Graph Exploration Algorithms

Dariusz Dereniowski

Faculty of Electronics, Telecommunications and Informatics
Gdańsk University of Technology

8th workshop on GRAph Searching, Theory & Applications, Anogia, Crete, Greece, April 12, 2017

Outline

- we consider only unknown graphs

Outline

- we consider only unknown graphs
- exploration of unlabeled networks (just to note)

Outline

- we consider only unknown graphs
- exploration of unlabeled networks (just to note)
- exploration of labeled graphs

Outline

- we consider only unknown graphs
- exploration of unlabeled networks (just to note)
- exploration of labeled graphs
  - piecemeal exploration of labeled graphs (more generally, energy constrained)

Outline

- we consider only unknown graphs
- exploration of unlabeled networks (just to note)
- exploration of labeled graphs
  - piecemeal exploration of labeled graphs (more generally, energy constrained)
  - exploration time of labeled graphs

Unlabeled networks

- vertices have no identifiers,
- edges have port numbers

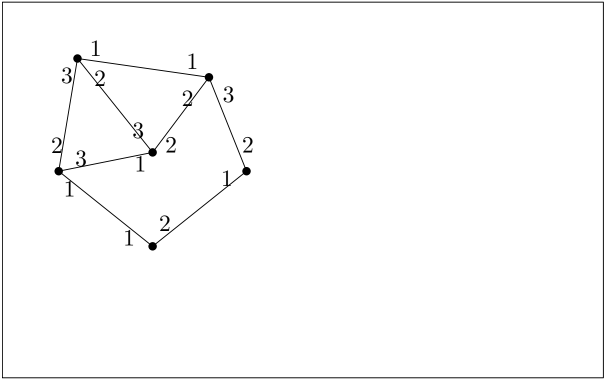

Unlabeled networks

- vertices have no identifiers,
- edges have port numbers

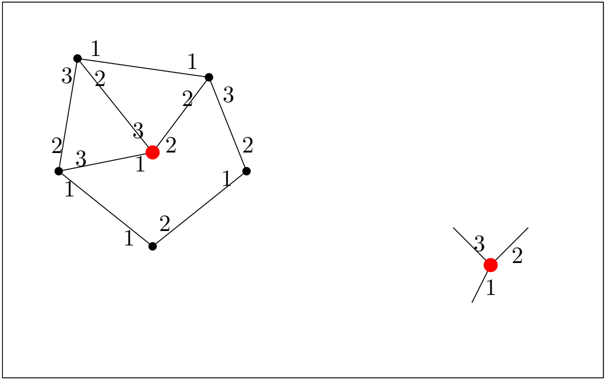

Unlabeled networks

- vertices have no identifiers,
- edges have port numbers

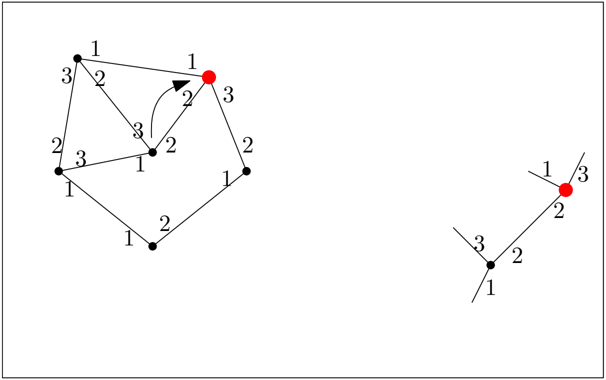

Unlabeled networks

- vertices have no identifiers,
- edges have port numbers

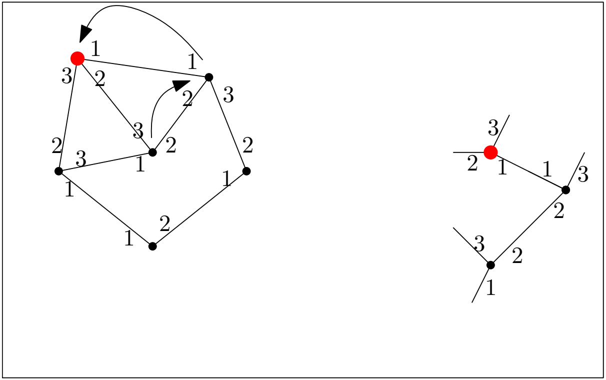

Unlabeled networks

- vertices have no identifiers,
- edges have port numbers

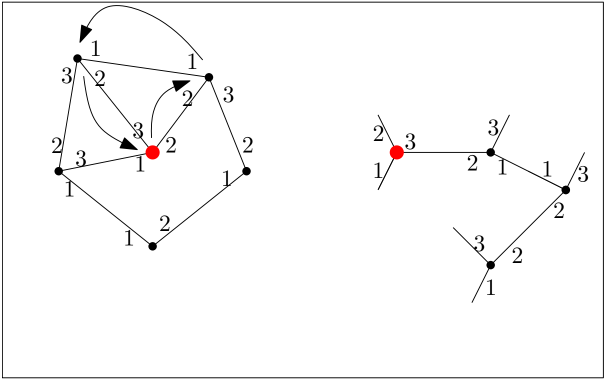

Port labeled networks — typical considerations

- feasibility of the task (symmetry breaking)

Port labeled networks — typical considerations

- feasibility of the task (symmetry breaking)
- agents with little memory (e.g., exploration by finite automaton)

Port labeled networks — typical considerations

- feasibility of the task (symmetry breaking)
- agents with little memory (e.g., exploration by finite automaton)
- communication

