Can we find a short path, without a map?

Can we parallelize maze-solving?

Romain Cosson

Based on joint works with Laurent Massoulié & Laurent Viennot

PhD Seminar

Keywords: collaborative exploration, trees, competitive analysis

20 slides / 20 mins

Online framework: the science of decision-making

Online problem = Information arrives over time. Examples :

• Linear Search	*(Where should I search for my lost wallet?)*	**Starter**
• Learning with expert advice	*(Who should I trust?)*	
• Metrical Task Systems	*(When should I move out to another city?)*	
• Secretary problem	*(When should I stop a hiring process?)*	
• Bandit problems	*(Should I be safe (exploitation) or be bold (exploration)?)*	
• List Update Problem	*(How should I organize my bookshelf?)*	
• Collective Tree Exploration	*(How to solve a maze, with a team?)*	**Main Dish**
• Layered Graph Traversal	*(How to find a short path, without a map?)*	

A simple example: **Linear Search**

Where should I look for my lost wallet?

- Introduced by Bellman in the 1950s
  - *you* start from $0 \in \mathbb{Z}$ in an infinite street, your Cost = #steps
  - *wallet* is lost at $x \in \mathbb{Z}$ at distance $|x| = \mathrm{OPT}$ (unknown)
  - Your strategy has competitive ratio α if you always find your wallet with at most

$$\mathbf{Cost} \leq \boldsymbol{\alpha} \cdot \mathbf{OPT}$$

The doubling strategy is 9-competitive

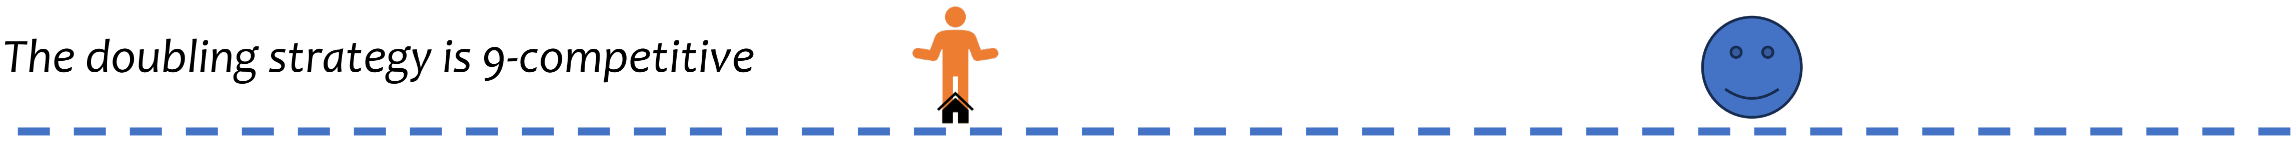

Another example : **Learning with Expert advice**

Who should I trust when there are many self-claimed « experts »?

- Introduced independently in many fields in the late 20th century.
  - Each day $t \in \{1, \dots, T\}$ there are n experts forecasting rain ($y_{t,n} = 1$) or sun ($y_{t,n} = 0$)
  - OPT = number of mistakes made by the most accurate forecaster
  - Cost = the number of mistakes that you make
  - Your strategy has regret $R(n, T)$ if it satisfies

$$\mathbf{Cost} \leq \mathbf{OPT} + \boldsymbol{R(n, T)}$$

The multiplicative weights strategy has $R(n, T) = \mathcal{O}(\sqrt{T \log n})$ *regret*

Recap: on online problems

- OPT : the optimal possible cost, if you had all the information !
- Cost : the actual cost you pay, $\mathrm{Cost} \geq \mathrm{OPT}$
- A strategy acheives :
  - **Competitive ratio** $\boldsymbol{\alpha}$ if

$$\mathrm{Cost} \leq \boldsymbol{\alpha} \times \mathrm{OPT}$$

  - **Regret** $\boldsymbol{R}$ if

$$\mathrm{Cost} \leq \mathrm{OPT} + \boldsymbol{R}$$

Main dish !

- **Layered Graph Traversal**
  - *Can you find the shortest path, when you don't have a map?*
  - Papadimitriou and Yannakakis, 1991 (online algorithms)

- **Collective Tree Exploration**
  - *Is maze-solving parallelizable?*
  - Fraigniaud, Gasieniec, Kowalski and Pelc, 2004 (distributed algorithms)

Layered Tree Traversal

Can you find the shortest path,
Without a map?

