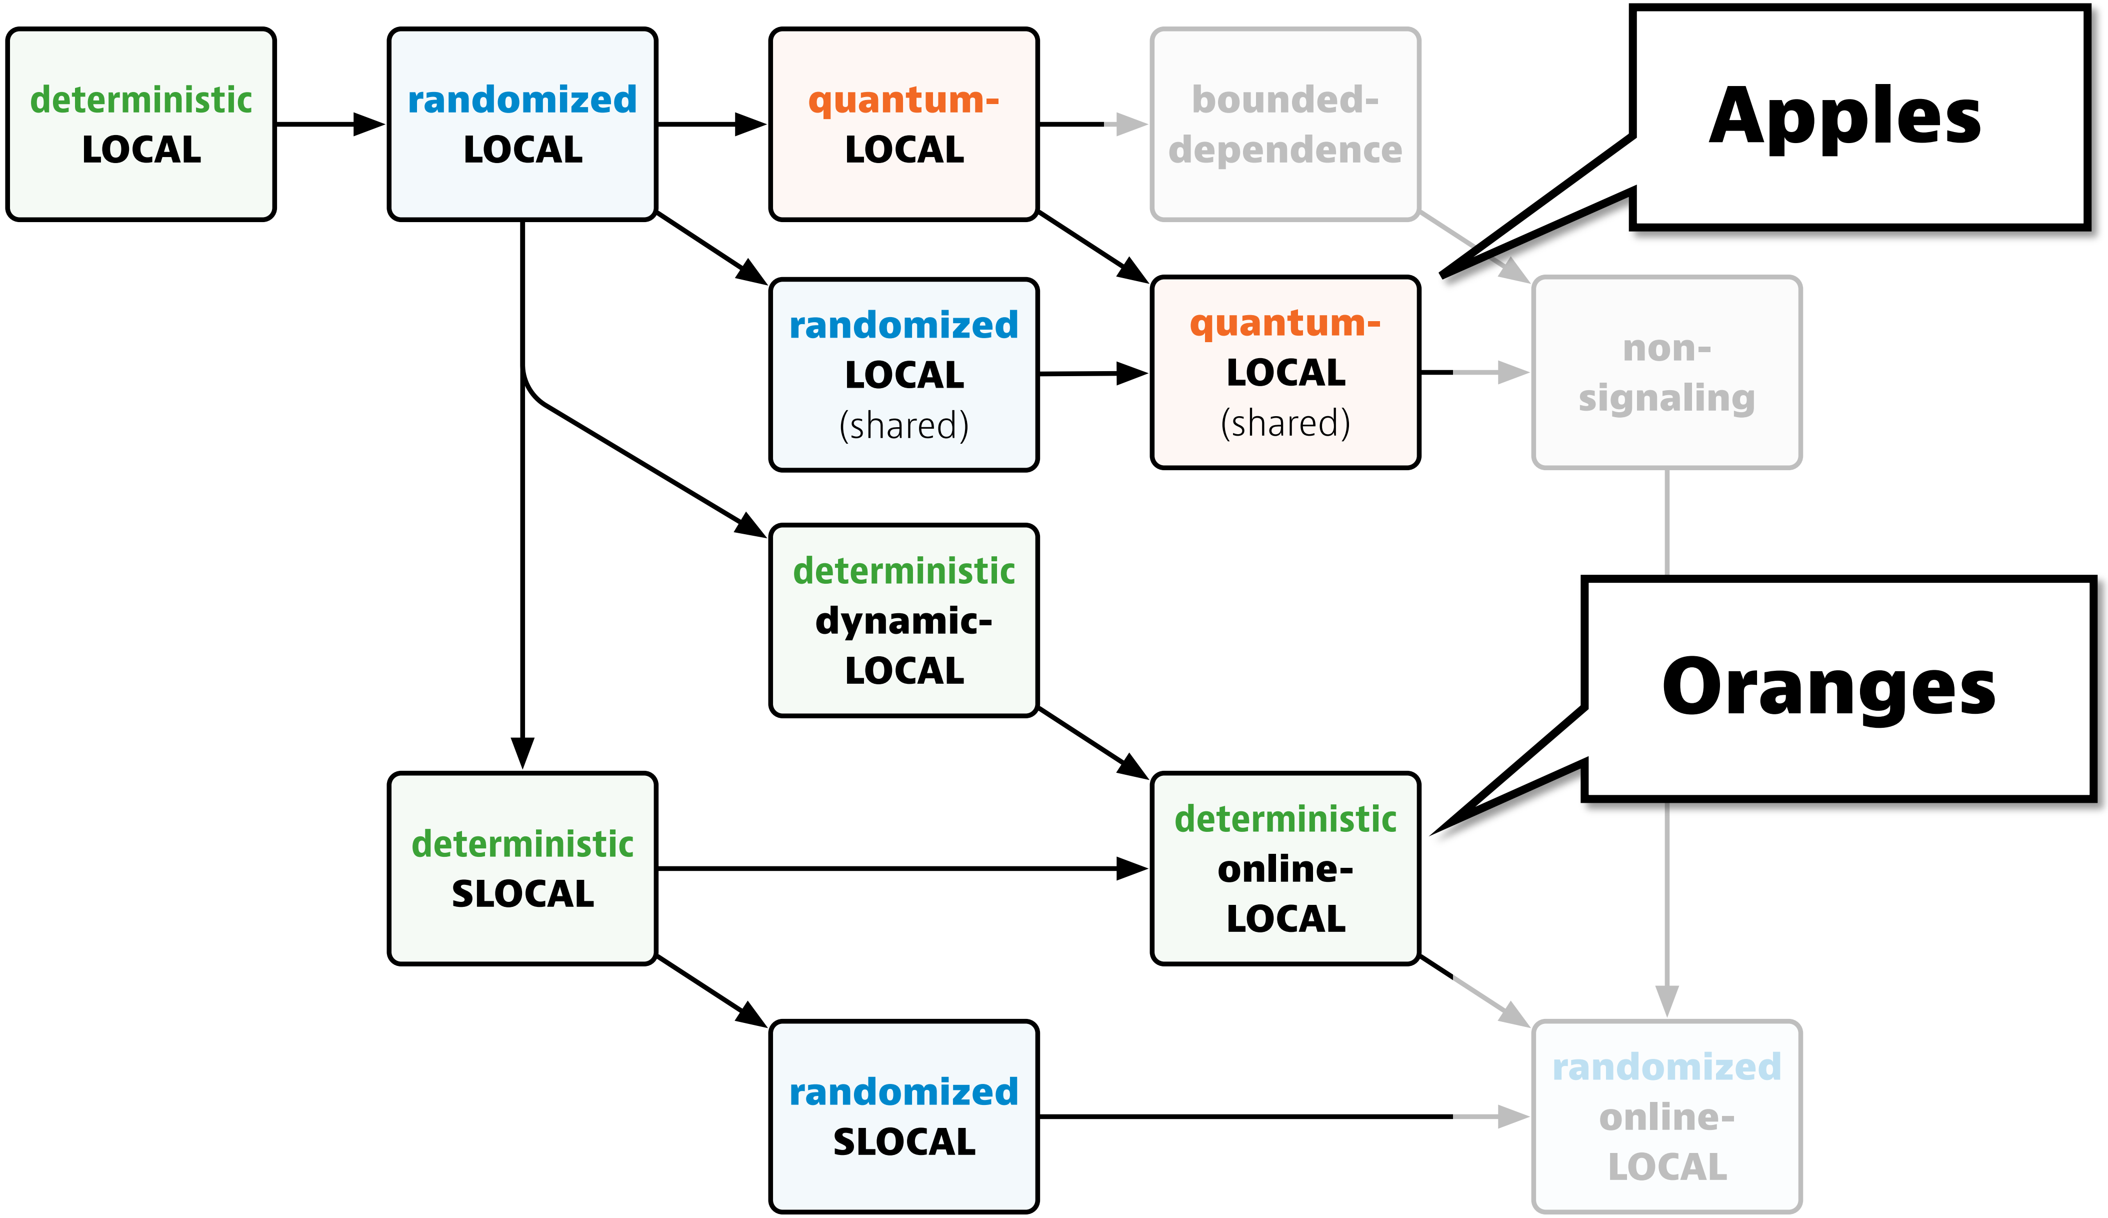

deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
Apples
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
Oranges
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

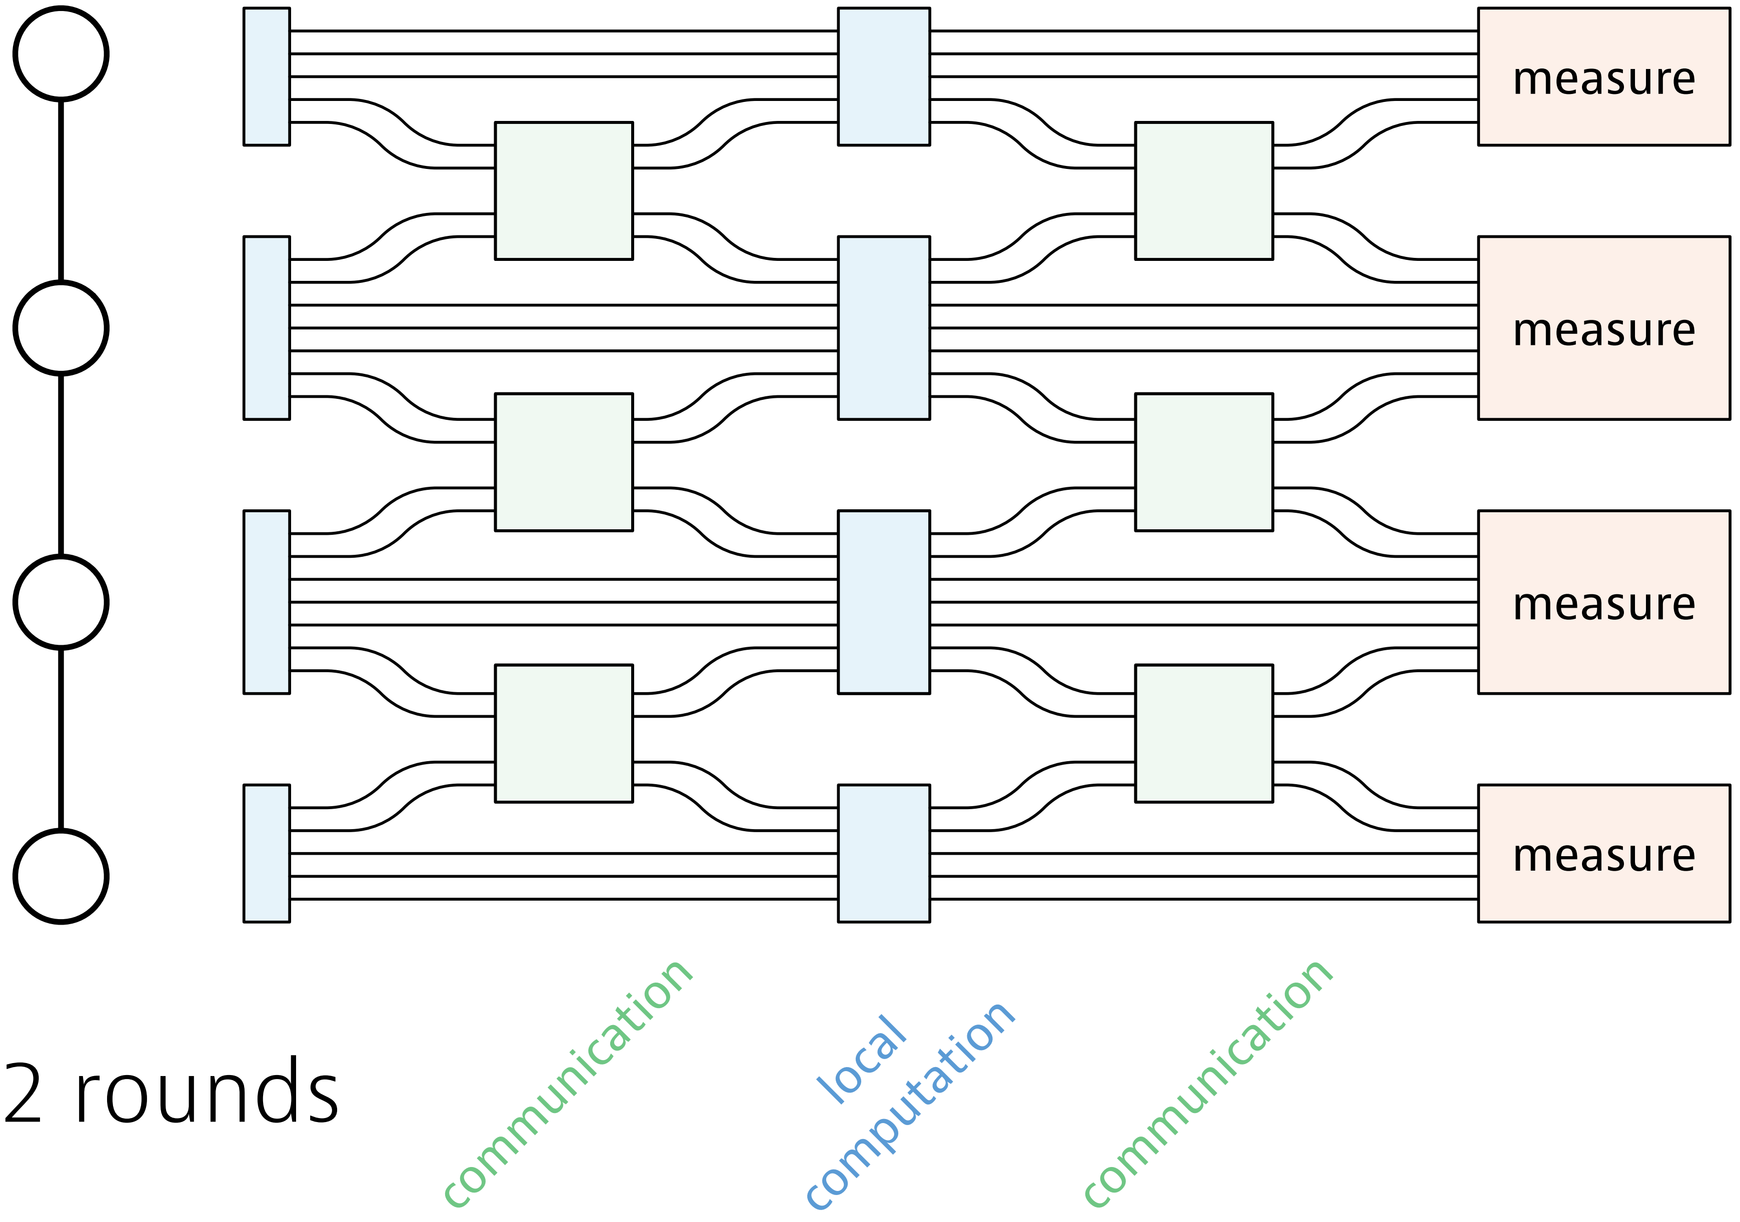
measure
measure
measure
measure
communication
local
computation
communication