Port labeled networks — typical considerations

- feasibility of the task (symmetry breaking)
- agents with little memory (e.g., exploration by finite automaton)
- communication
- impact of knowledge

Port labeled networks — typical considerations

- feasibility of the task (symmetry breaking)
- agents with little memory (e.g., exploration by finite automaton)
- communication
- impact of knowledge
  - Note: size of advice for an arbitrary n-node graph is $\Theta(n \log n)$ for connected monotone edge search [Nisse & Soguet'07] ($\log n$ bits of advice are provided to vertices having *whiteboards*)

Unknown labeled graphs

- vertices have unique identifiers
- agent is able to distinguish incident edges

Exploring unknown labeled graphs

Theorem (Panaite & Pelc'99)

There exists an exploration algorithm with penalty[1] $3n$ *for any* n*-node graph.*

[1] the reference value is the number of edges

Exploring unknown labeled graphs

Theorem (Panaite & Pelc'99)

There exists an exploration algorithm with penalty[1] $3n$ *for any* n*-node graph. Note:*

- *DFS does not work (penalty* $\Omega(m)$*)*

[1] the reference value is the number of edges

Exploring unknown labeled graphs

Theorem (Panaite & Pelc'99)

There exists an exploration algorithm with penalty[1] $3n$ *for any* n*-node graph. Note:*

- *DFS does not work (penalty* $\Omega(m)$*)*
- *greedy approach does not work (penalty* $\omega(n)$*)*

[1] the reference value is the number of edges

Piecemeal exploration

- single agent

Piecemeal exploration

- single agent
- the graph is labeled (nodes are distinguishable) and unknown to the agent

Piecemeal exploration

- single agent
- the graph is labeled (nodes are distinguishable) and unknown to the agent
- agent initially located at a homebase

Piecemeal exploration

- single agent
- the graph is labeled (nodes are distinguishable) and unknown to the agent
- agent initially located at a homebase
- agent has a battery of limited size B: it needs to return to the homebase to recharge after at most B edge traversals)

Piecemeal exploration

- single agent
- the graph is labeled (nodes are distinguishable) and unknown to the agent
- agent initially located at a homebase
- agent has a battery of limited size B: it needs to return to the homebase to recharge after at most B edge traversals)
- minimize the number of trips (i.e., recharging events)
  - (closely related model to *tethered* agents)

Piecemeal exploration — offline version (complexity)

Simple reduction from 3-partition (Instance: $S = \{a_1, \ldots, a_{3m}\}$. Q.: is there a partition of S into m sets of the same sum W?)

Piecemeal exploration — offline version (complexity)

Simple reduction from 3-partition (Instance: $S = \{a_1, \ldots, a_{3m}\}$.
Q.: is there a partition of S into m sets of the same sum W?)

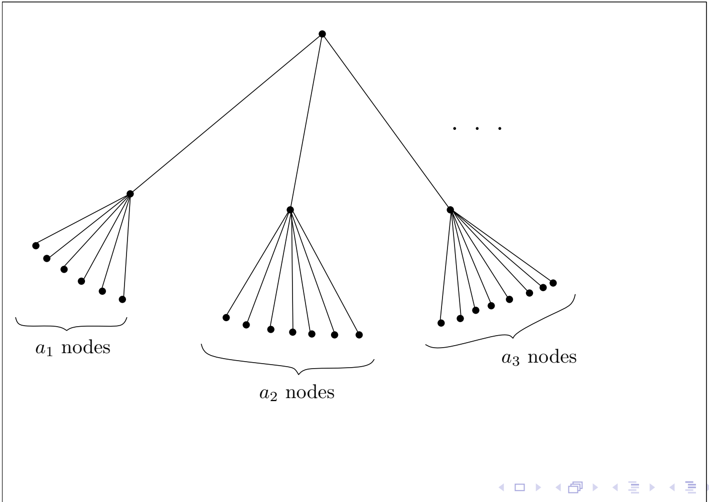

Piecemeal exploration — offline version (complexity)

Simple reduction from 3-partition (Instance: $S = \{a_1, \ldots, a_{3m}\}$. Q.: is there a partition of S into m sets of the same sum W?)

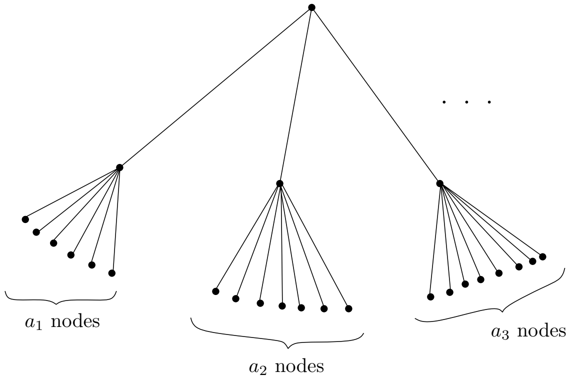

Set battery size $B := 2W + 6$.

Piecemeal exploration (short survey)

Theorem (Awerbuch et al.'99)

There exists a $O(m + n^{1+o(1)})$-time piecemeal exploration algorithm with battery size $(2+\alpha)r$ in any undirected graph, where r is the radius of the graph and $\alpha > 0$ is some constant.