Can you find the shortest path, without a map?

Notation : n = #nodes = #edges+1 and D = depth and w = width

- **« the width » :** max # nodes at given combinatorial depth
- **« a layer » :** set of nodes at the given combinatorial depth

<u>**Online Problem**</u> : Layers are revealed one after the other !

- $\mathbf{OPT} = \boldsymbol{D}$ so an « $\boldsymbol{\alpha(w)}$**-competitive path** » has length

$$\text{Cost} \leq \alpha(w)D$$

Example with width w = 3, depth D = 5

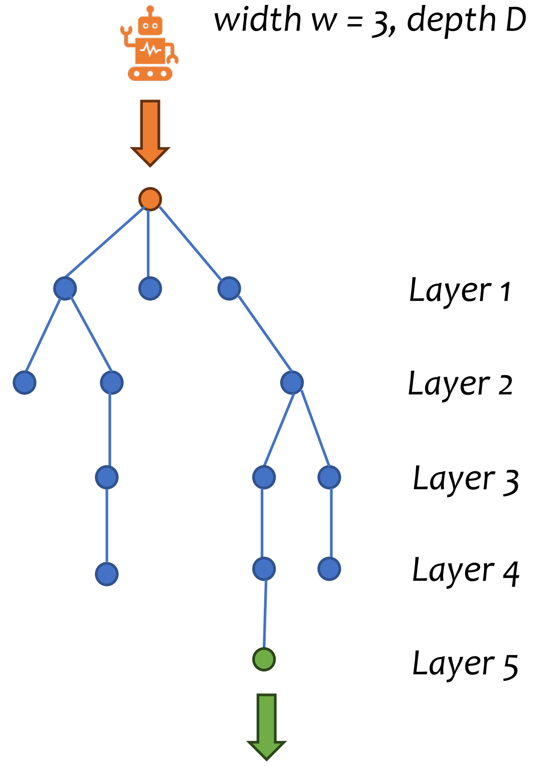

15 moves, i.e. 3-competitive

Is layered-feedback realistic?

width w

length D

The unweighted variant

✓ **Observation:** In **unweighted** tree $n = wD$

✓ Depth-First Search is thus $O(w)$ - competitive

✓**Question:** Can we do better than Depth-First Search?

Our work [CM]

✓ **Yes**! There is a $O(\sqrt{w})$-competitive strategy

✓ *For a more general formulation of the problem*

✓ *Uses random choices!*

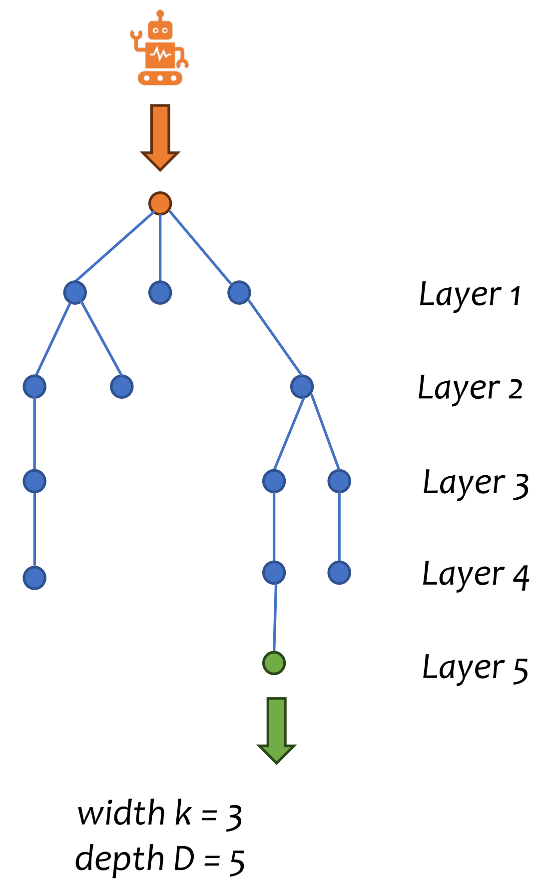

Definition and Motivation

Collective Tree Exploration

Is maze-solving parallelizable?

Is solving a (tree) maze parallelizable?