2 rounds

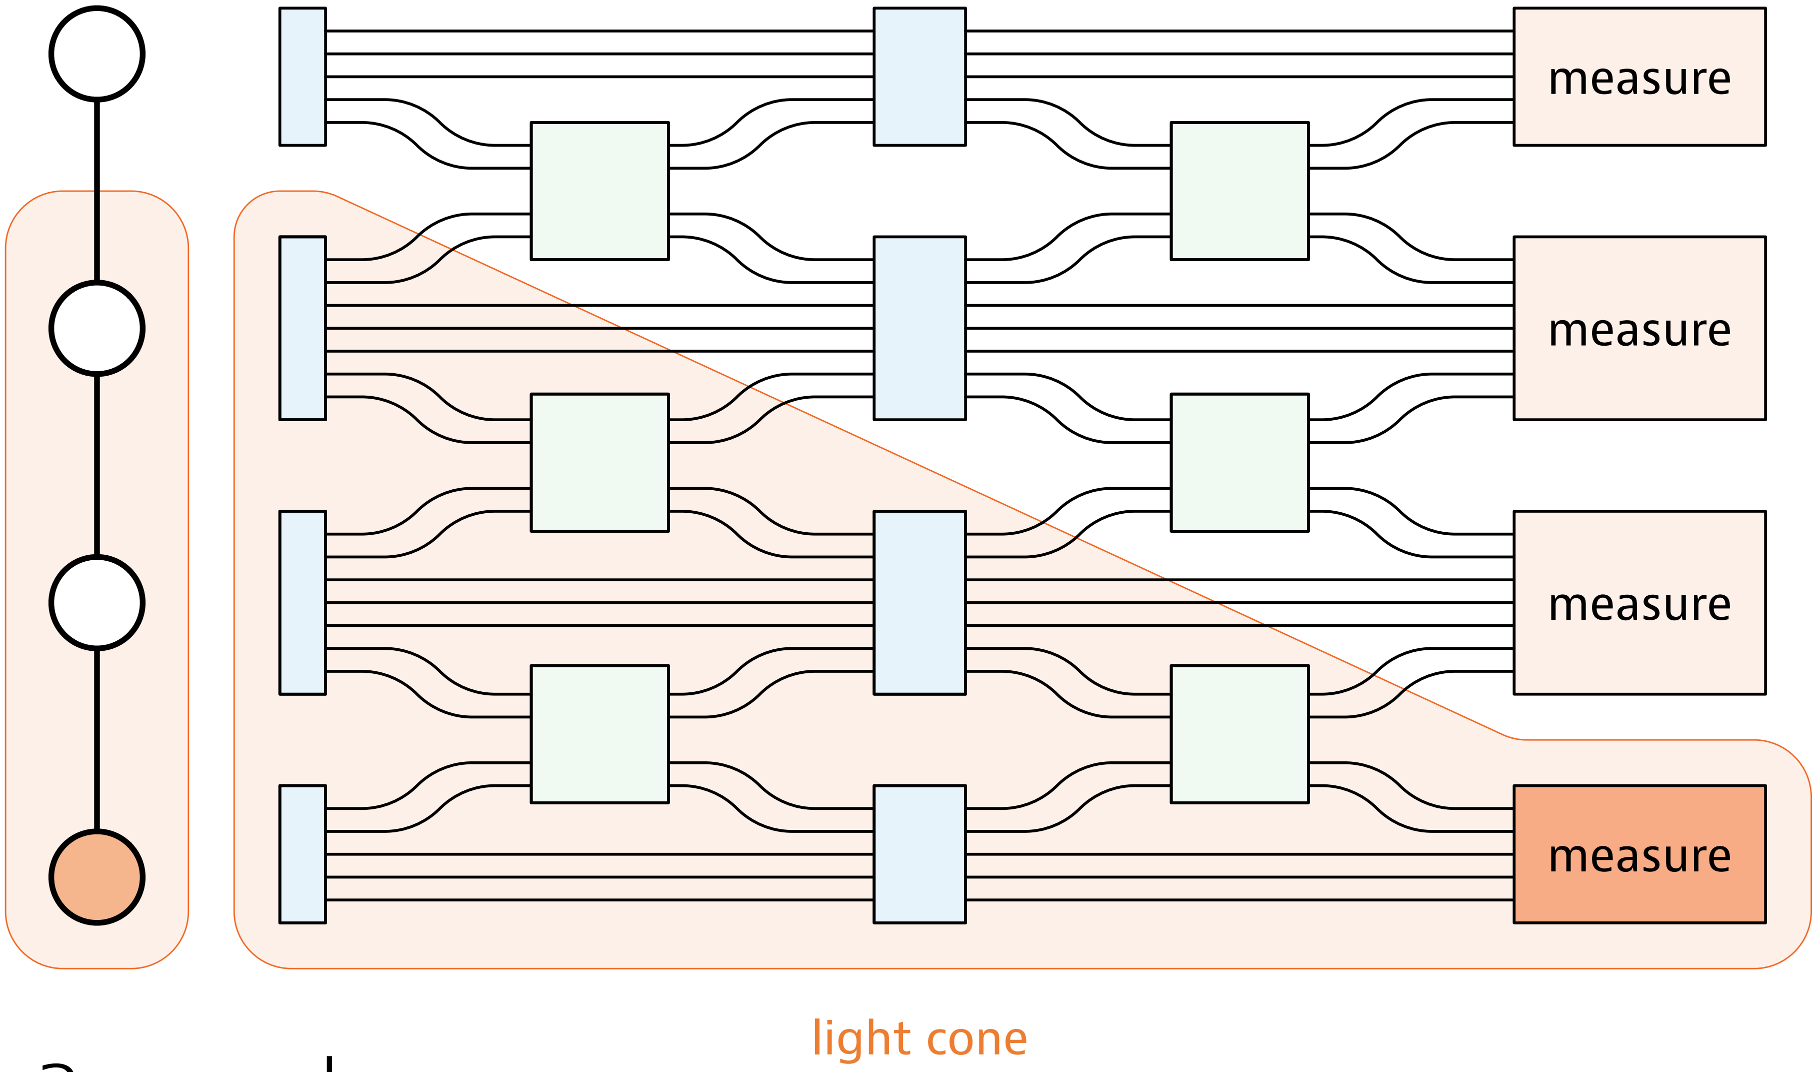

2 rounds

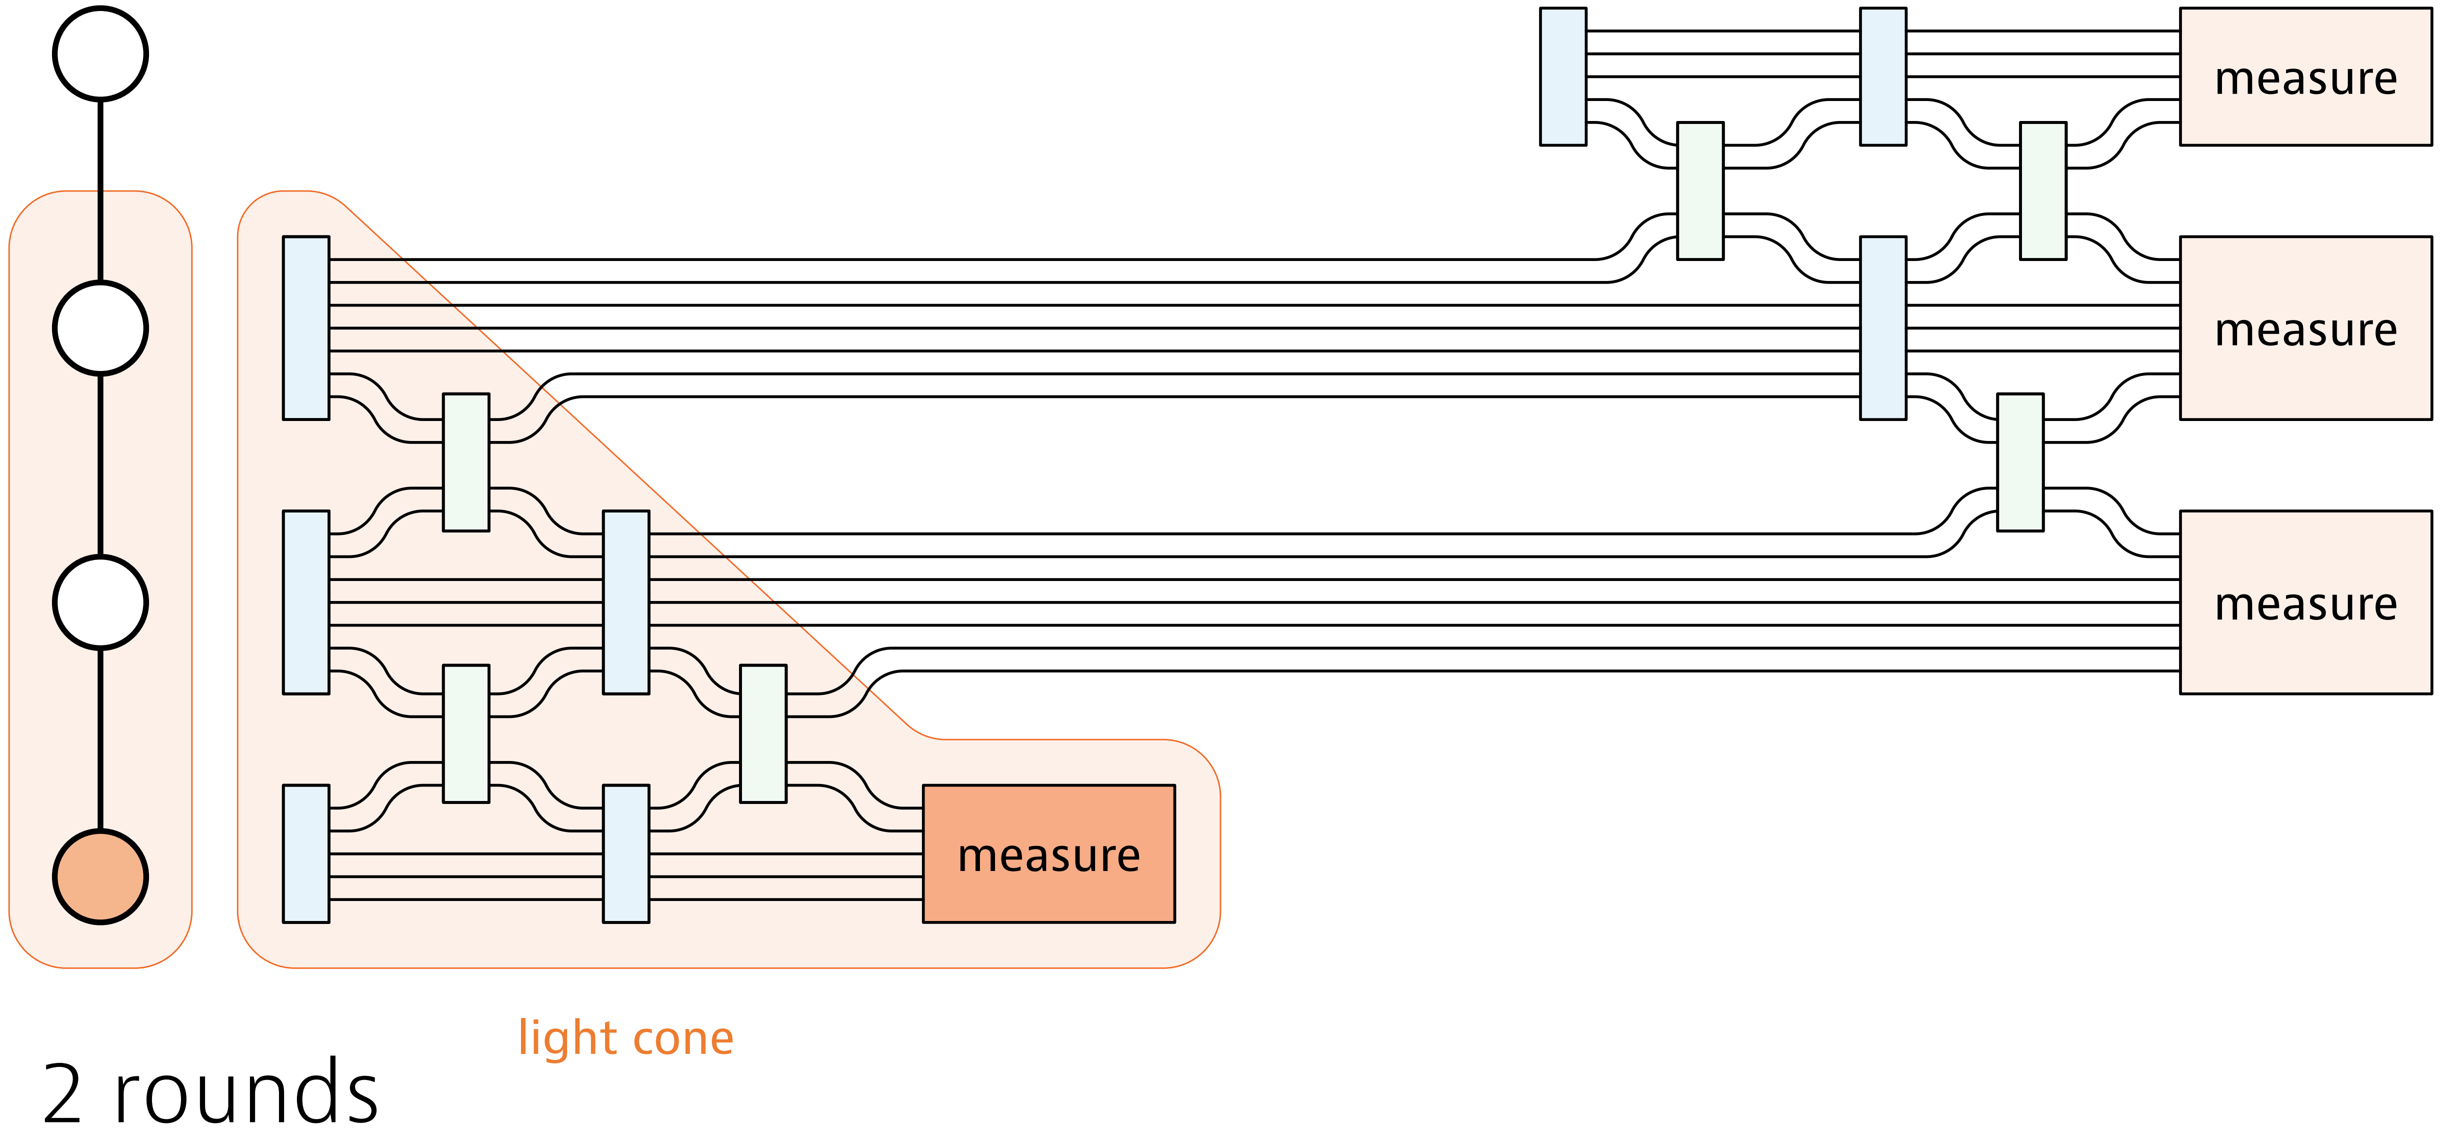
measure
measure
measure
measure
light cone
2 rounds

Non-signaling model

- Quantum LOCAL can't violate causality
- Key idea: **define** a model so that it can do **anything** except violating causality

Non-signaling model

Definition (*non-signaling distribution*):

- fix any **set of nodes *X*** …

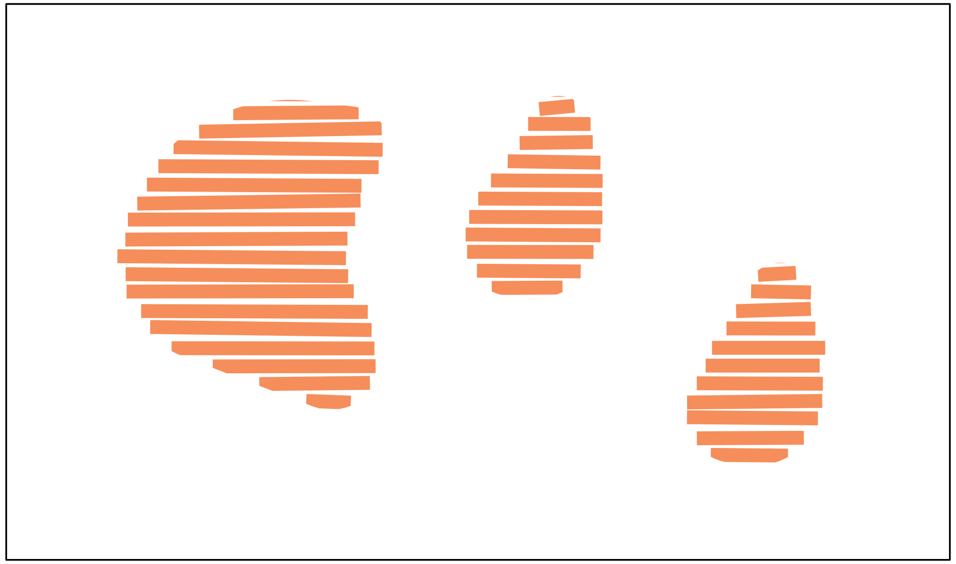

Gavoille, Kosowski, Markiewicz 2009
Arfaoui, Fraigniaud 2014

Non-signaling model

Definition (*non-signaling distribution*):

- fix any **set of nodes *X***
- changes in the input **more than *T* hops away** from ***X*** do not influence the output distribution of ***X***

Gavoille, Kosowski, Markiewicz 2009
Arfaoui, Fraigniaud 2014

Classical probability theory

Classical (randomized) distributed algorithms

↓

Quantum distributed algorithms

Weird quantum things

↓

Non-signaling "algorithms"

Classical probability theory

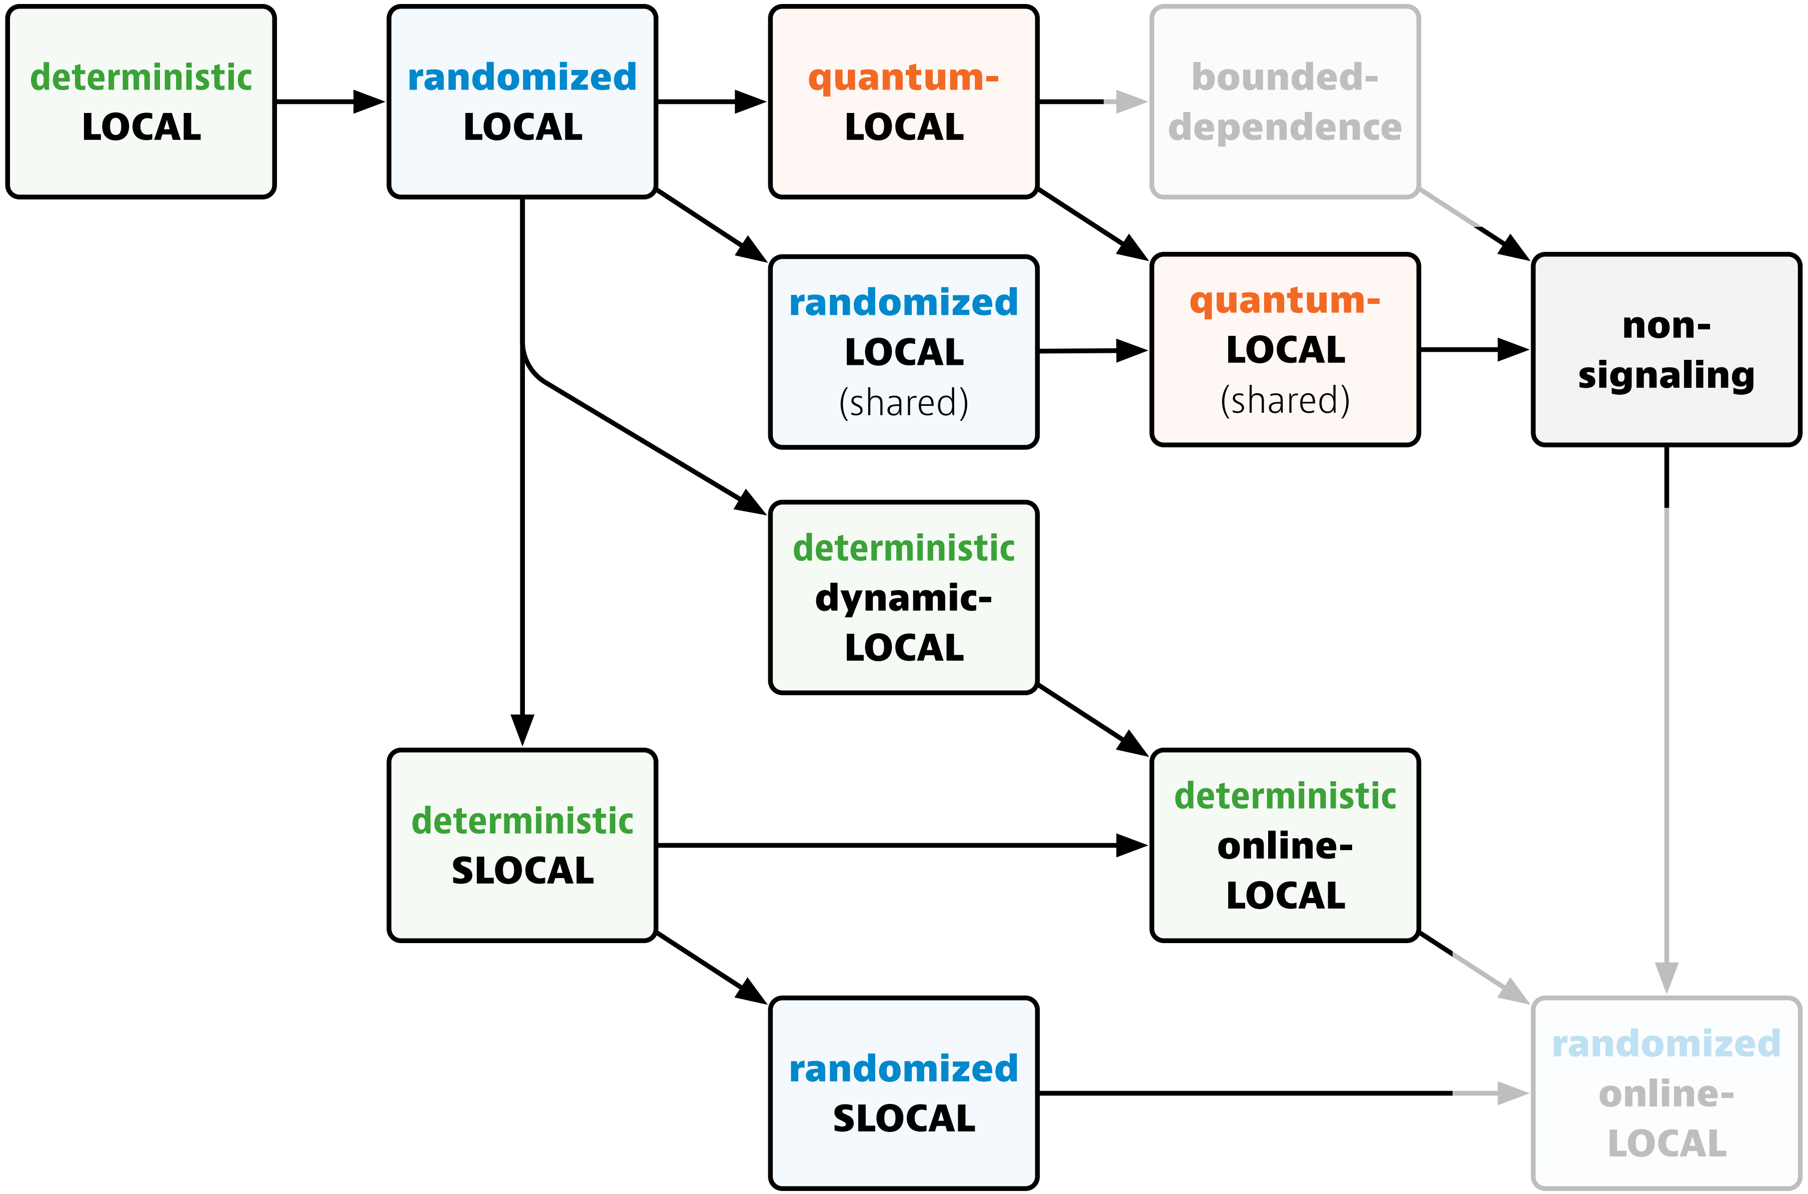
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