Piecemeal exploration (short survey)

Theorem (Awerbuch et al.'99)

There exists a $O(m + n^{1+o(1)})$*-time piecemeal exploration algorithm with battery size* $(2+\alpha)r$ *in any undirected graph, where* r *is the radius of the graph and* $\alpha > 0$ *is some constant.*

Theorem (Awerbuch & Kobourov'98)

There exists a $O(m + n\log^2 n)$*-time piecemeal exploration algorithm with battery size* $(2+\alpha)r$ *in any undirected graph, where* r *is the radius of the graph and* $\alpha > 2$ *is some constant.*

Piecemeal exploration (short survey)

Theorem (Awerbuch et al.'99)

There exists a $O(m + n^{1+o(1)})$-time piecemeal exploration algorithm with battery size $(2+\alpha)r$ in any undirected graph, where r is the radius of the graph and $\alpha > 0$ is some constant.

Theorem (Awerbuch & Kobourov'98)

There exists a $O(m + n\log^2 n)$-time piecemeal exploration algorithm with battery size $(2+\alpha)r$ in any undirected graph, where r is the radius of the graph and $\alpha > 2$ is some constant.

Theorem (Duncan, Kobourov & Kumar'06)

There exists a $O(m/\alpha)$-time piecemeal exploration algorithm with battery size $2(1+\alpha)r$ in any undirected graph, where r is the radius of the graph and $\alpha > 0$ is some constant.

Energy constrained exploration — the model

- as before: unknown labeled graphs

Energy constrained exploration — the model

- as before: unknown labeled graphs
- multiple agents, each being able to perform at most B edge traversals

Energy constrained exploration — the model

- as before: unknown labeled graphs
- multiple agents, each being able to perform at most B edge traversals
- equivalent to piecemeal exploration when we insist that each agents needs to return to the homebase; different without this assumption

Energy constrained exploration — the model

- as before: unknown labeled graphs
- multiple agents, each being able to perform at most B edge traversals
- equivalent to piecemeal exploration when we insist that each agents needs to return to the homebase; different without this assumption
- we aim at a stronger algorithm that uses local communication

Energy constrained exploration — the model

- as before: unknown labeled graphs
- multiple agents, each being able to perform at most B edge traversals
- equivalent to piecemeal exploration when we insist that each agents needs to return to the homebase; different without this assumption
- we aim at a stronger algorithm that uses local communication
- approach that sometimes works: start with global communication and then patch your solution

Energy constrained exploration — optimization criteria

What we optimize? Two examples are:

Energy constrained exploration — optimization criteria

What we optimize? Two examples are:

- Example 1: keep the number of agents optimal but increase the battery size (the battery size becomes $(1+\alpha)B$; try to keep $\alpha > 0$ as small as possible) [our first example below; with assumption that each agent returns to the homebase]

Energy constrained exploration — optimization criteria

What we optimize? Two examples are:

- Example 1: keep the number of agents optimal but increase the battery size (the battery size becomes $(1+\alpha)B$; try to keep $\alpha > 0$ as small as possible) [our first example below; with assumption that each agent returns to the homebase]
- Example 2: keep the battery size B but increase the number of agents [our second example below; without returning to the homebase]