Notation : n = #nodes = #edges+1 and D = depth

- **With 1 agent:** right-hand on wall (RHW, aka DFS) $\leq 2n$ moves
- **With 2 agents:** right+left-hand on wall (RHW+LHW) $\leq n$ moves each
- **What about $k \geq 3$ agents?** Moving synchronously at each round
- Exploration <-> Finding exit

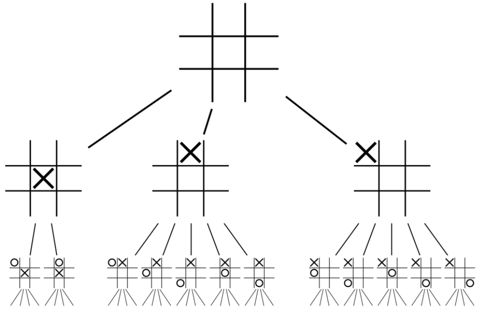

Tic-tac-toe tree game

Silver coin of Knossos (Crete, 400 BC)
Berlin state museum

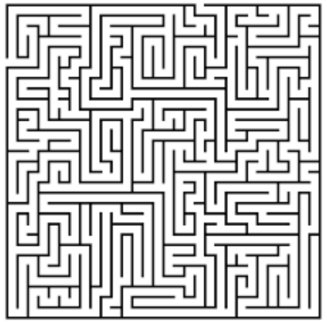

A computer-generated (tree) maze


VLAM!
?!
VOUS NE SORTIREZ JAMAIS D'ICI, ÉTRANGERS! CE TOMBEAU SERA VOTRE TOMBEAU!

« Collective Tree Exploration »

[FGKP 2004] Fraigniaud, Gasieniec, Kowalski, Pelc *Collective Tree Exploration*

- **Goal:** Traverse all edges of unknown tree $\mathrm{T} = (V, E)$ with n nodes and depth D
  - ✓ With $k \in \mathbb{N}$ agents moving synchronously at each round

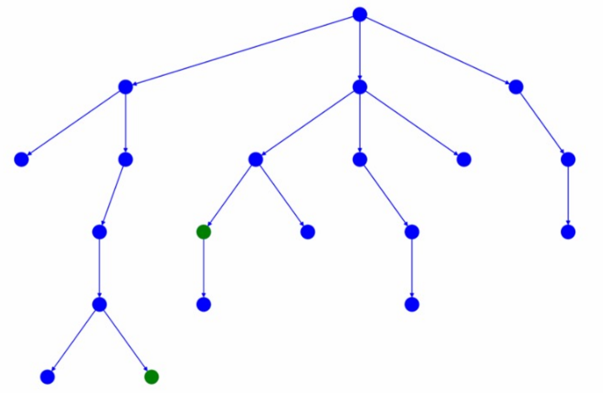

A tree with $n = 20$ and $D = 6$

- **Communication Models:**
  - ✓ *Centralized* (full « complete » communication)
  - ✓ *Distributed* (restricted « write-read » communication)

- **Main result** [FGKP 2004] :
  - ✓ Distributed algorithm SPLIT in which explorers split evenly at intersections
  - ✓ $\text{Runtime}_{\text{SPLIT}} \leq O\left(\frac{2n}{\log k} + D\right)$

The *Competitive Ratio* approach

What runtime could we hope for, had we know the tree in advance?

➢ **Offline variant :** OPT = Exploration time if tree were known in advance

$D \leq \mathrm{OPT}$ and $\frac{n}{k} \leq \mathrm{OPT}$ in fact, $\frac{n}{k} + D \approx \mathrm{OPT}$

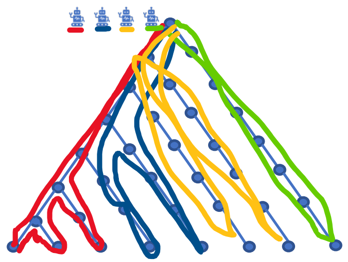

➢ **Consequence :** SPLIT has a **competitive ratio** in $O\left(\frac{k}{\log k}\right)$

$$\mathrm{Runtime}_{\mathrm{SPLIT}} = \mathcal{O}\left(\frac{n}{\log k} + D\right) \leq O\left(\frac{k}{\log k}\right)\mathrm{OPT}$$

Recent results

✓ Continuous **analysis** of online algorithms (online convex optimization)