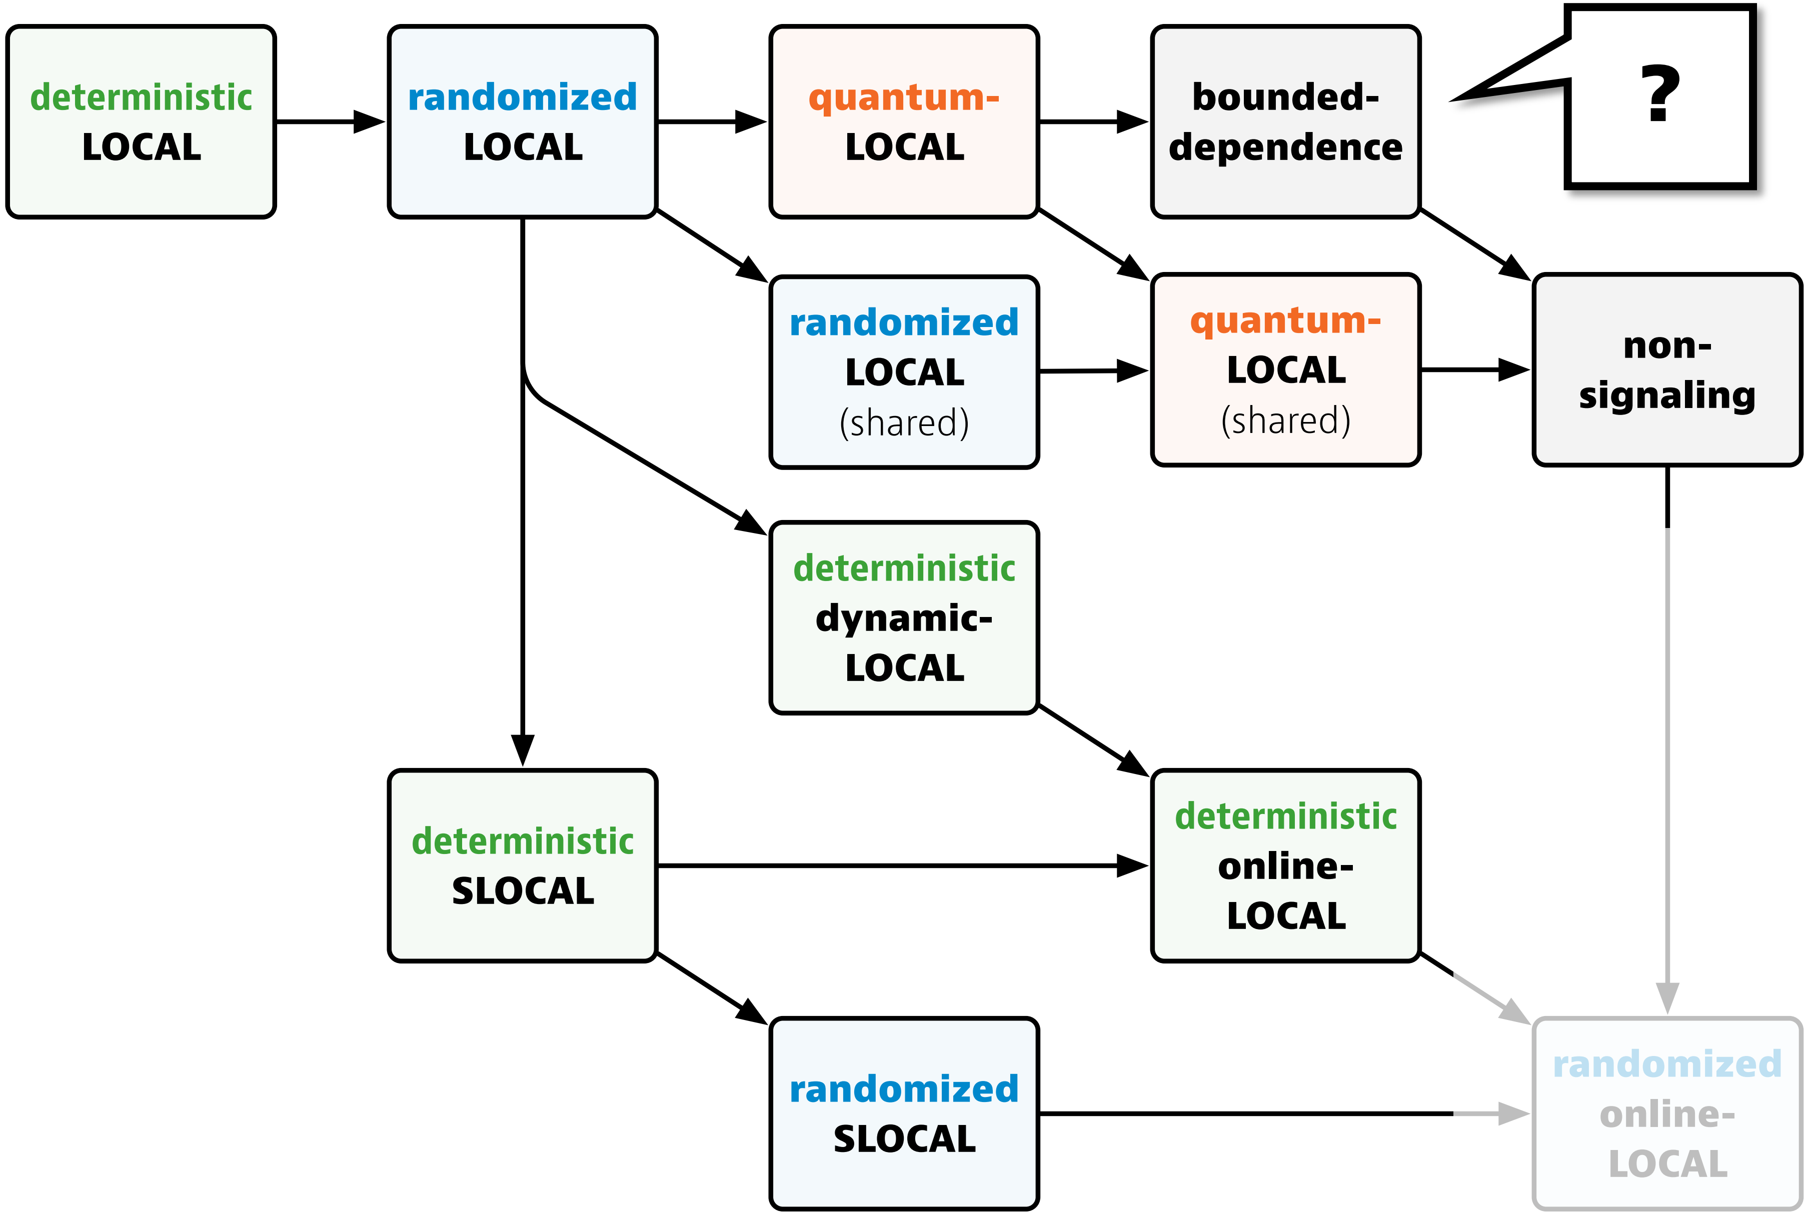
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
?
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

Bounded dependence

- **Finitely dependent distributions**
  - X and Y far from each others → independent
  - usually "far" = some constant
- For clarity, we call it here **bounded dependence model** when "far" = $T(n)$

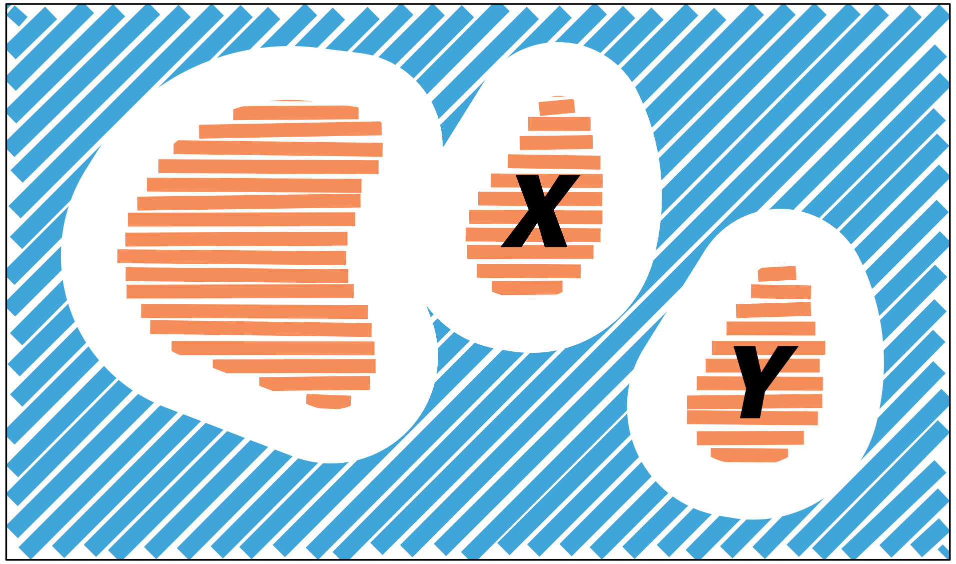

Bounded dependence

Quantum-LOCAL without shared quantum state:

Output of $T(n)$-round algorithm

↓

Not just non-signaling but also bounded dependence for distance ≈ $T(n)$

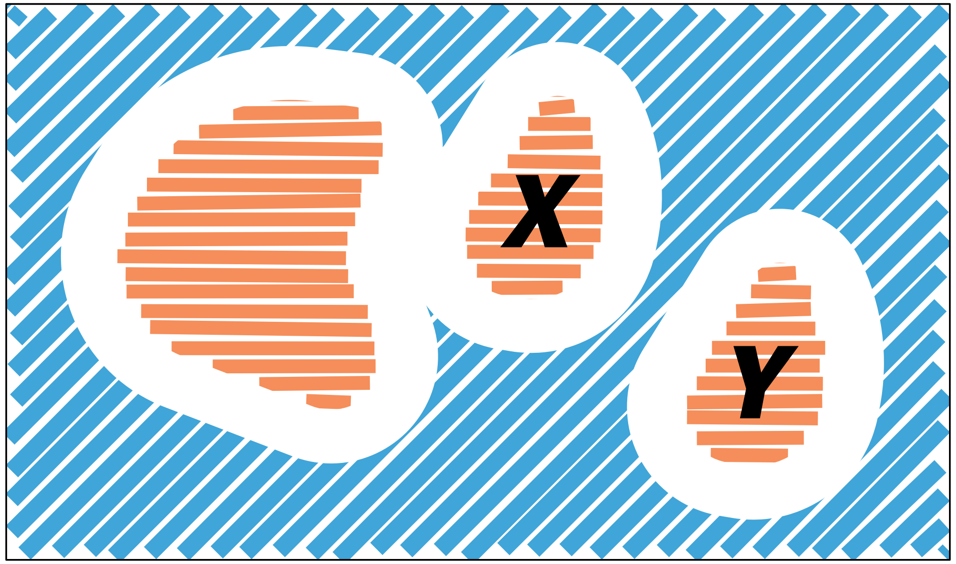

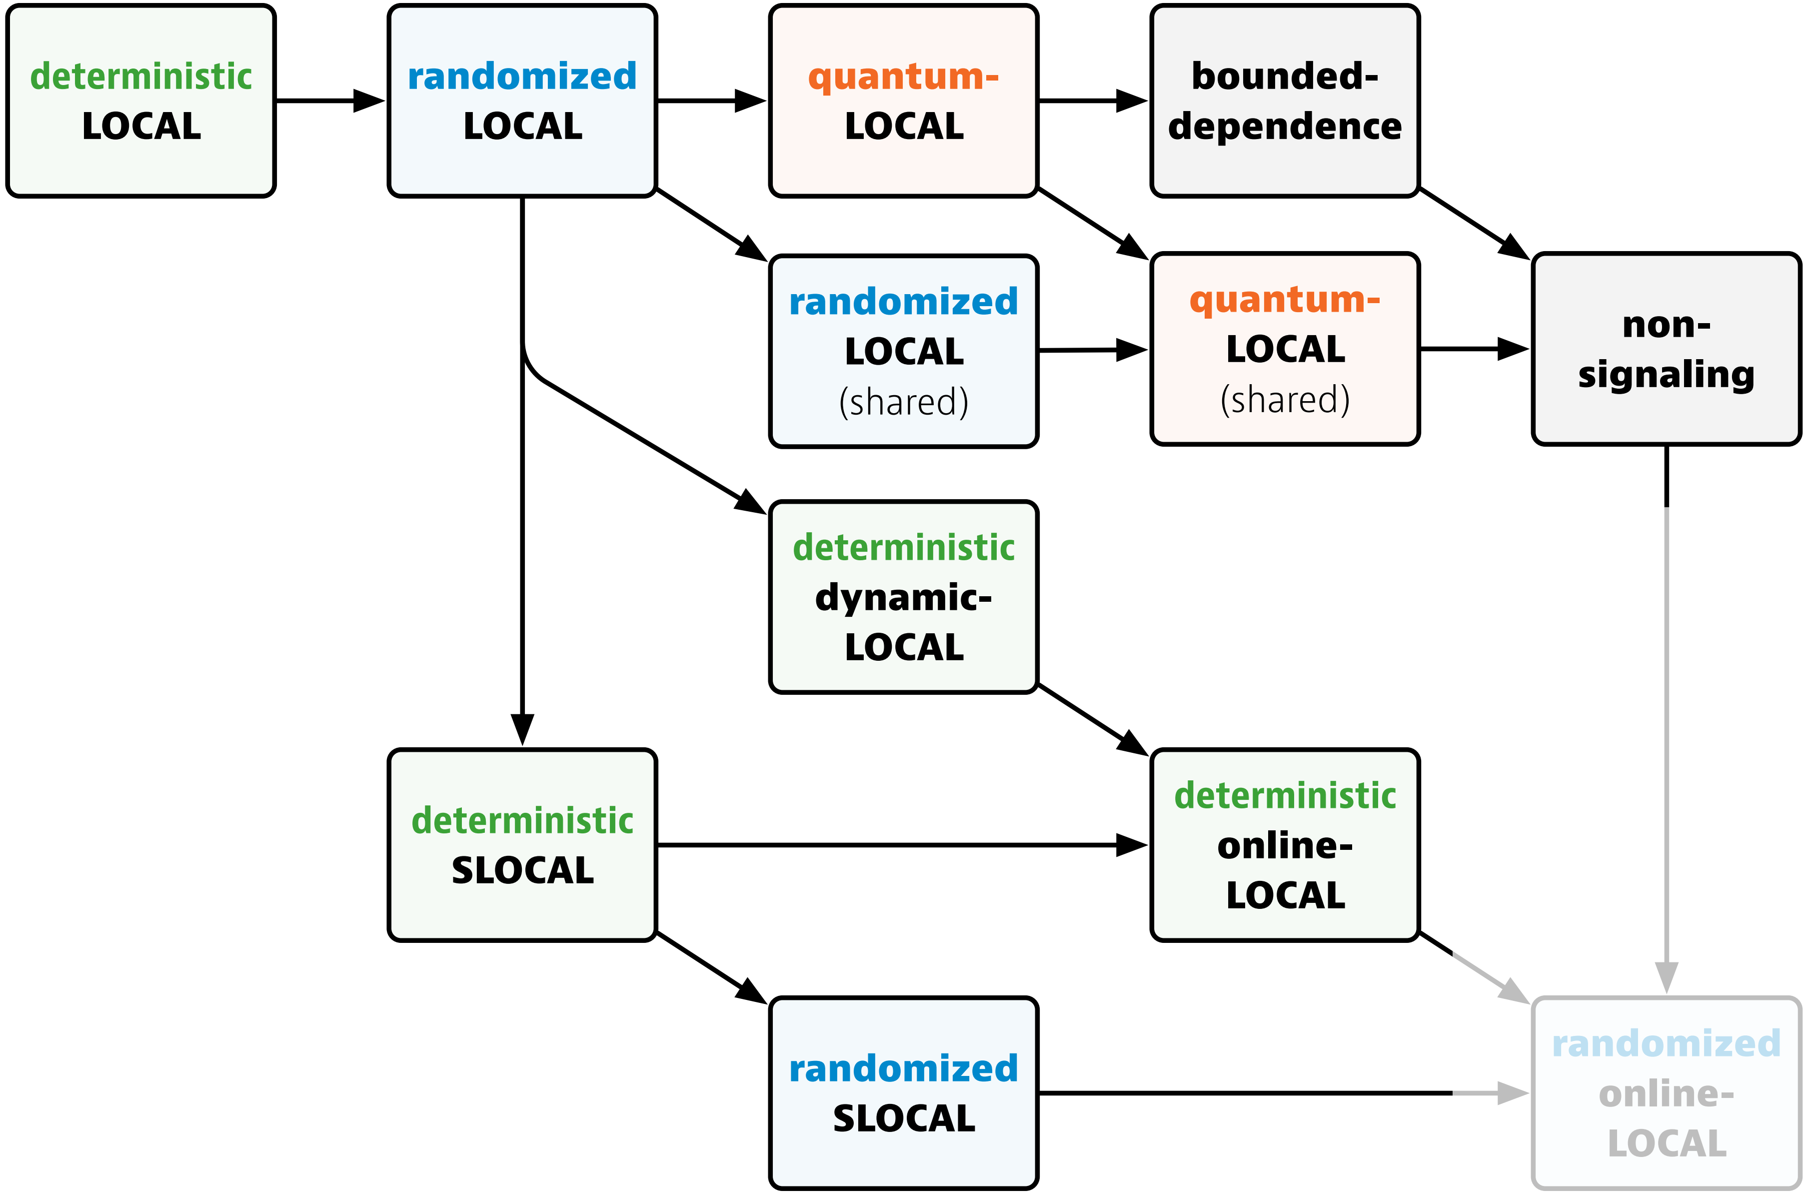
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

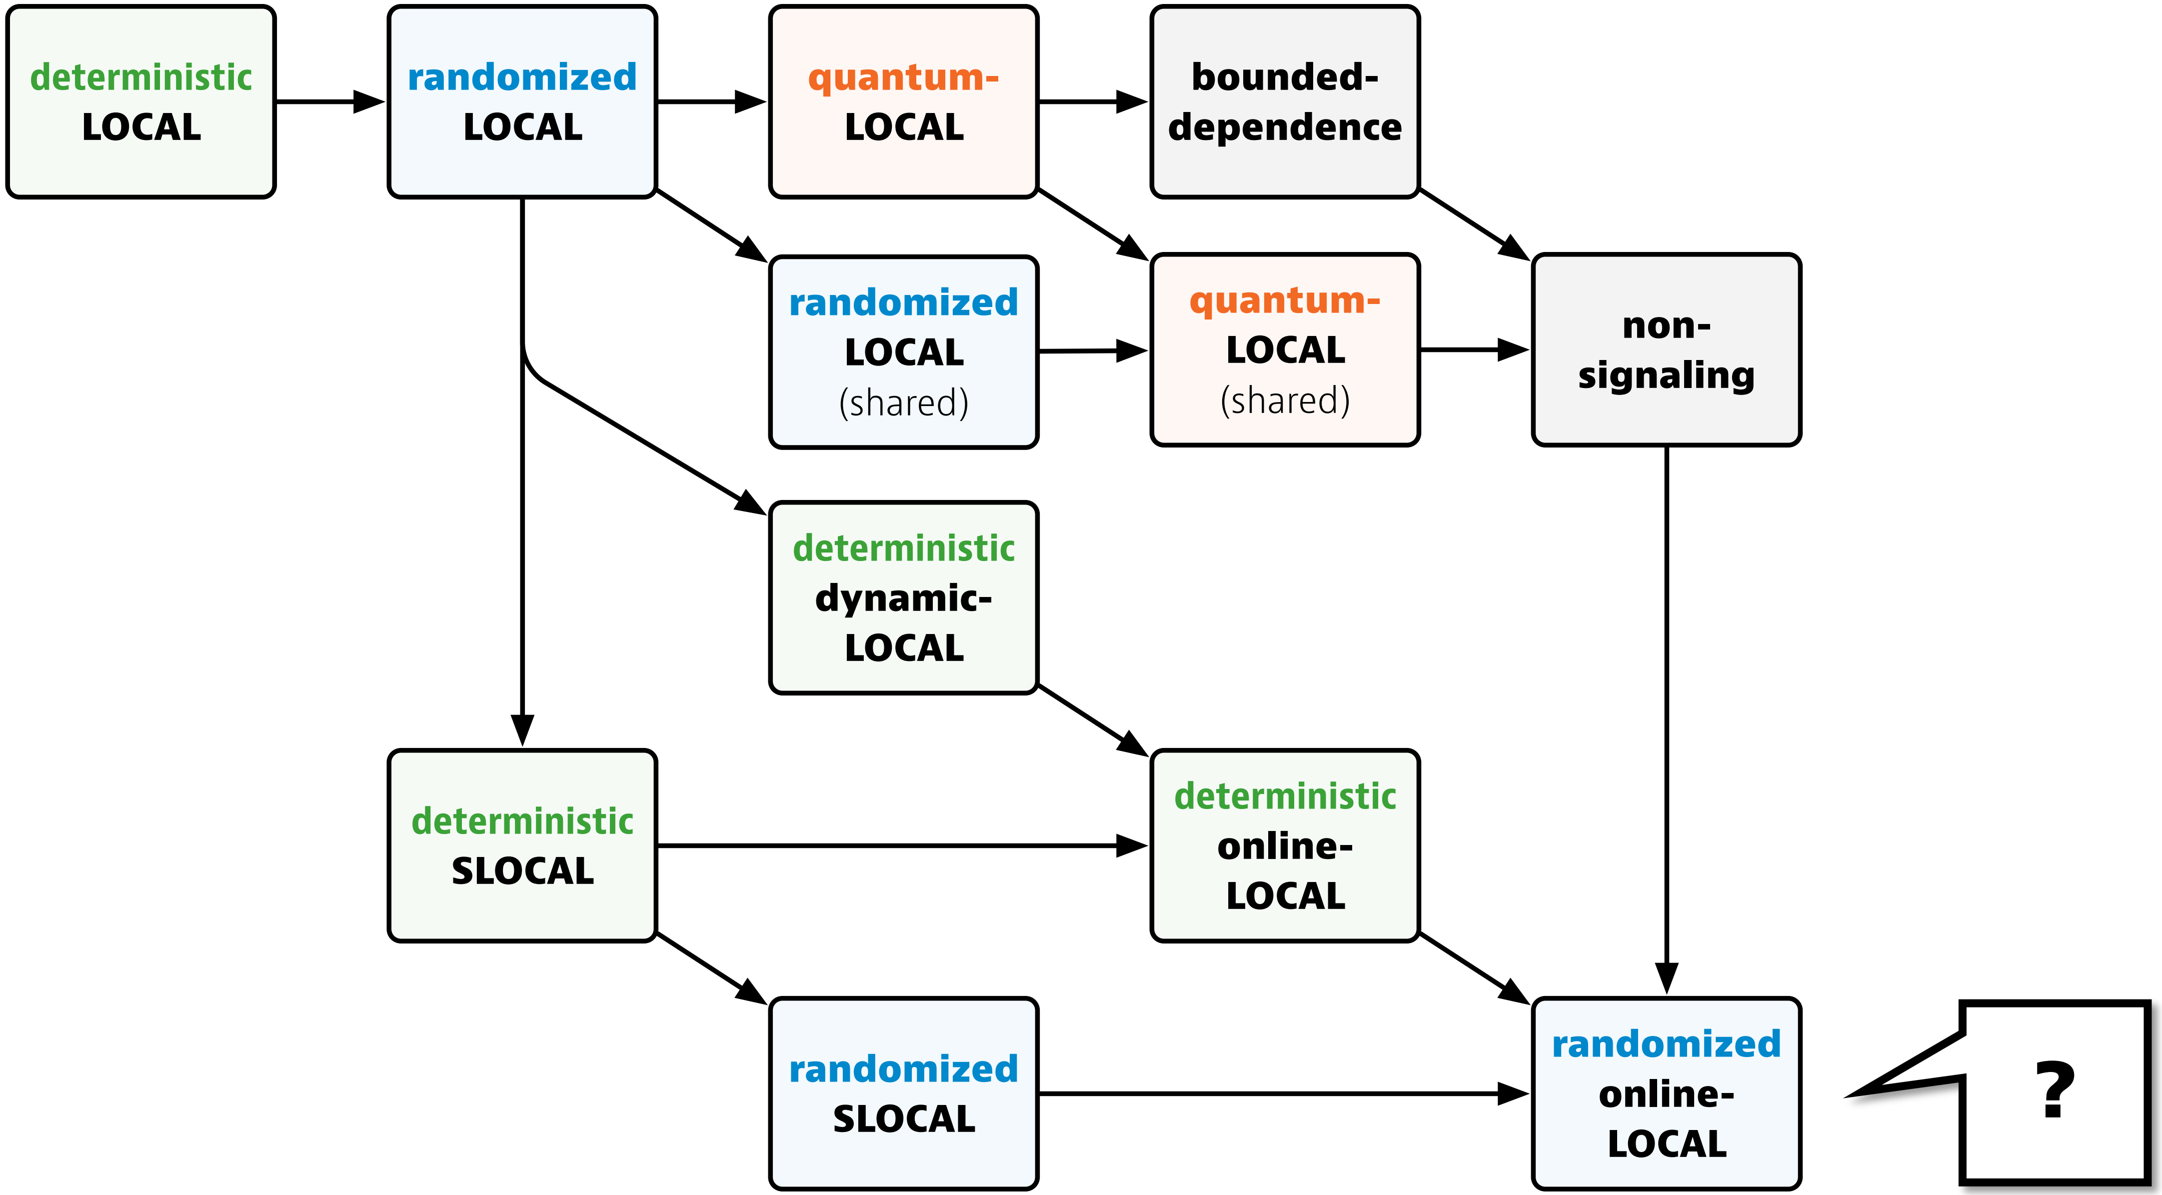
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL
?