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

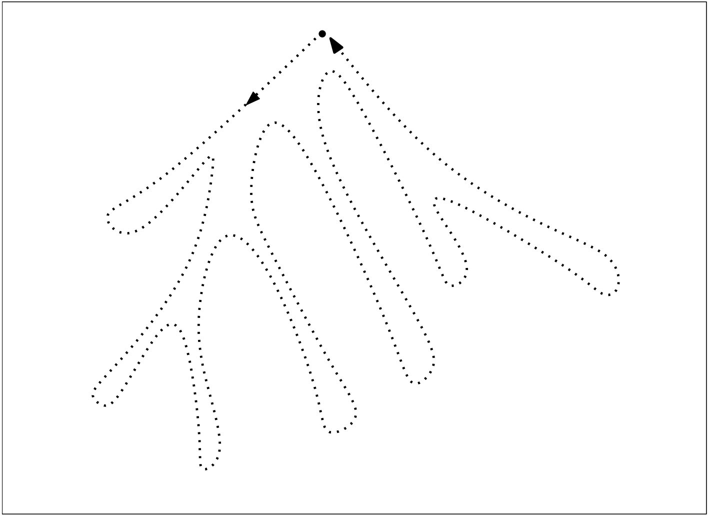

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

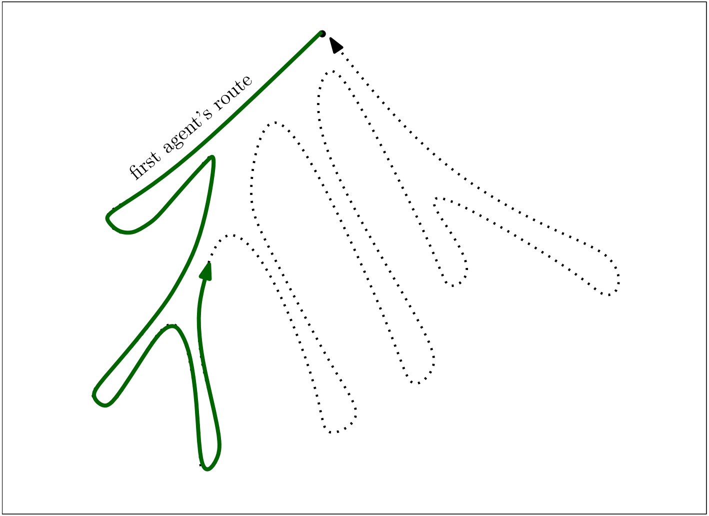

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

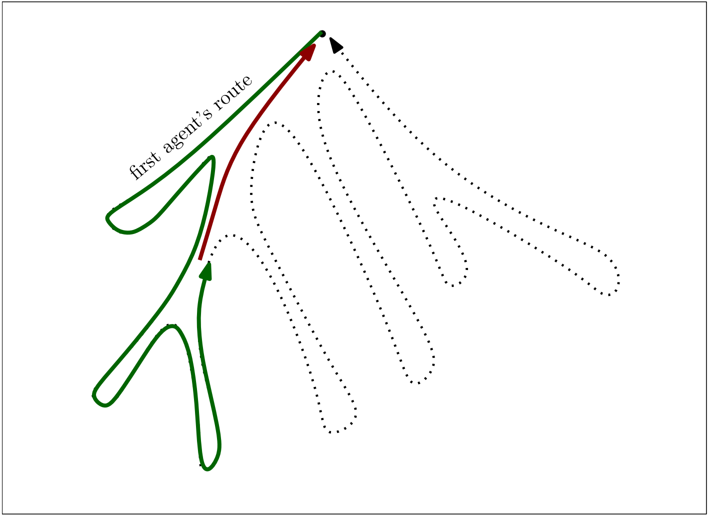

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

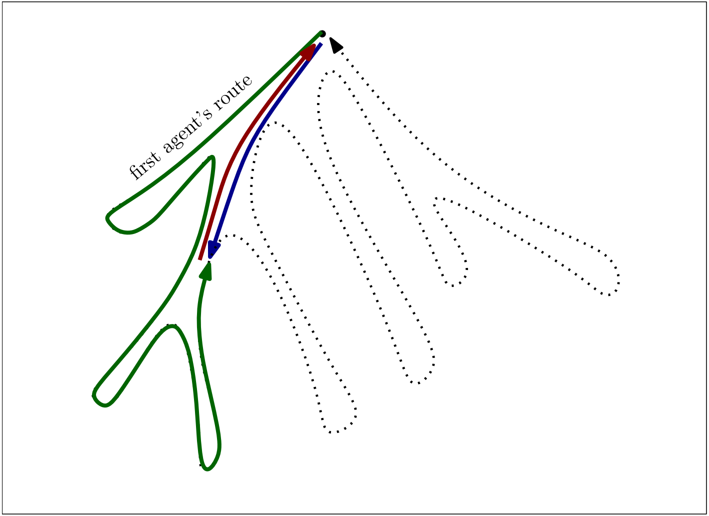

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

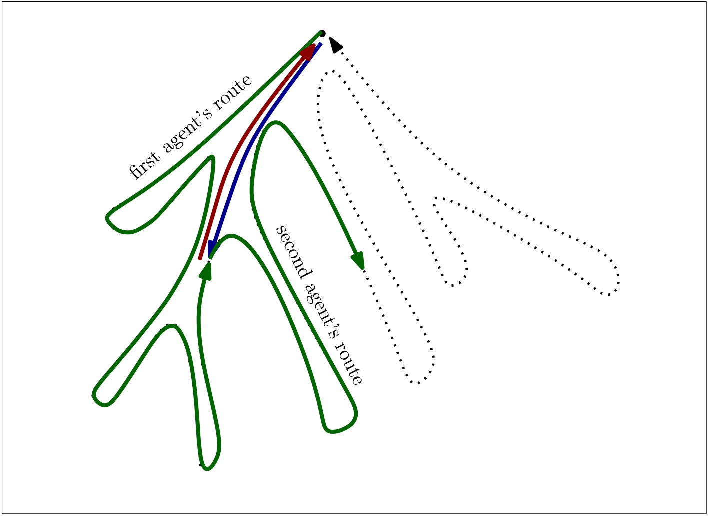

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

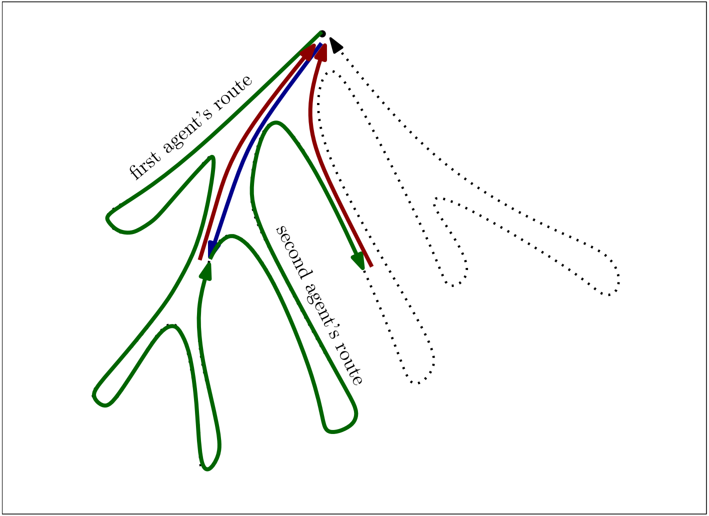

Example 1 [Dynia, Korzeniowski & Schindelhauer'06]

(Piecemeal exploration, approximate by increasing the battery size)

Theorem (Dynia, Korzeniowski & Schindelhauer'06)

There exists a 8-competitive algorithm that explores any unknown input tree by energy constrained agents using local communication.