✓New idea: the explorers behave like electrons in an **electric network**

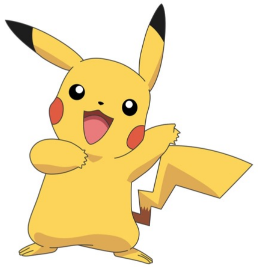

Latest results [C., Massoulié]

- Regret: $\boldsymbol{O(kD)}$
- Competitive Ratio: $O(\sqrt{\boldsymbol{k}})$

Open Questions

Some Open Questions

Open Question 1: Is there a competitive collective **graph** exploration algo

(competitive ratio) $c(k)\left(\frac{m}{k} + D\right)$ or in $\frac{2m}{k} + f(k, D)$ (regret)

where m is # of edges, ***D*** **is graph diameter,** and $f(\cdot,\cdot)$ is some arbitrary function.

Open Question 2: Is there a (competitive) gap between ***distributed*** and ***centralized*** collective tree exploration?

Thank you !

Feel free to reach out

Office # C320

Why $\log k$ appears in SPLIT?

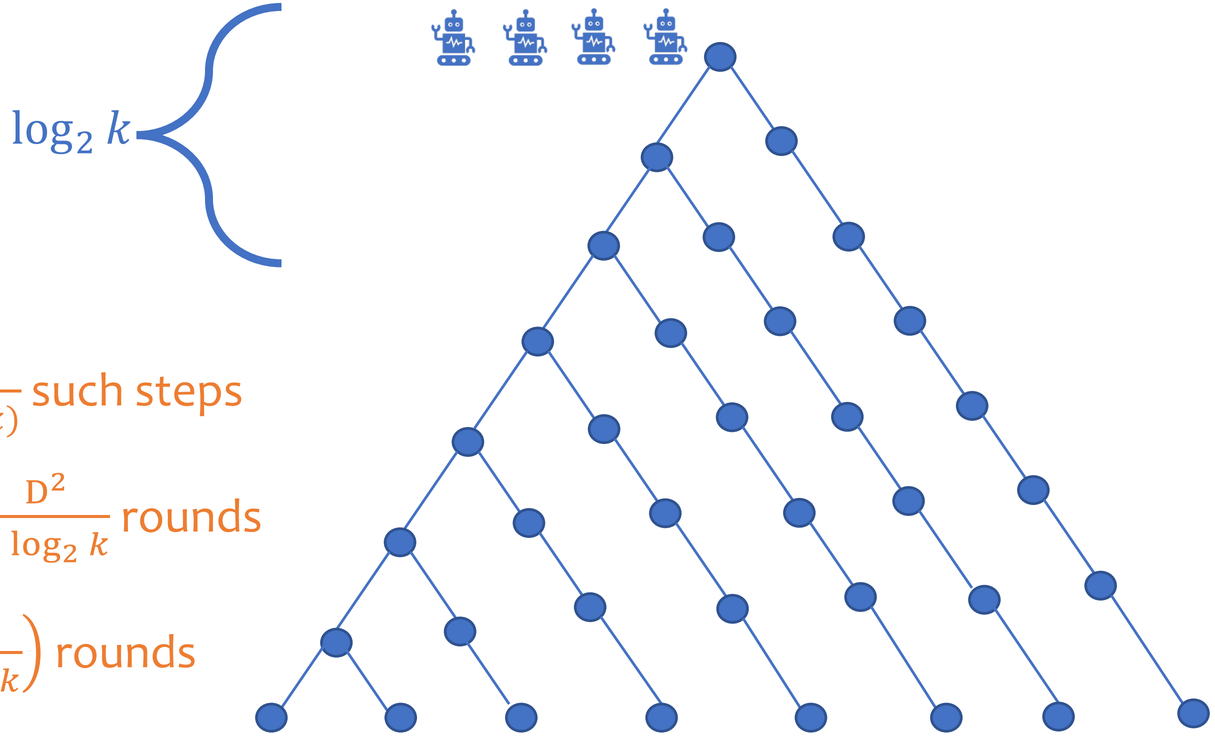

The « comb »

$$n \approx \frac{D^2}{2}$$

We need $\frac{D}{\log_2(k)}$ such steps

This means $\approx \frac{D^2}{2 \log_2 k}$ rounds

At least $\Omega\left(\frac{n}{\log_2 k}\right)$ rounds

[HKLT 2014] This lower-bound applies to any « greedy » algorithm.

def. « Greedy » = a robot never goes upwards, if there is an unexplored edge below