Rand. online-LOCAL

- Adversary fixes a **graph** + **order** in which nodes are revealed

Oblivious adversary

- For each node v
  - algorithm sees radius-T neighborhood of v
  - algorithm must choose the label of v
- Algorithm can **remember** everything, algorithm can use **randomness**

Rand. online-LOCAL

- **Trivial:**
  - randomize online-LOCAL can simulate deterministic online-LOCAL
- **Suprise:**
  - randomized online-LOCAL ***can simulate any non-signaling distribution*** (with the same asymptotic locality)

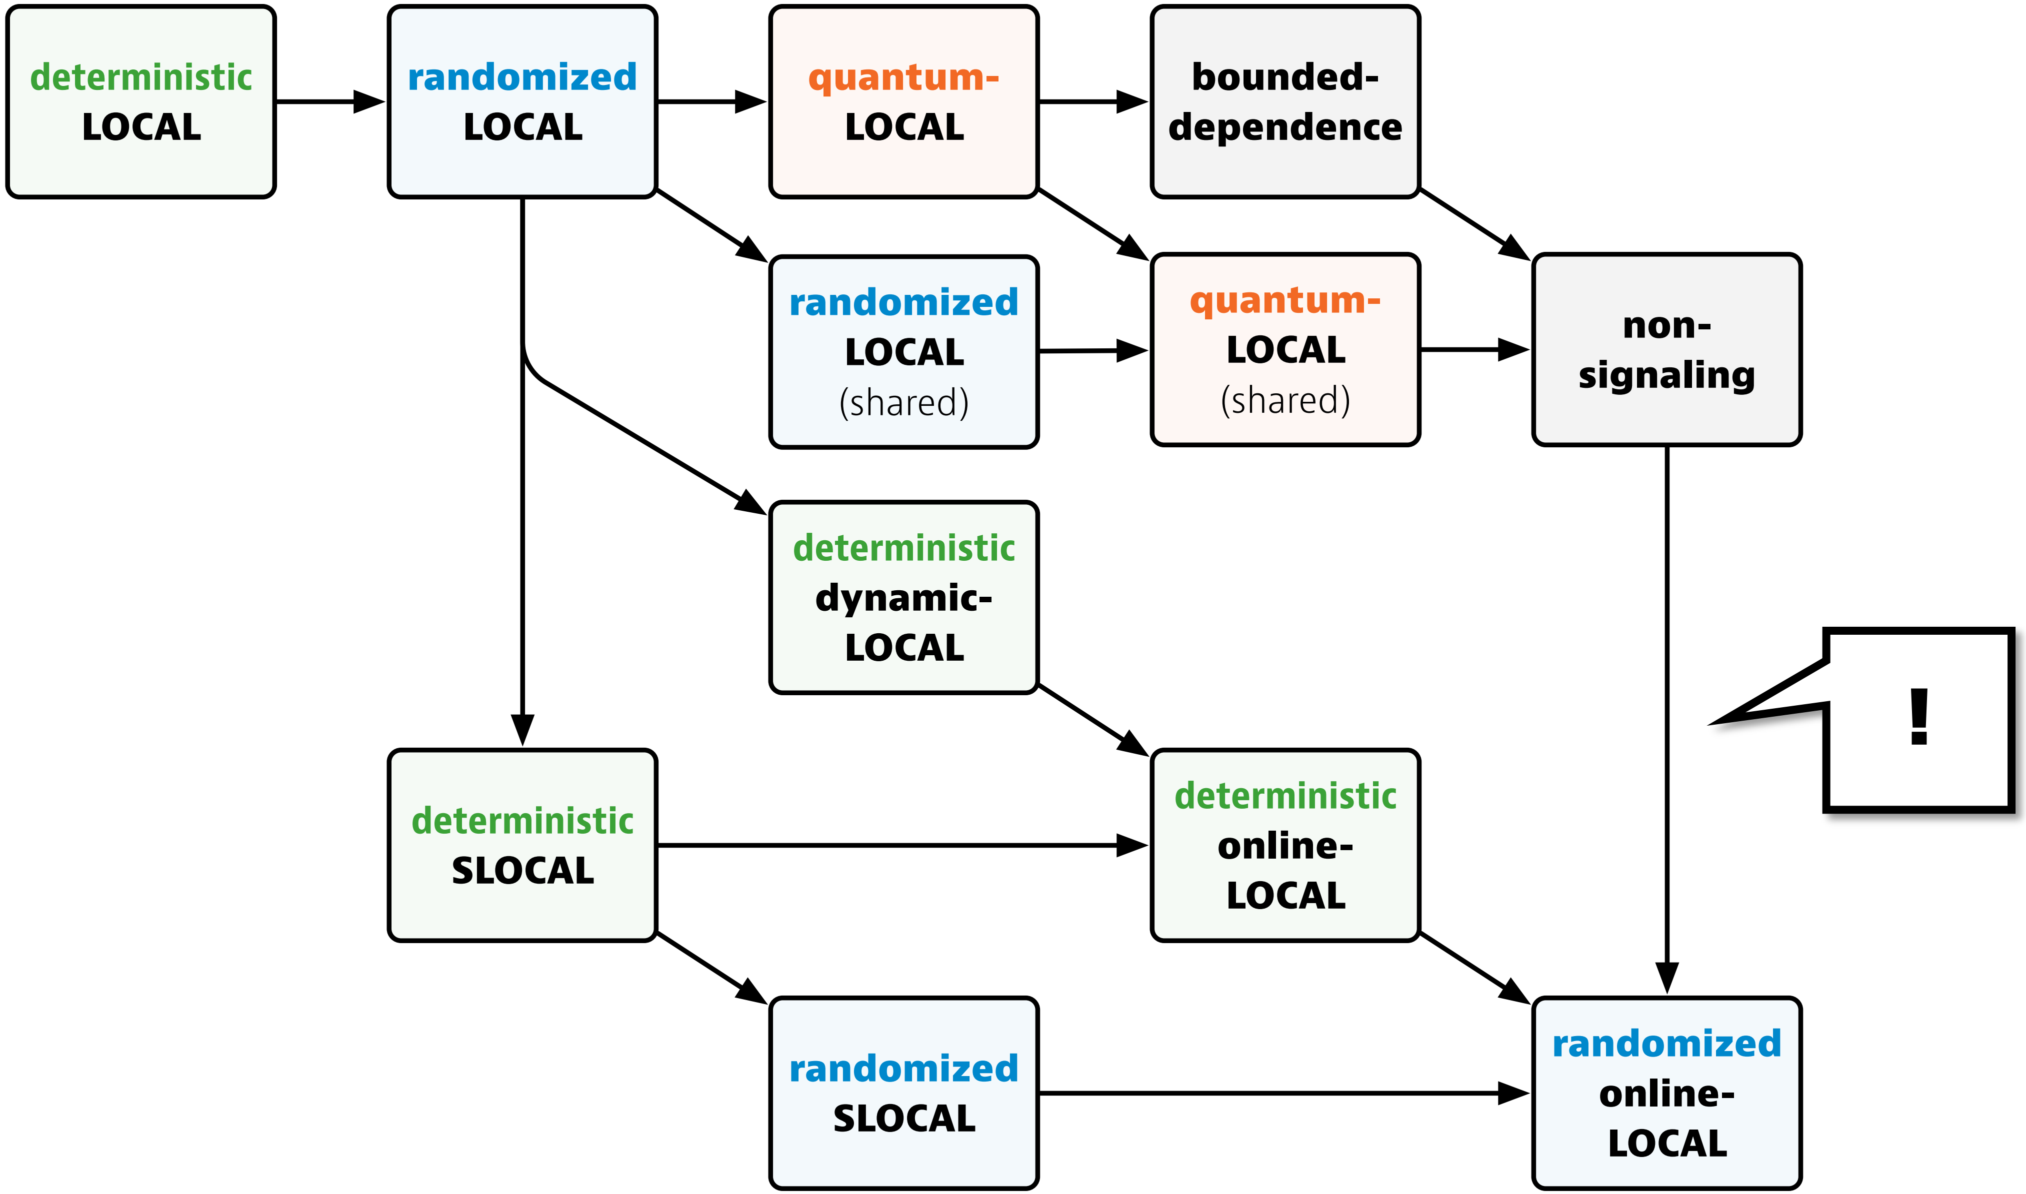
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
!
randomized SLOCAL
randomized online-LOCAL

Some separations

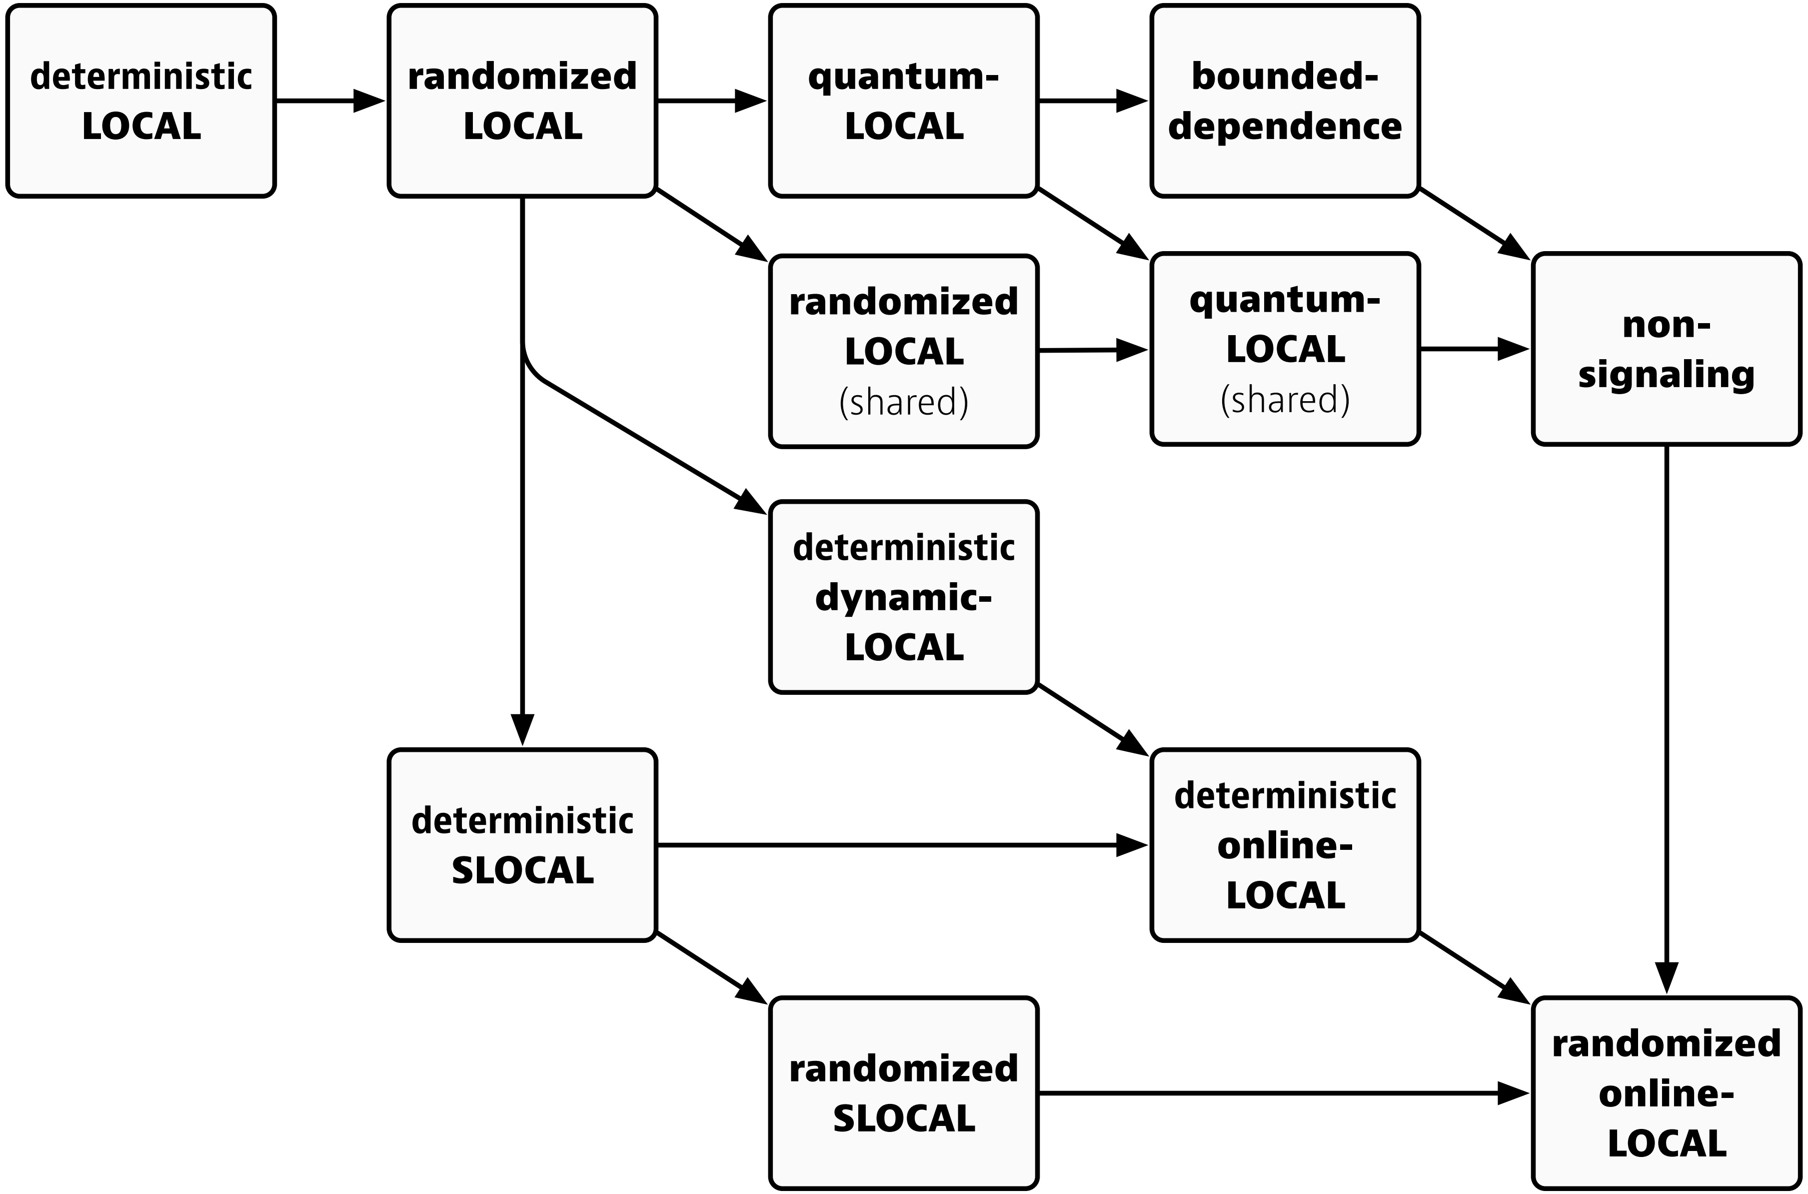
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