Example 2 (global communication) [Das, D. & Karousatou'14]

Example 2 (global communication) [Das, D. & Karousatou'14]

Take some $0 < \varepsilon < \frac{1}{4}$

Example 2 (global communication) [Das, D. & Karousatou'14]

Take some $0 < \varepsilon < \frac{1}{4}$

Example 2 (global communication) [Das, D. & Karousatou'14]

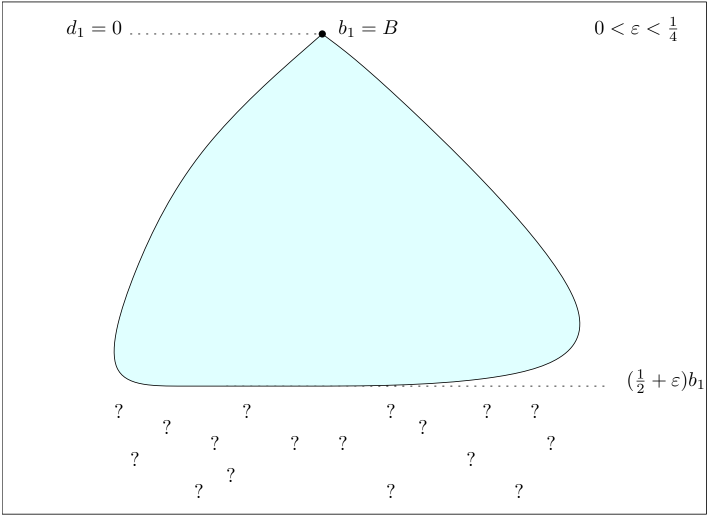

Example 2 (global communication) [Das, D. & Karousatou'14]

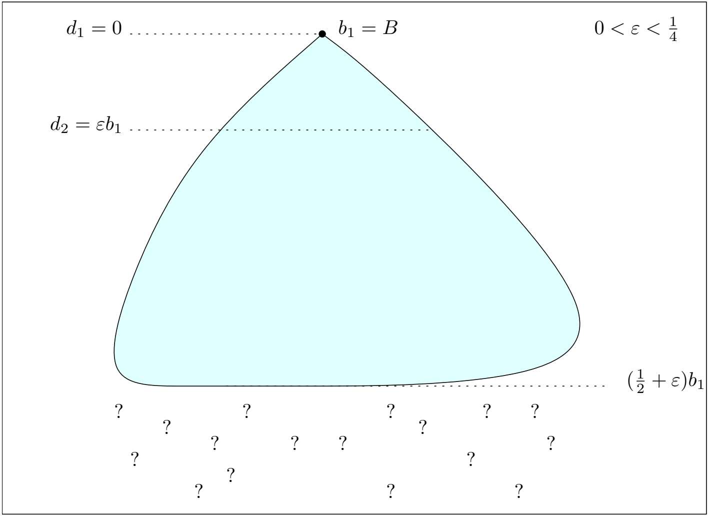

Example 2 (global communication) [Das, D. & Karousatou'14]

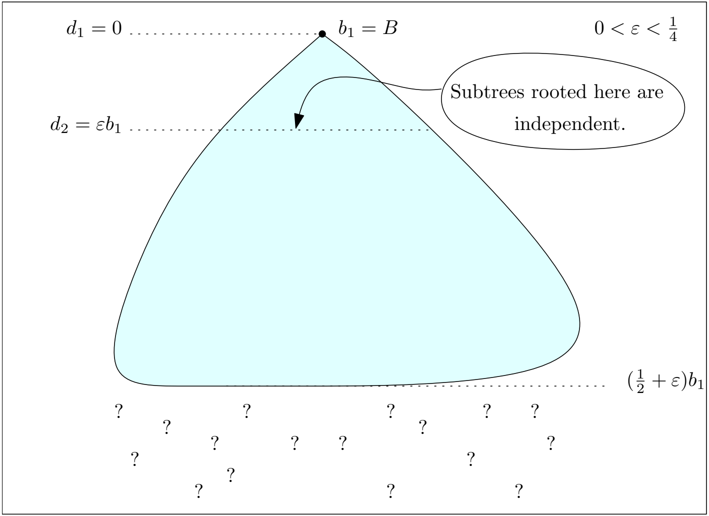

Example 2 (global communication) [Das, D. & Karousatou'14]

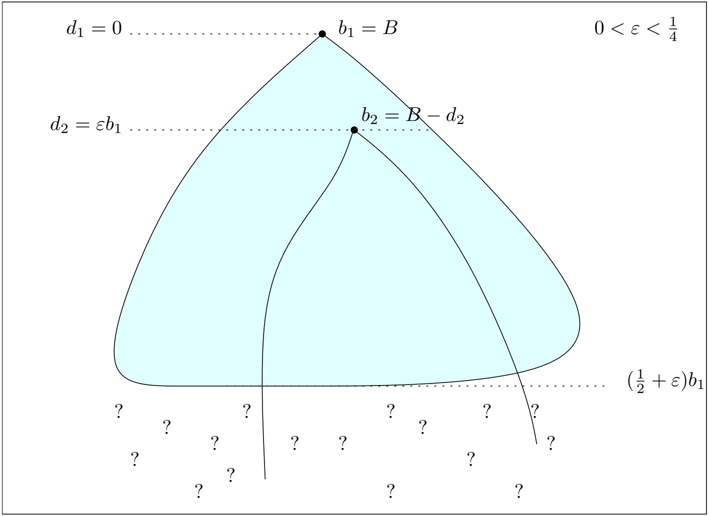

Example 2 (global communication) [Das, D. & Karousatou'14]

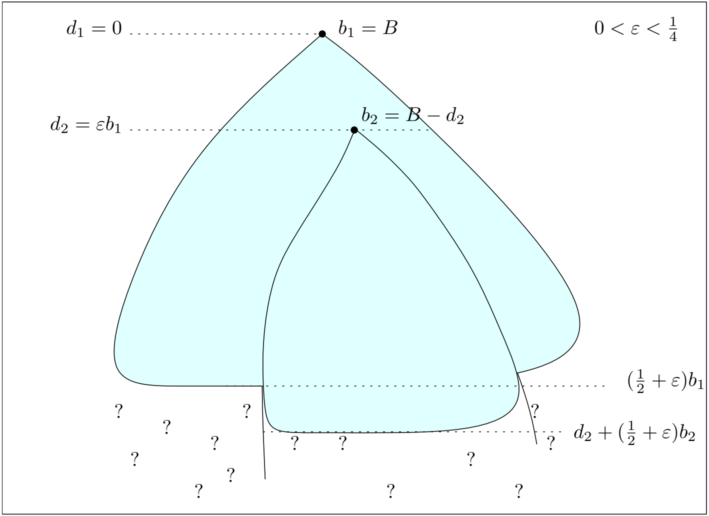

Example 2 (global communication) [Das, D. & Karousatou'14]

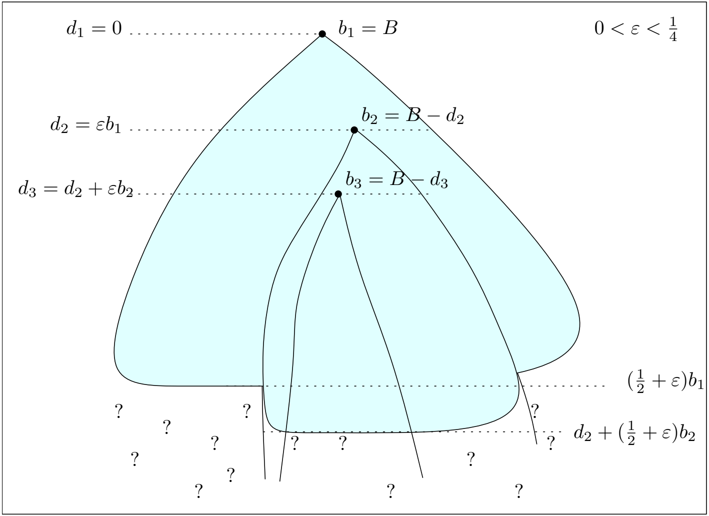

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Previous solution can be adopted for local communication:

- first modification (communication inside levels): instead of one agent performing a DFS we introduce a team (of fixed size) of agents

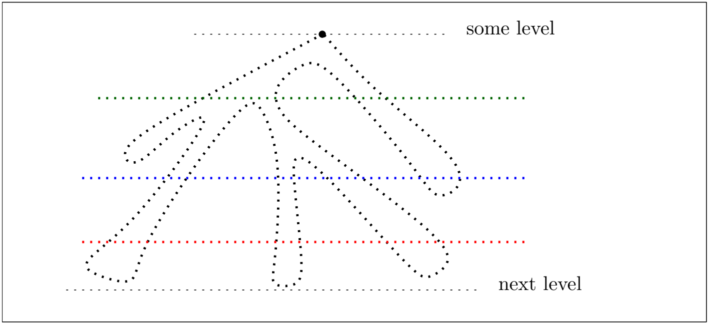

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Previous solution can be adopted for local communication:

- first modification (communication inside levels): instead of one agent performing a DFS we introduce a team (of fixed size) of agents

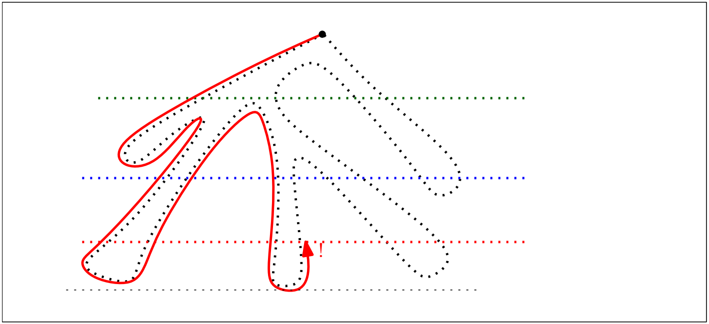

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Previous solution can be adopted for local communication:

- first modification (communication inside levels): instead of one agent performing a DFS we introduce a team (of fixed size) of agents

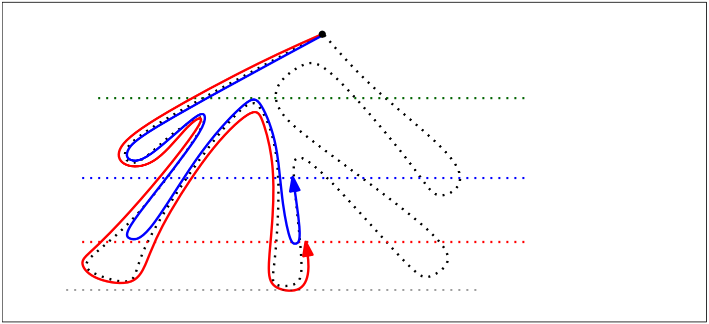

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Previous solution can be adopted for local communication:

- first modification (communication inside levels): instead of one agent performing a DFS we introduce a team (of fixed size) of agents

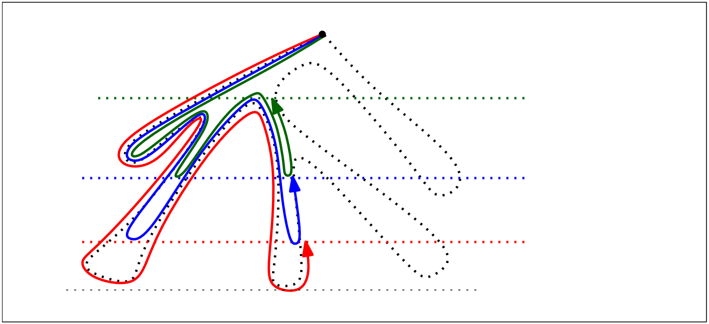

Example 2 (local communication) [Das, D. & Karousatou'14]

Note: local communication means that an agent only knows its own history and what it learned when meeting other agents.