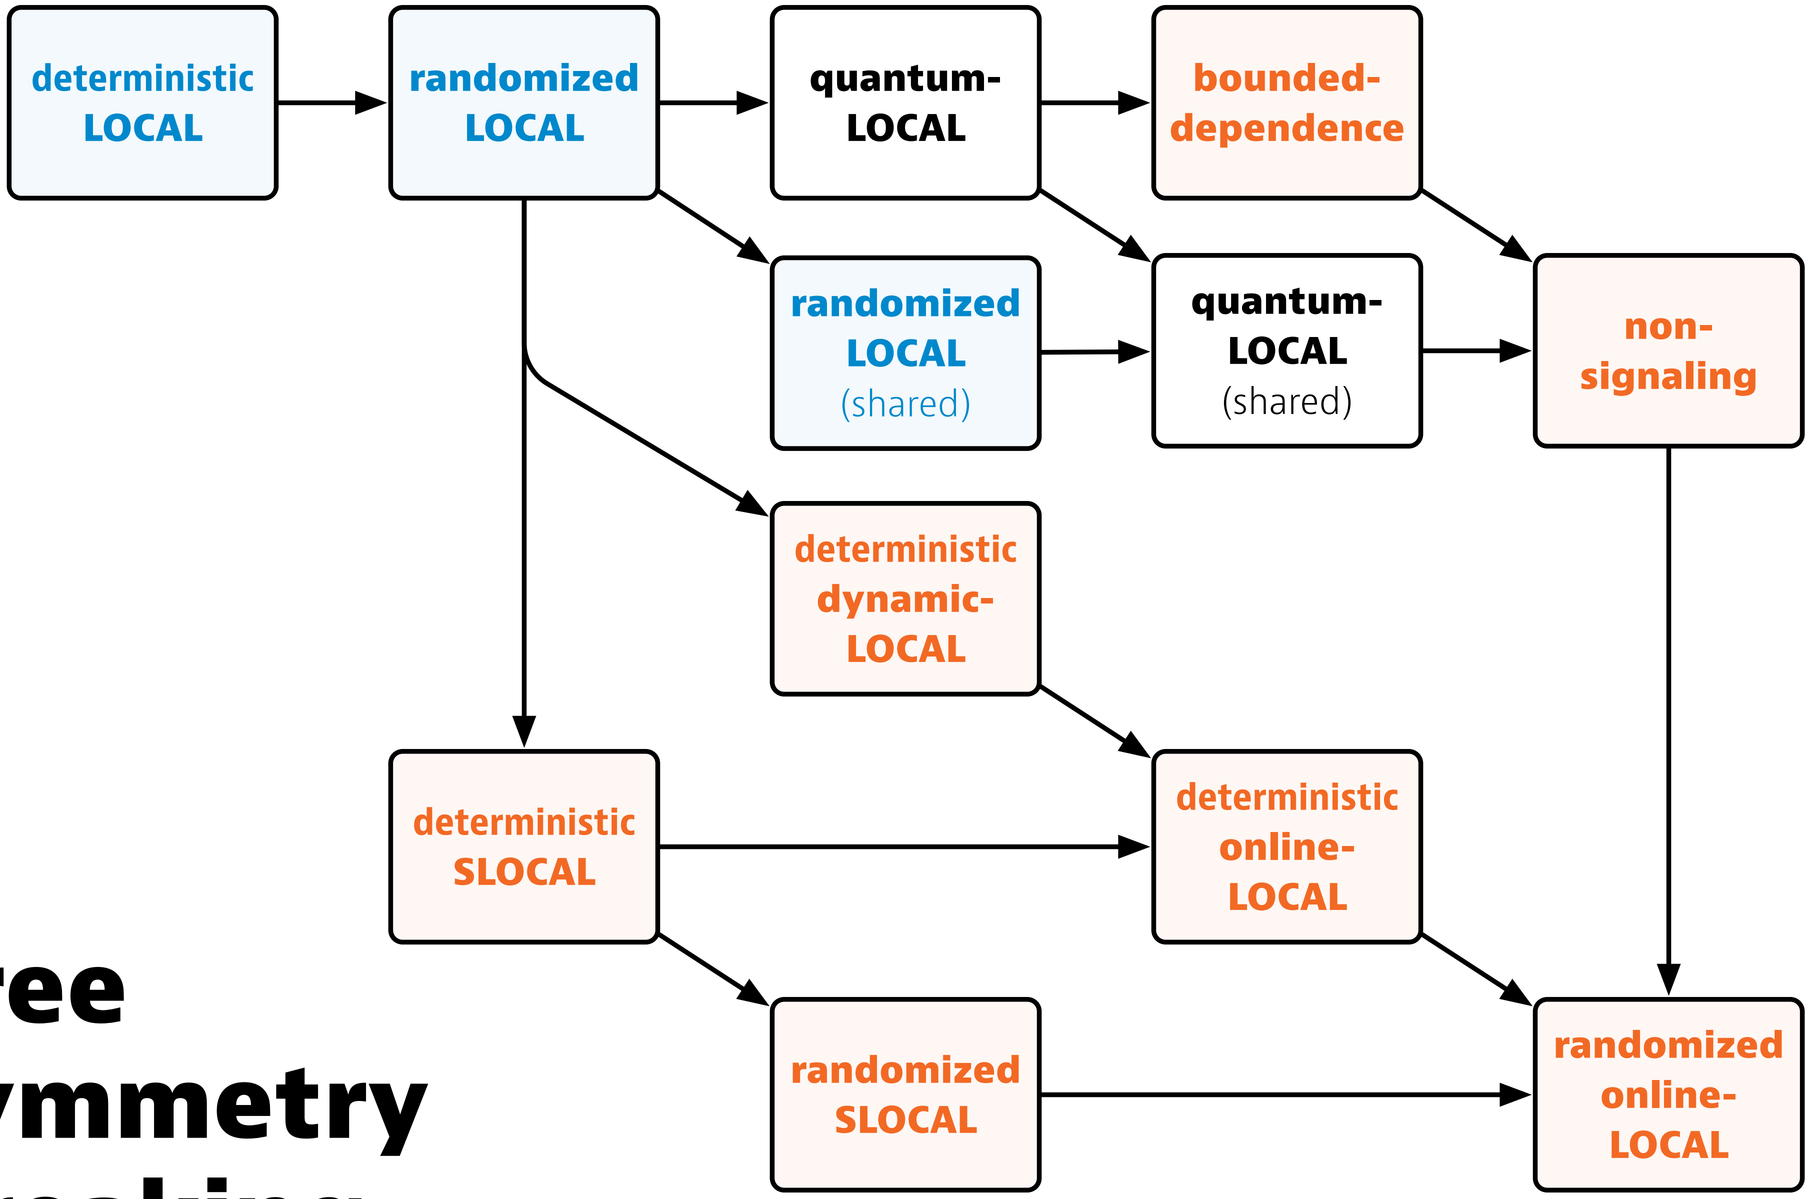

Free symmetry breaking

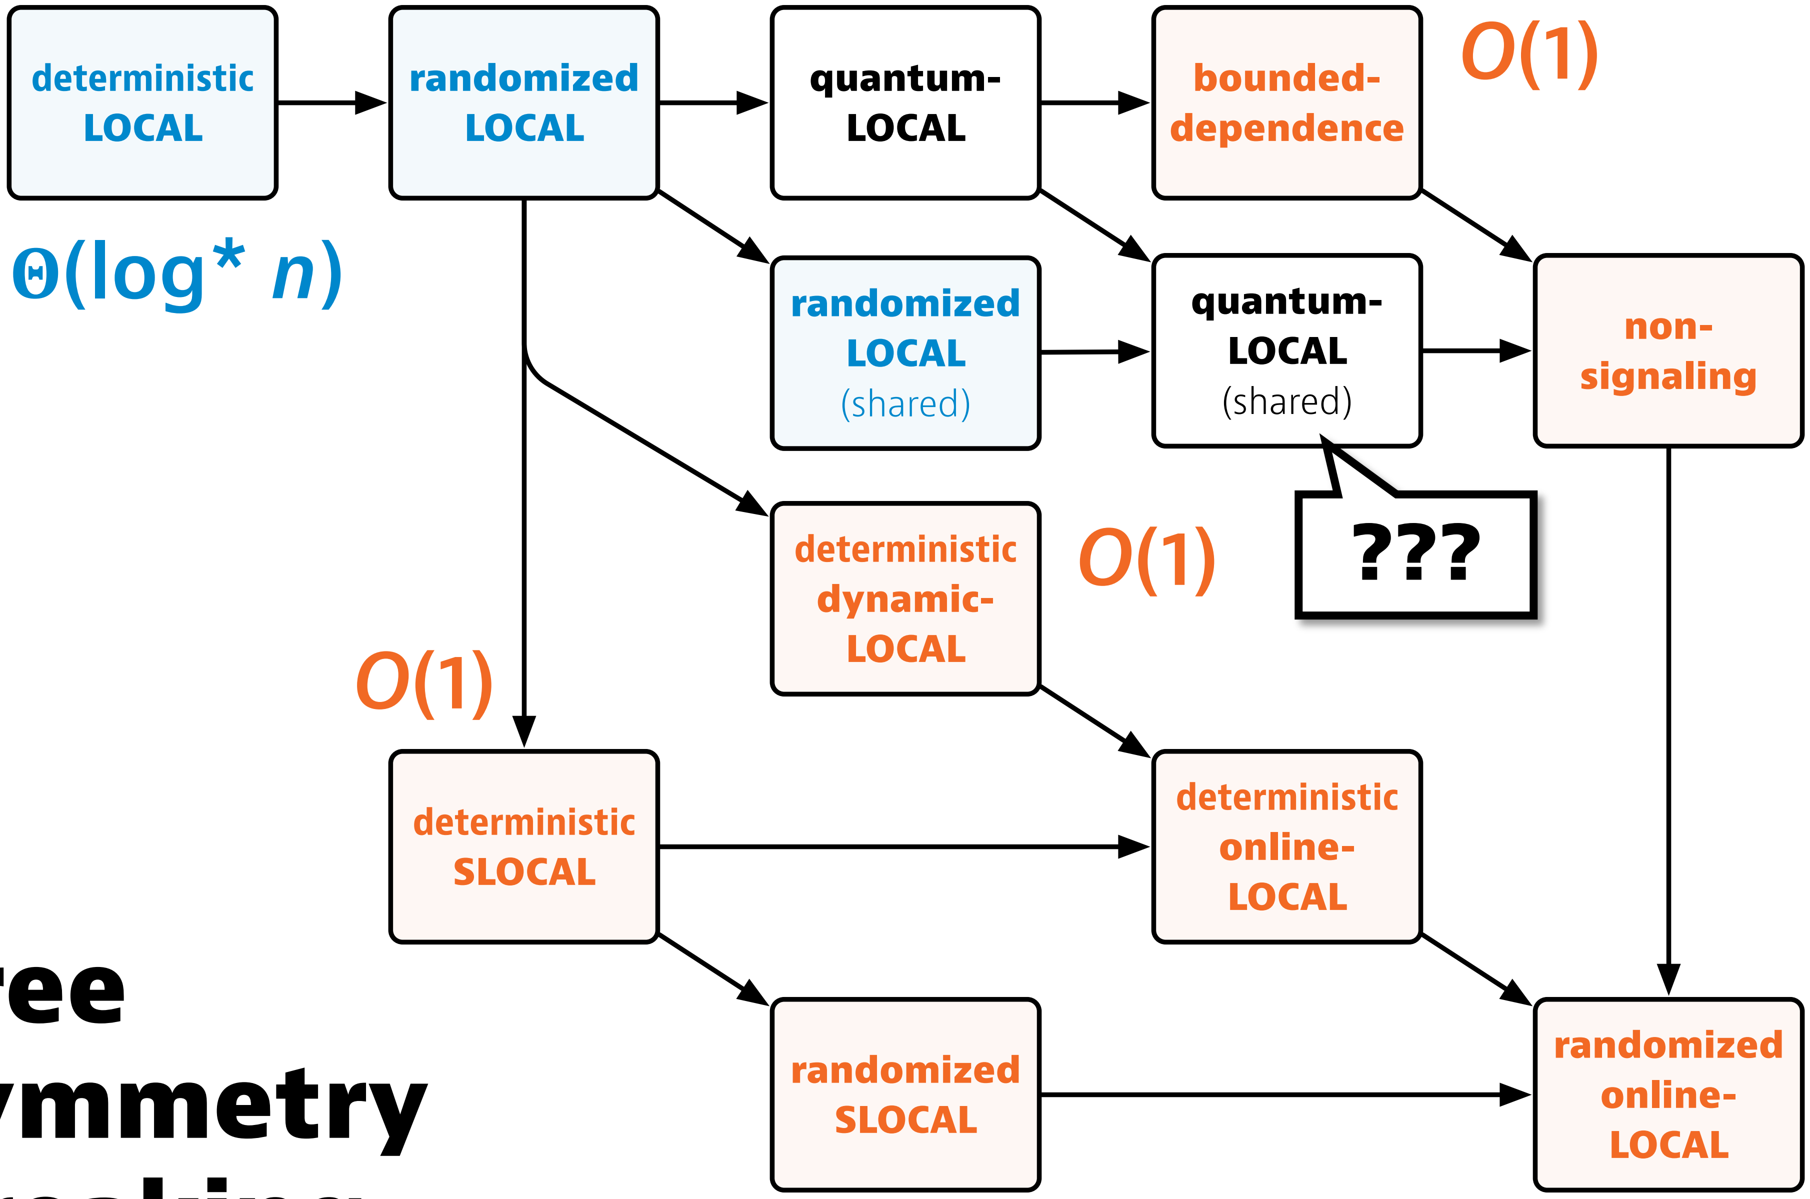

Free symmetry breaking

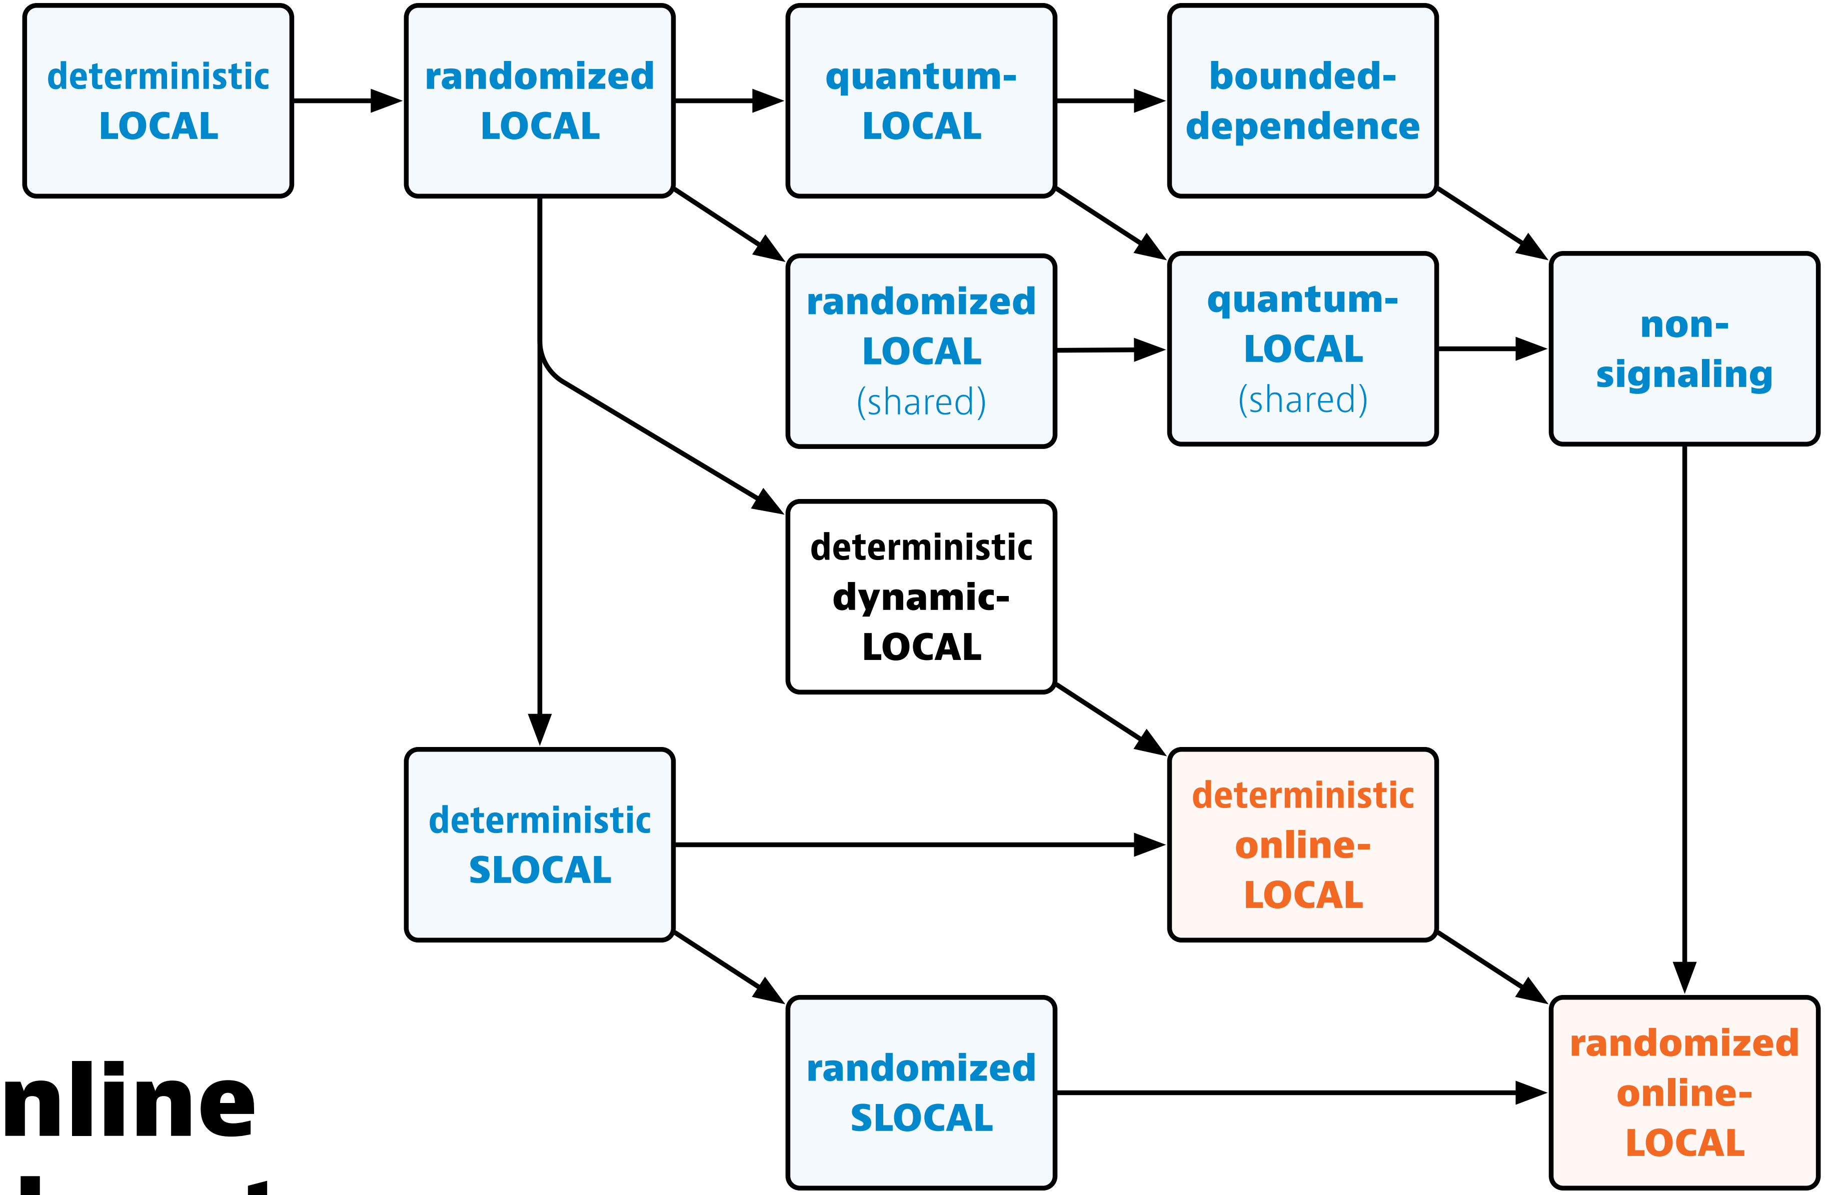

Online advantage

Online advantage

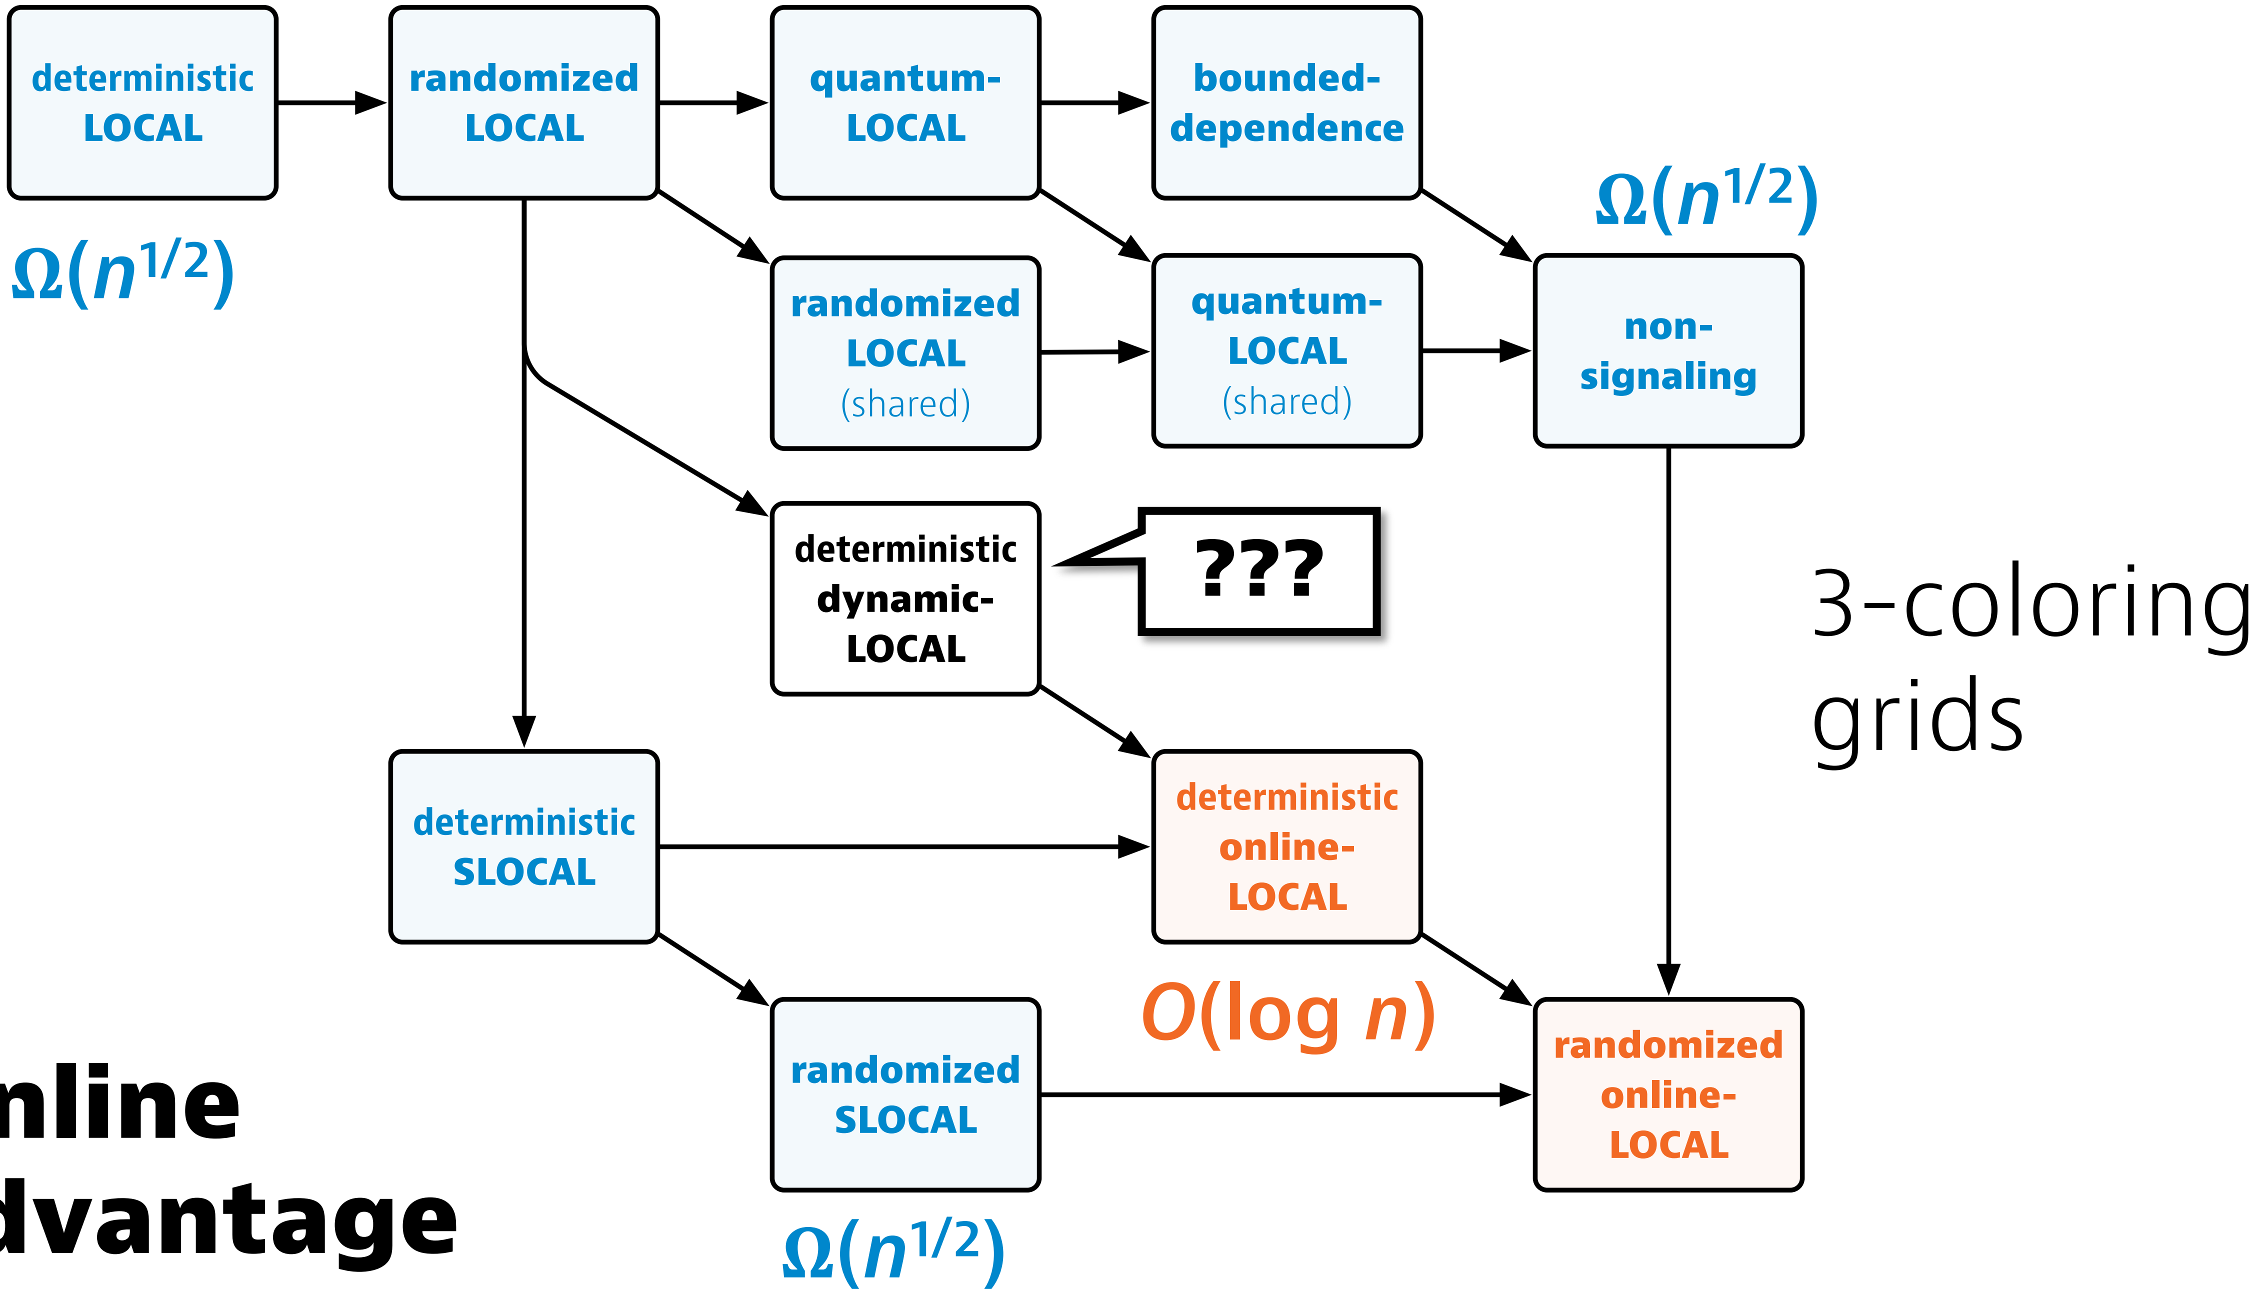

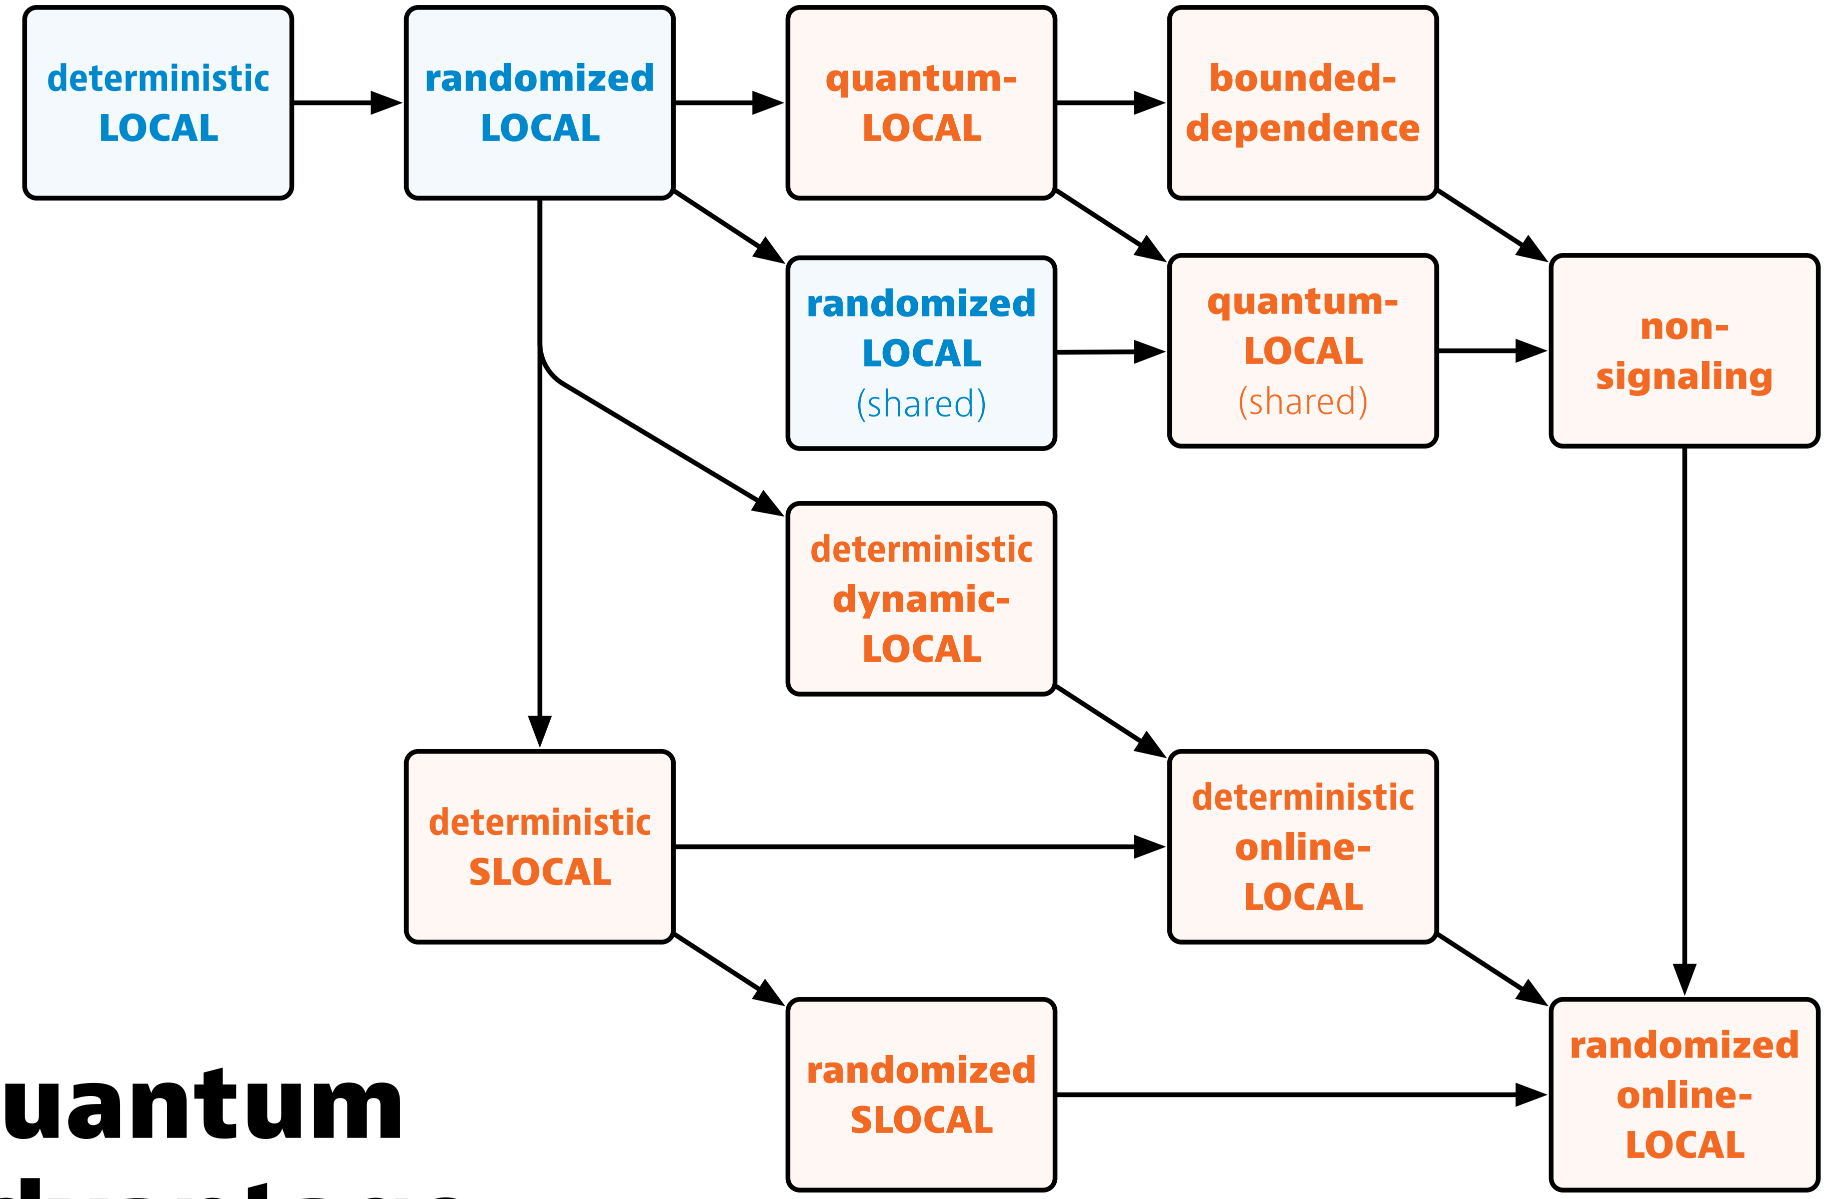

Quantum advantage

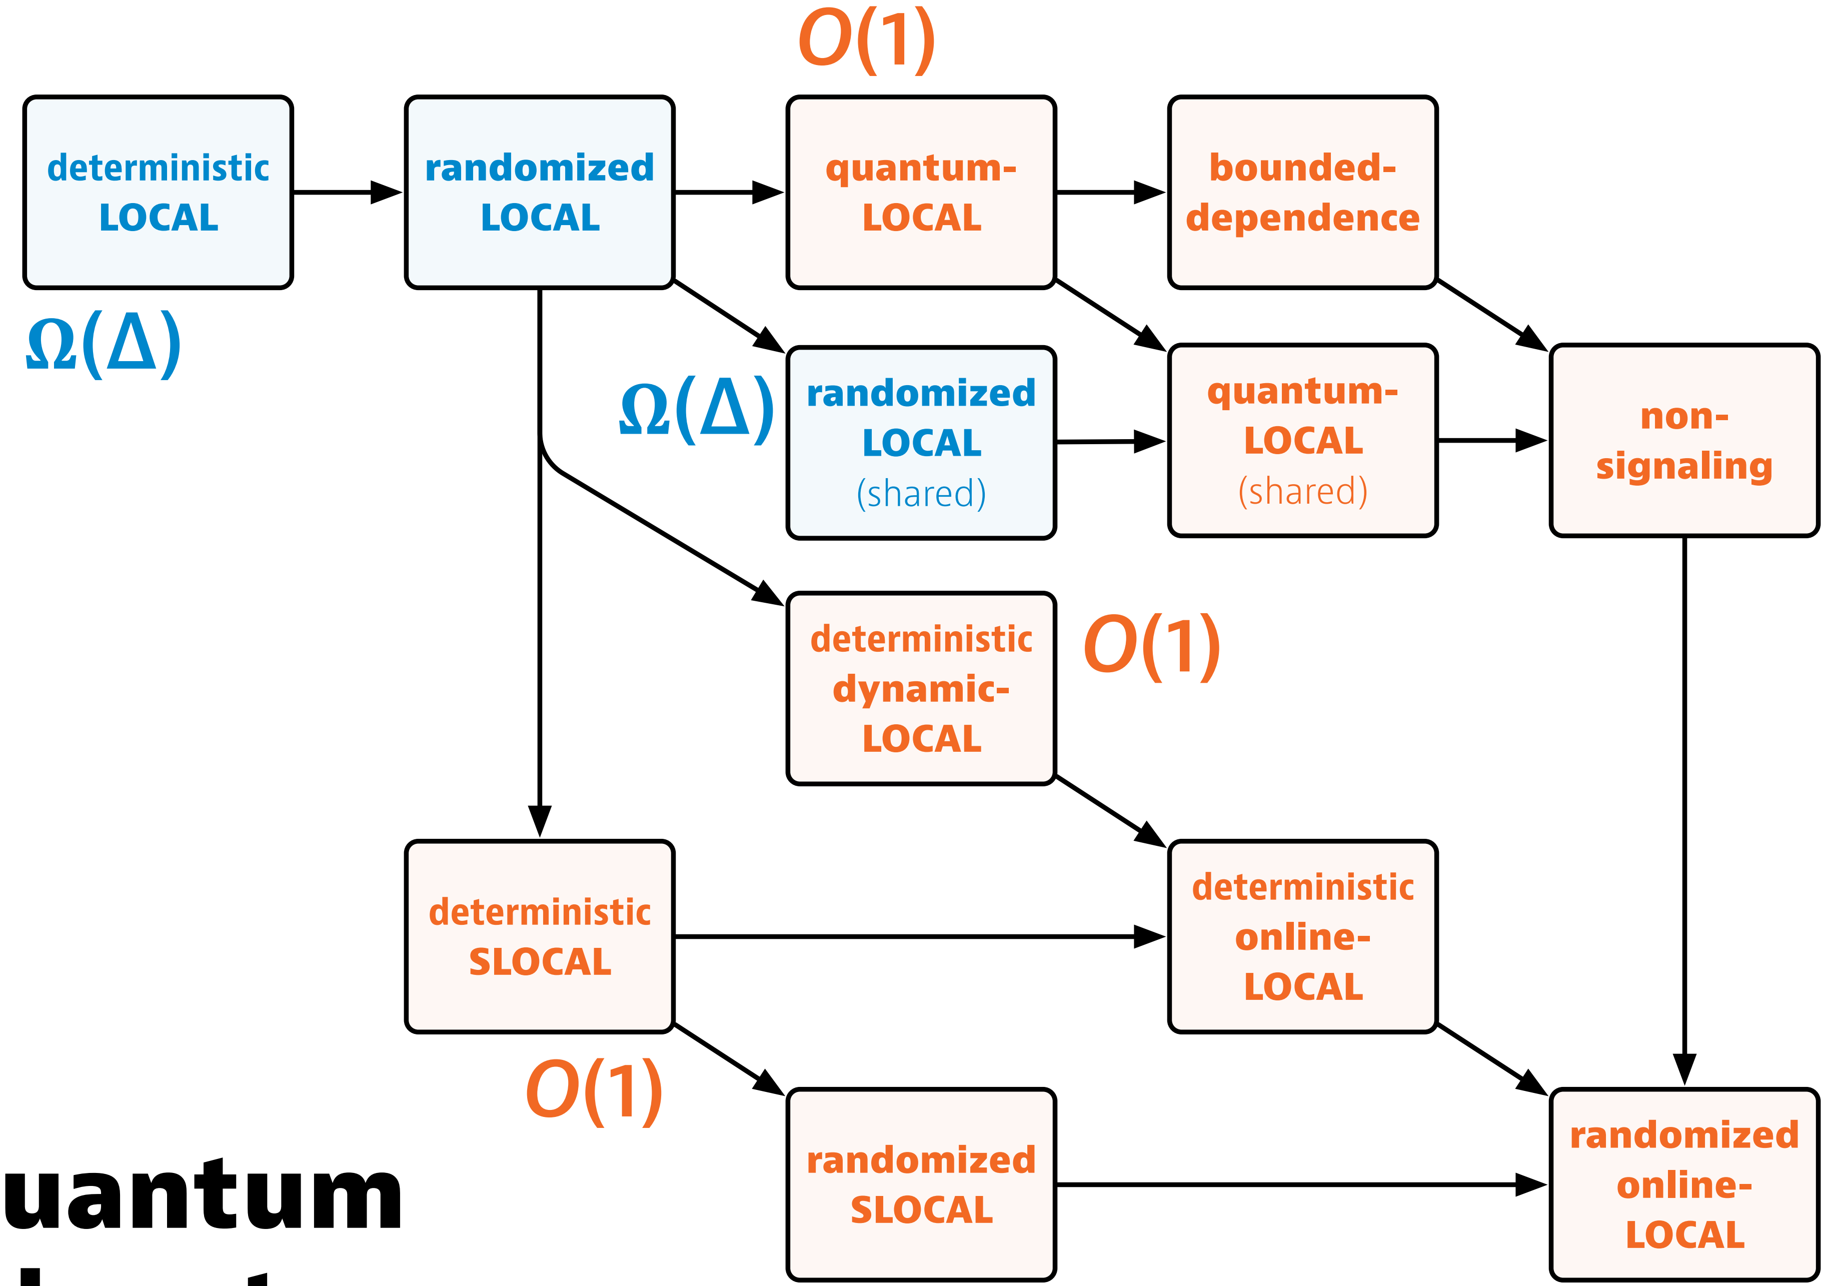

Iterated GHZ

Quantum advantage

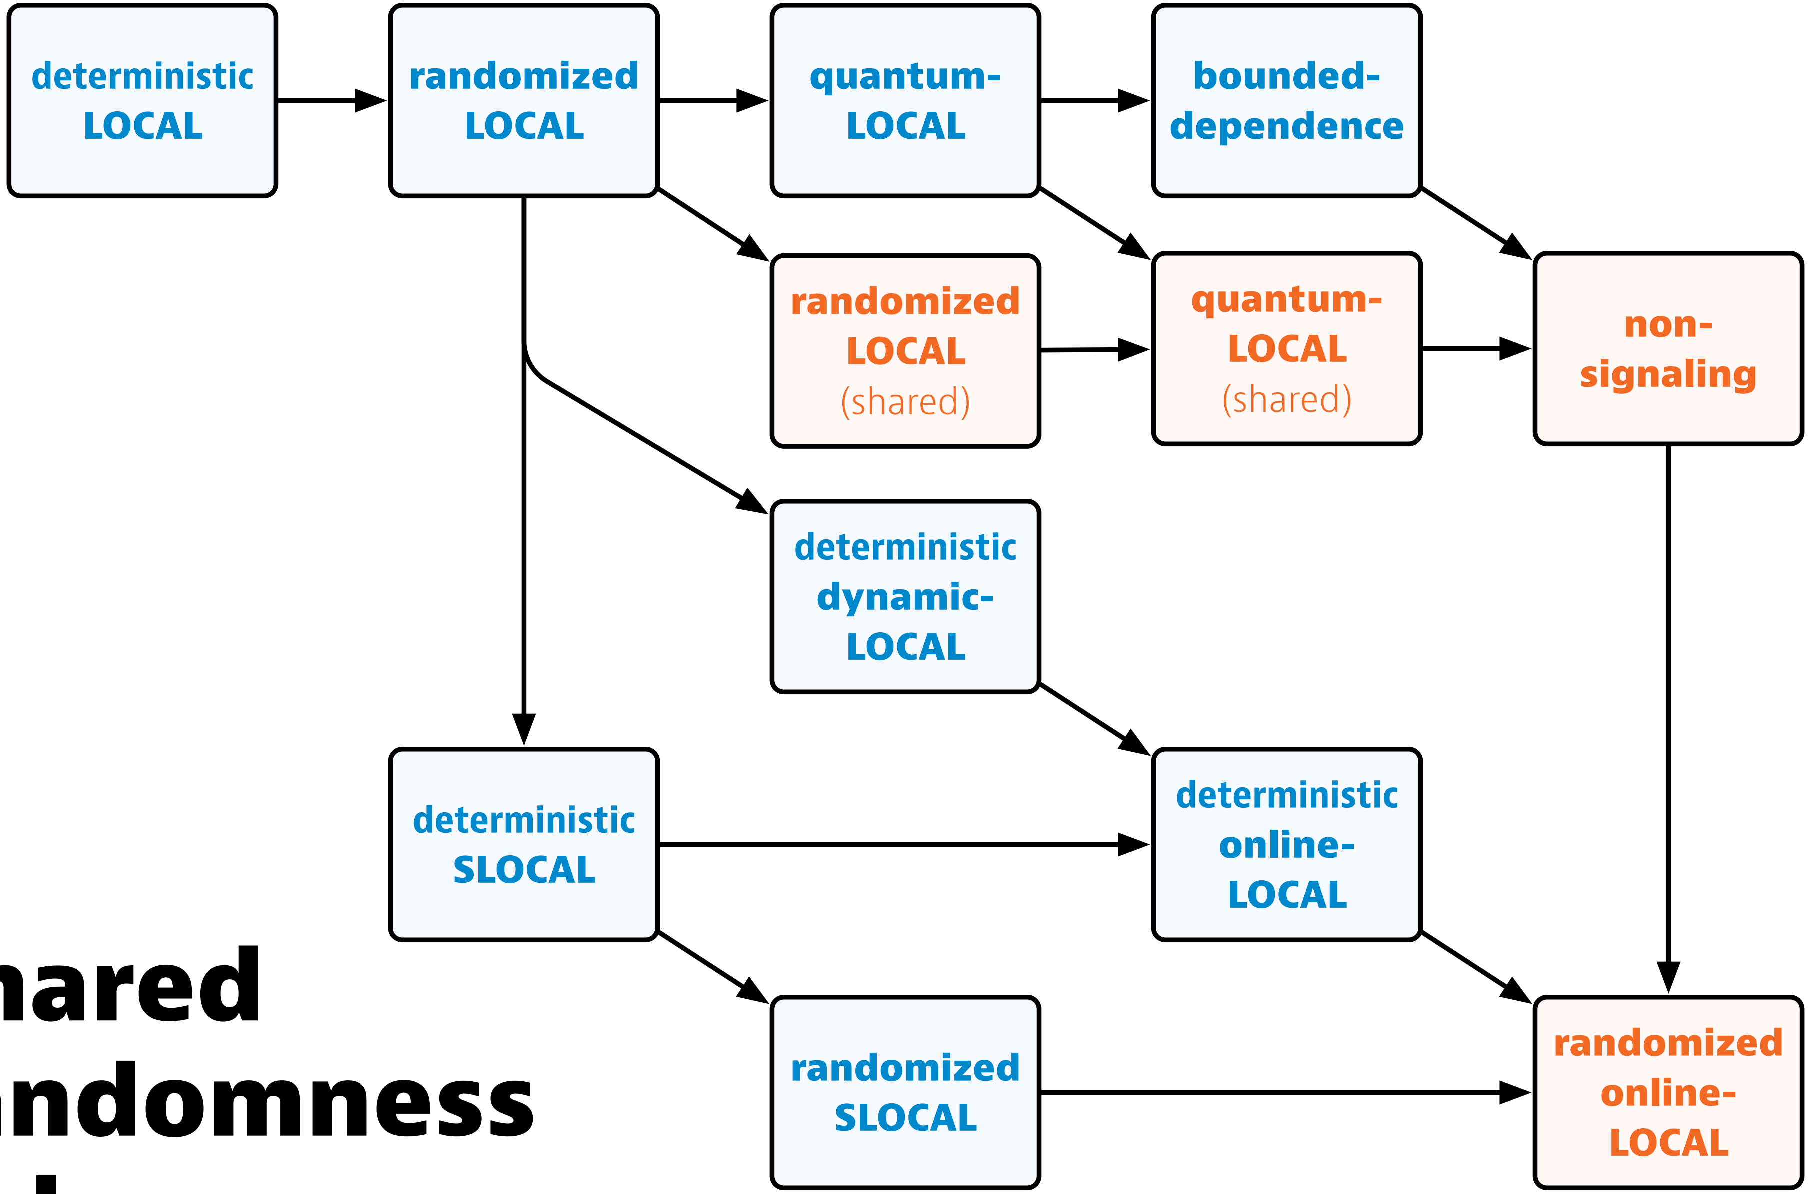

Shared randomness helps

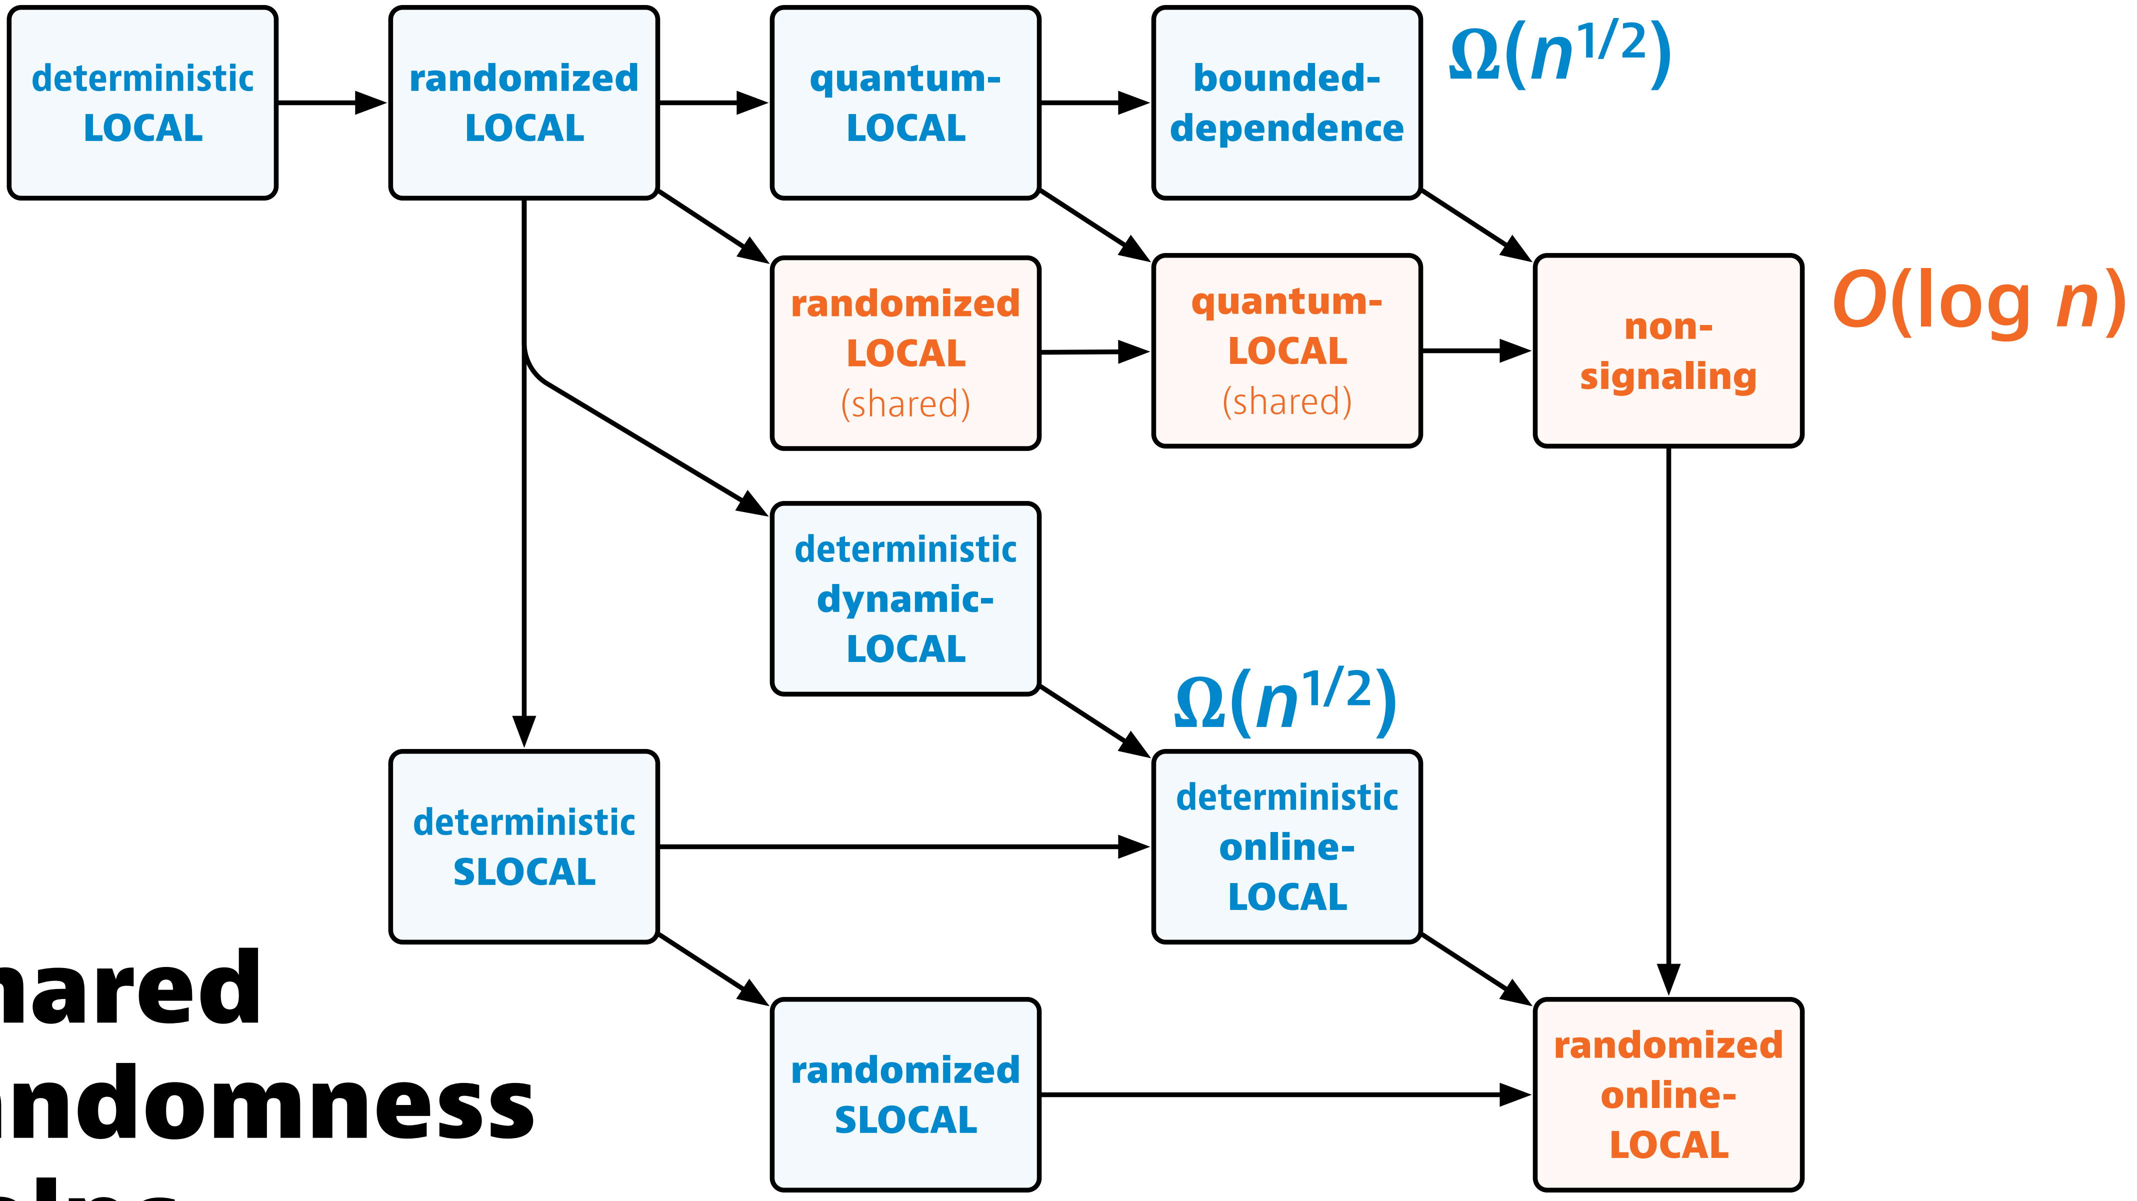

Shared randomness helps

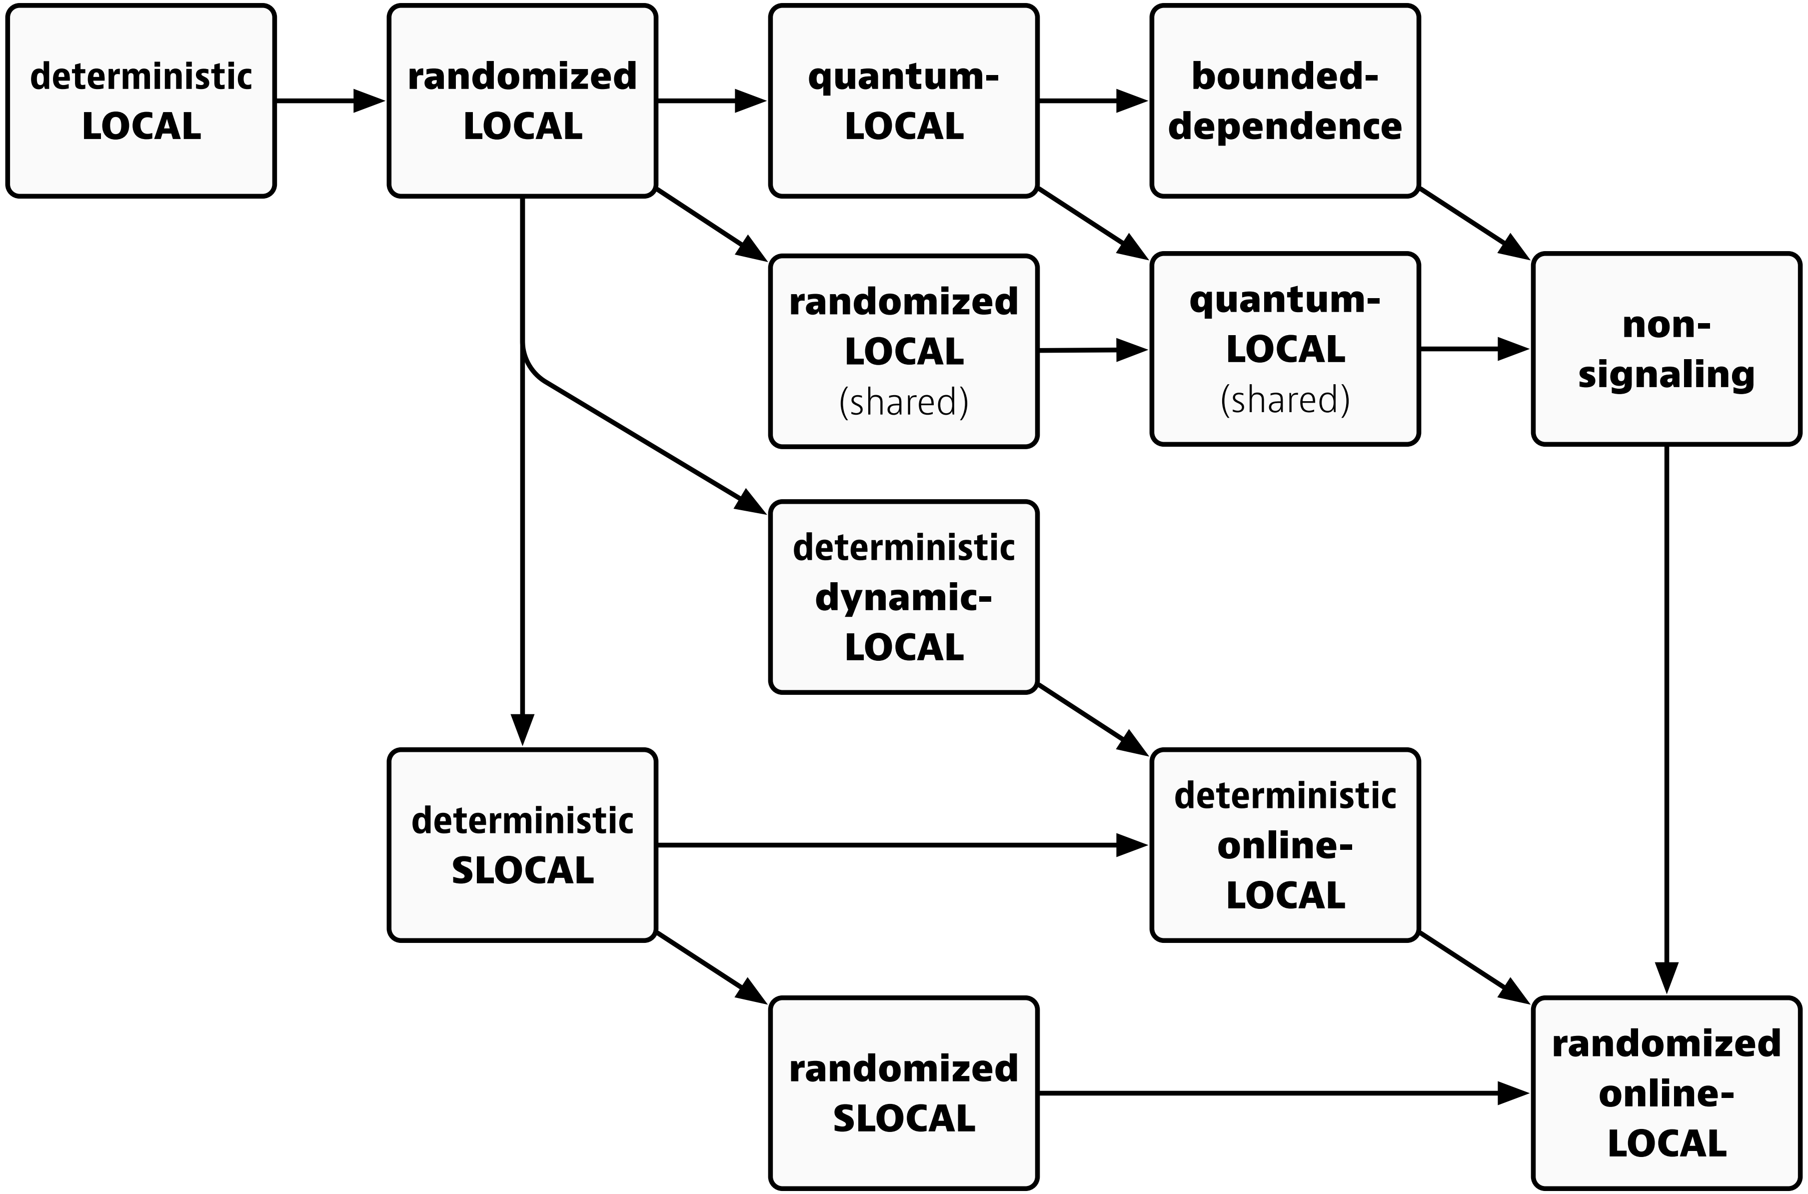
deterministic LOCAL
randomized LOCAL
quantum-LOCAL
bounded-dependence
randomized LOCAL (shared)
quantum-LOCAL (shared)
non-signaling
deterministic dynamic-LOCAL
deterministic SLOCAL
deterministic online-LOCAL
randomized SLOCAL
randomized online-LOCAL