Previous solution can be adopted for local communication:

- first modification (communication inside levels): instead of one agent performing a DFS we introduce a team (of fixed size) of agents

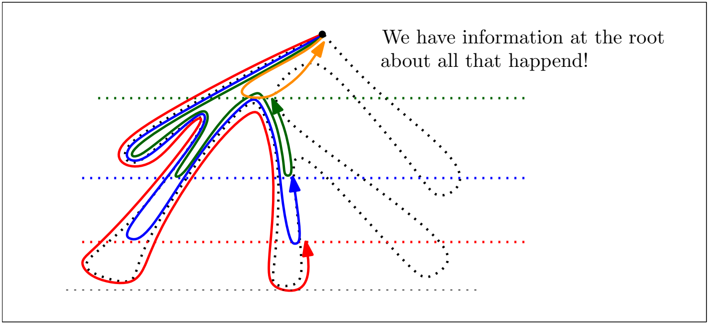

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent
  - the special agent is responsible for (1) reporting to its ancestor from previous level that exploration of the subtree rooted at v is completed, and

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent
  - the special agent is responsible for (1) reporting to its ancestor from previous level that exploration of the subtree rooted at v is completed, and
  - (2) redirecting agents coming from the root to the currently explored subtree

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent
  - the special agent is responsible for (1) reporting to its ancestor from previous level that exploration of the subtree rooted at v is completed, and
  - (2) redirecting agents coming from the root to the currently explored subtree
  - agents use time to communicate as follows

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent
  - the special agent is responsible for (1) reporting to its ancestor from previous level that exploration of the subtree rooted at v is completed, and
  - (2) redirecting agents coming from the root to the currently explored subtree
  - agents use time to communicate as follows
  - every few time steps ($\Theta(B)$ steps) a new agent 'appears' at the root which is then redirected through consecutive levels down the tree by the special agents.

Example 2 (local communication) [Das, D. & Karousatou'14]

Previous solution can be adopted for local communication:

- second modification (communication between levels):
  - at each node v of a given level there is a special agent
  - the special agent is responsible for (1) reporting to its ancestor from previous level that exploration of the subtree rooted at v is completed, and
  - (2) redirecting agents coming from the root to the currently explored subtree
  - agents use time to communicate as follows
  - every few time steps ($\Theta(B)$ steps) a new agent 'appears' at the root which is then redirected through consecutive levels down the tree by the special agents.

Theorem (Das, D. & Karousatou'14)

There exists a $O(\log B)$-competitive algorithm that explores any unknown input tree by energy constrained agents using local communication.