Some equivalences

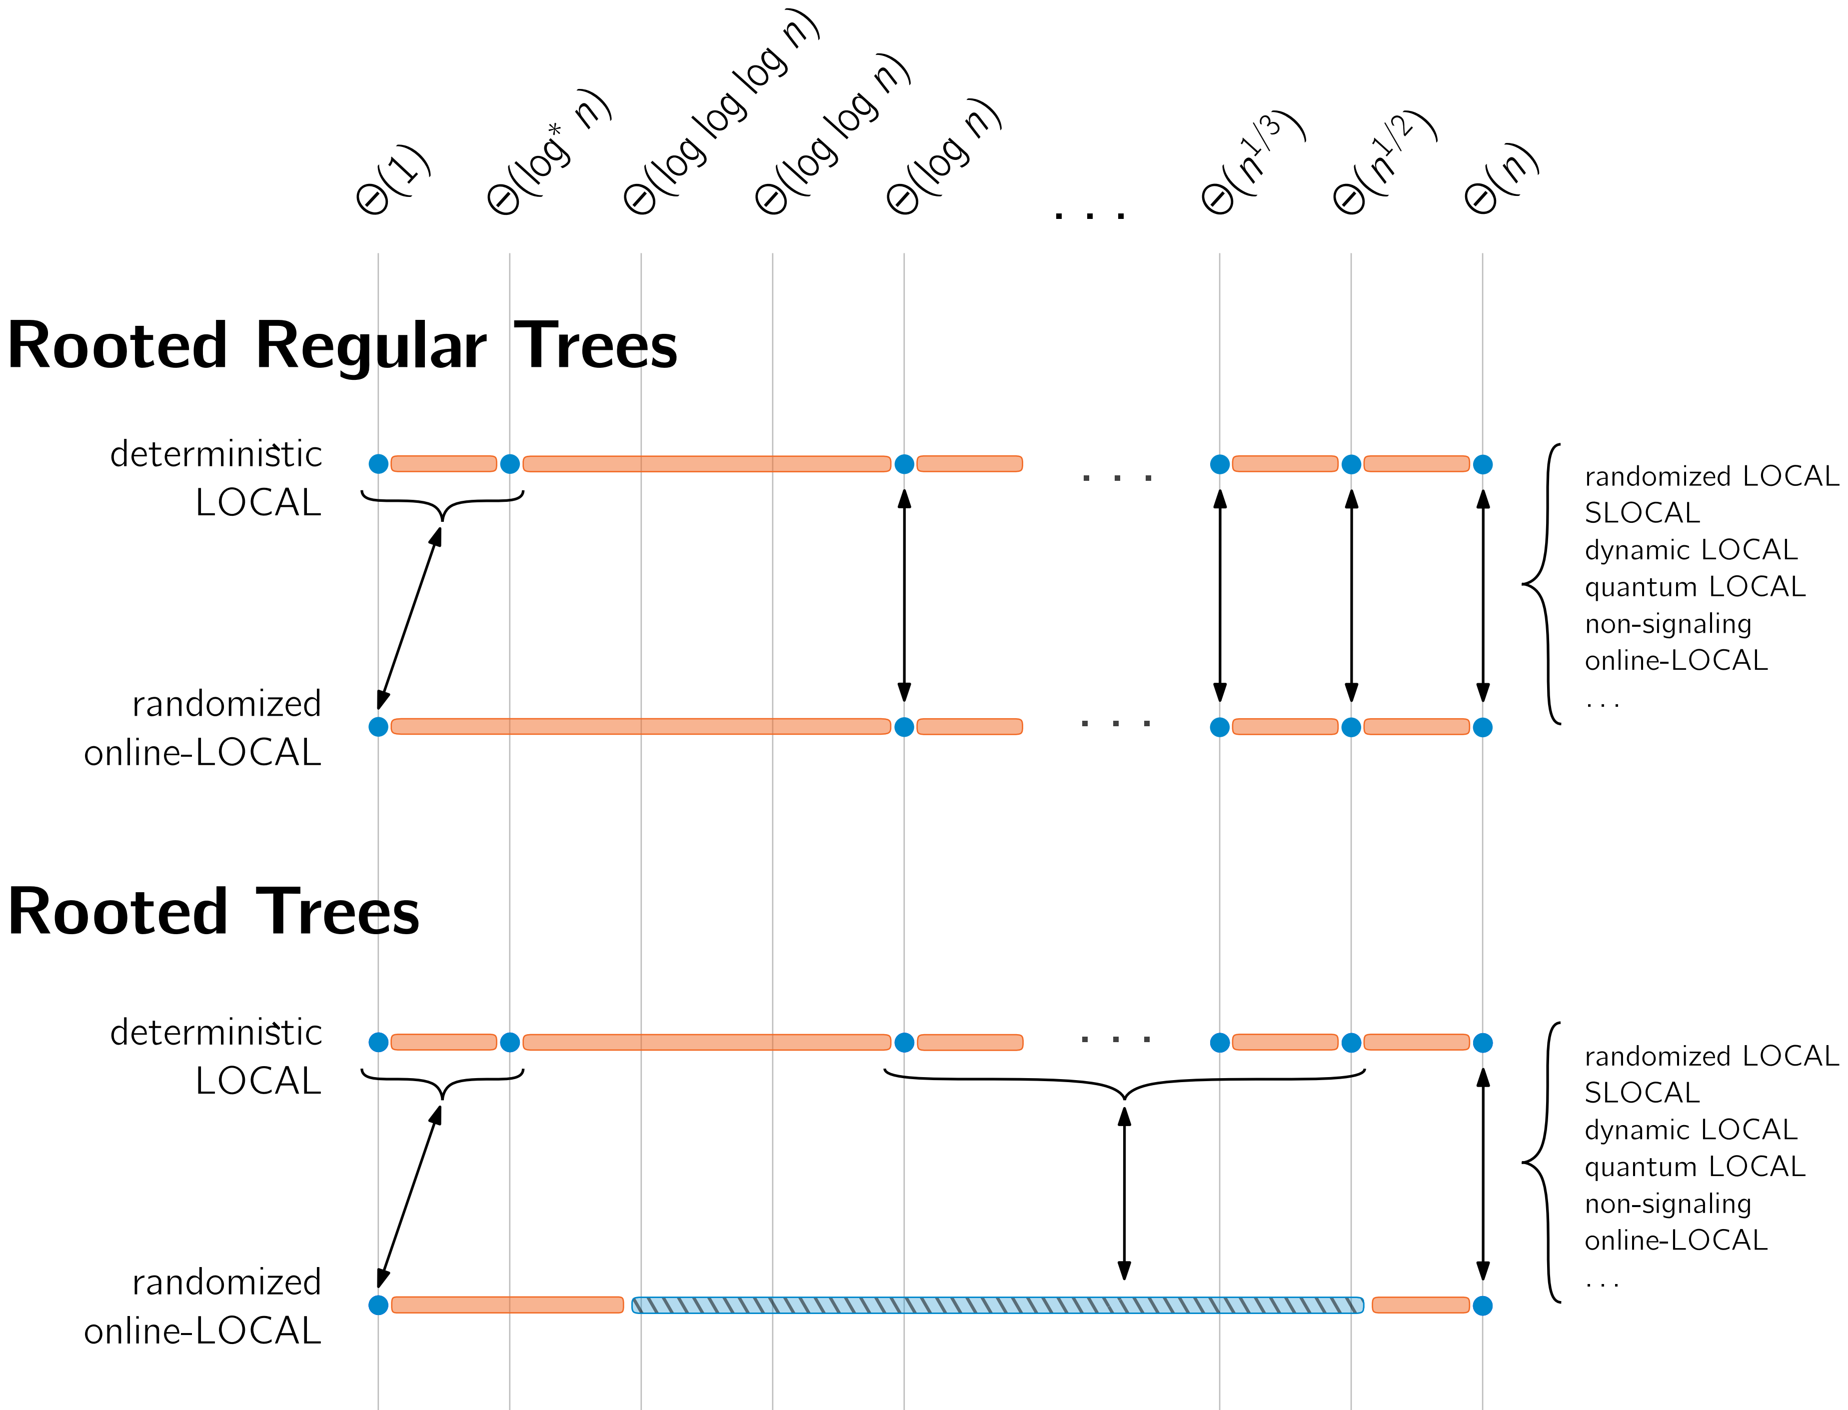

Θ(1)
Θ(log* n)
Θ(log log log n)
Θ(log log n)
Θ(log n)
...
Θ(n^{1/3})
Θ(n^{1/2})
Θ(n)
Rooted Regular Trees
deterministic LOCAL
randomized online-LOCAL
randomized LOCAL
SLOCAL
dynamic LOCAL
quantum LOCAL
non-signaling
online-LOCAL
...
Rooted Trees
deterministic LOCAL
randomized online-LOCAL
randomized LOCAL
SLOCAL
dynamic LOCAL
quantum LOCAL
non-signaling
online-LOCAL
...

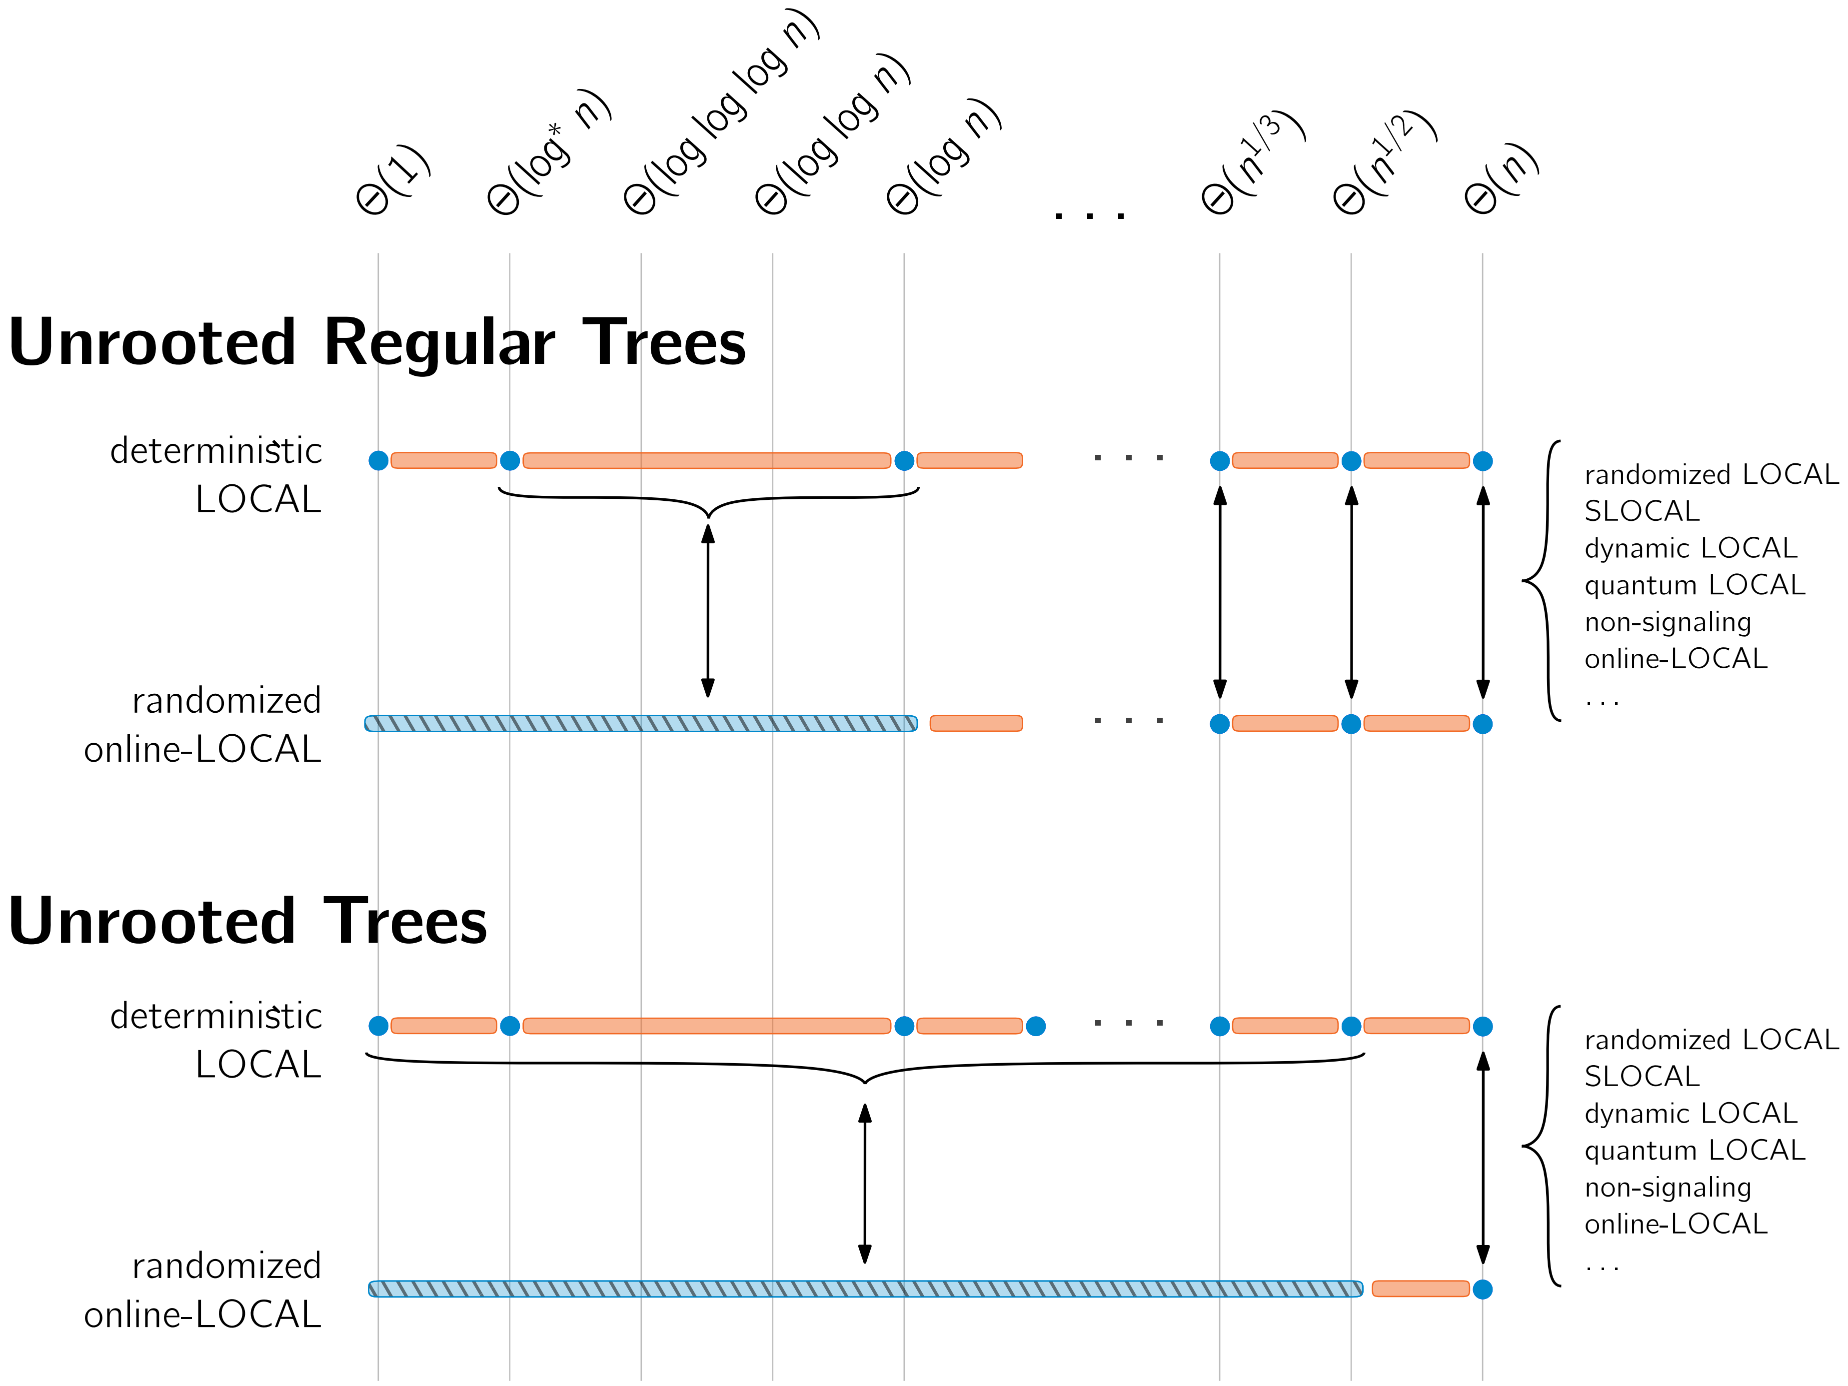

Θ(1)
Θ(log* n)
Θ(log log log n)
Θ(log log n)
Θ(log n)
. . .
Θ(n^1/3)
Θ(n^1/2)
Θ(n)
Unrooted Regular Trees
deterministic LOCAL
randomized online-LOCAL
randomized LOCAL
SLOCAL
dynamic LOCAL
quantum LOCAL
non-signaling
online-LOCAL
. . .
Unrooted Trees
deterministic LOCAL
randomized online-LOCAL
randomized LOCAL
SLOCAL
dynamic LOCAL
quantum LOCAL
non-signaling
online-LOCAL
. . .

Lots of open questions

Open questions

- **Does dynamic-LOCAL every help with LCLs?**
  - what is the complexity of ***3-coloring grids*** in dynamic-LOCAL?
  - what is the complexity of ***3-coloring bipartite graphs*** in dynamic-LOCAL?

Open questions

- **Can you make online-LOCAL algorithms promise-free?**
  - currently we need some kind of a ***promise***:
    "3-coloring grids", *"3-coloring in bipartite graphs"*
  - promise-free online-LOCAL algorithms don't seem to be any stronger than SLOCAL algorithms?

Open questions

- **Does global memory ever help in trees?**
  - could we ***simulate online-LOCAL algorithms in SLOCAL model*** in unrooted trees?
  - could we simulate dynamic-LOCAL algorithms?
  - or can we separate these for some LCL in trees?

Open questions

- **How to prove lower bounds in dynamic-LOCAL or online-LOCAL?**
  - e.g. how to show that you ***cannot 3-color trees with locality 10***?

Open questions

- **Is there any function-of-*n* distributed quantum advantage for LCLs?**
  - recall: we can show a function-of-Δ separation for a family of LCLs
  - can you ever have function-of-*n* quantum advantage for a single LCL?
  - example: can you ***color cycles*** in $O(1)$ rounds in ***quantum-LOCAL***?

Open questions

- **Lower bounds for sinkless orientation**
  - deterministic LOCAL: ≈ log *n*
  - randomized LOCAL: ≈ log log *n*
  - deterministic SLOCAL: ≈ log log *n*
  - randomized SLOCAL: ≈ log log log *n*
  - dynamic-LOCAL: ***no lower bounds!***
  - online-LOCAL: ***no lower bounds!***
  - quantum-LOCAL: ***no lower bounds!***

Open questions

- **New lower-bound proof techniques?**
  - ***existential graph-theoretic arguments*** do not work e.g. in online-LOCAL
  - ***round elimination*** do not work e.g. in quantum-LOCAL
  - current best hope: new ***simulation results***?

- arXiv:2109.06593 **Locality in online, dynamic, sequential, and distributed graph algorithms** (*ICALP 2023*)
- arXiv:2307.09444 **No distributed quantum advantage for approximate graph coloring** (*STOC 2024*)
- arXiv:2403.01903 **Online locality meets distributed quantum computing**
- arXiv:2407.05445 **Shared randomness helps with local distributed problems**
- arXiv:2409.13795 **Local problems in trees across a wide range of distributed models** (*OPODIS 2024*)
- arXiv:2411.03240 **Distributed quantum advantage for local problems**