A lower bound for energy constrained exploration

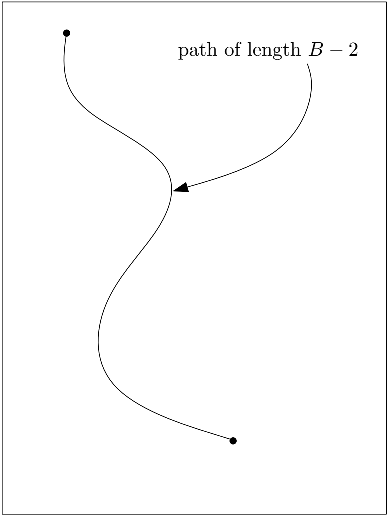

A lower bound for energy constrained exploration

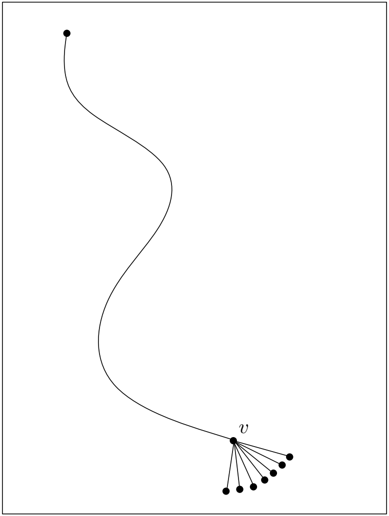

A lower bound for energy constrained exploration

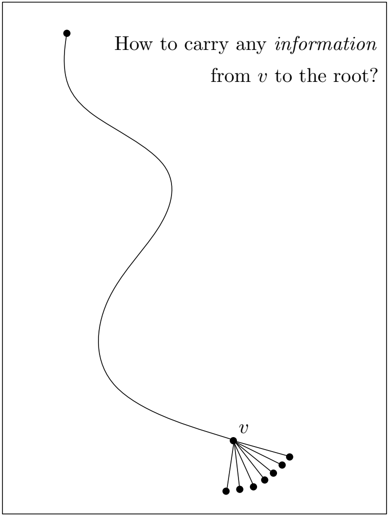

A lower bound for energy constrained exploration

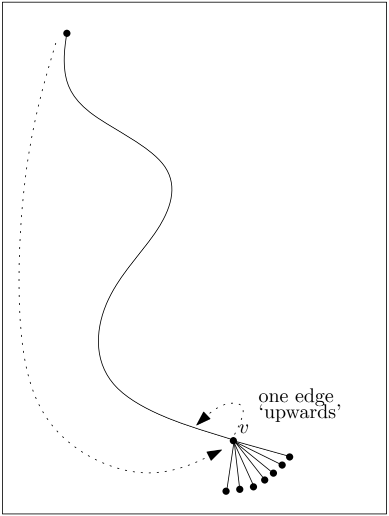

A lower bound for energy constrained exploration

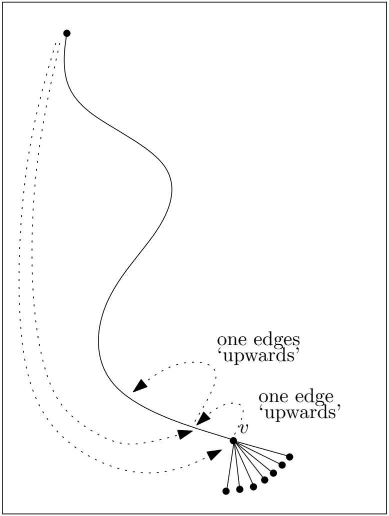

A lower bound for energy constrained exploration

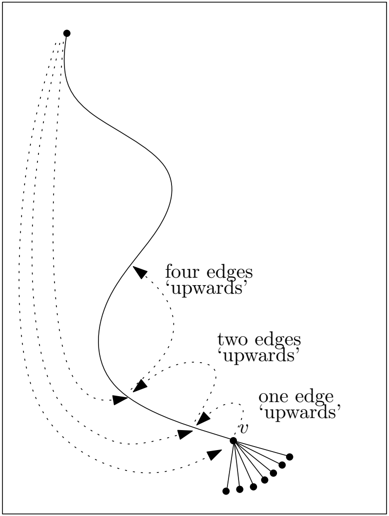

A lower bound for energy constrained exploration

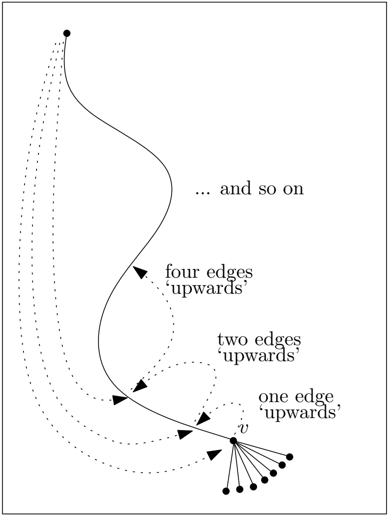

- we need $\log B$ agents to carry *information* from v to the root

A lower bound for energy constrained exploration

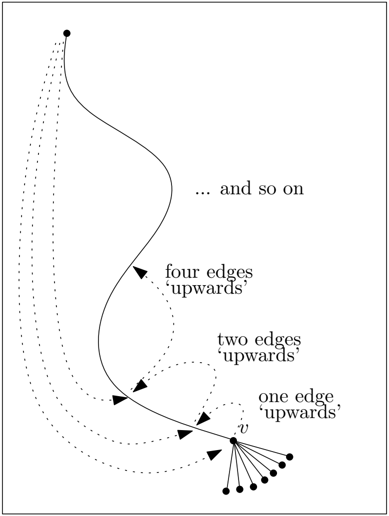

- we need $\log B$ agents to carry *information* from v to the root
- this information is the number of leaves

A lower bound for energy constrained exploration

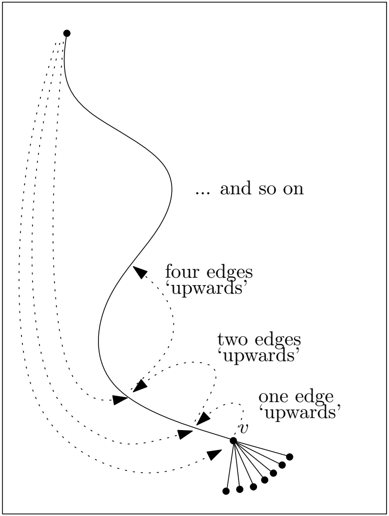

- we need $\log B$ agents to carry *information* from v to the root
- this information is the number of leaves
- Claim: $\Omega(\log B)$ agents are necessary for some graphs.

Exploration time

- a team of k robots start at the root of a tree
- the goal is to explore the tree

Unknown tree exploration — a short survey

- time $O(D + n/\log k)$ using whiteboards at nodes [Fraigniaud et al.'06]
  - this gives competitive ratio of $O(k/\log k)$ w.r. offline optimal $O(D + n/k)$
  - intermediate algorithmic step via global communication

Unknown tree exploration — a short survey

- time $O(D + n/\log k)$ using whiteboards at nodes [Fraigniaud et al.'06]
  - this gives competitive ratio of $O(k/\log k)$ w.r. offline optimal $O(D + n/k)$
  - intermediate algorithmic step via global communication
- time $O(n/k + D^{k-1})$ using whiteboards at nodes [Brass & Cabrera-Mora'11]
  - improvement only for small diameter and $k = O(\log_D n)$

Unknown tree exploration — a short survey

- time $O(D + n/\log k)$ using whiteboards at nodes [Fraigniaud et al.'06]
  - this gives competitive ratio of $O(k/\log k)$ w.r. offline optimal $O(D + n/k)$
  - intermediate algorithmic step via global communication
- time $O(n/k + D^{k-1})$ using whiteboards at nodes [Brass & Cabrera-Mora'11]
  - improvement only for small diameter and $k = O(\log_D n)$
- exploration in time $O(D)$ with a polynomial number of agents [D. et al.'13]
  - time $D(1 + \frac{1}{c-1} + o(1))$ using Dn^c for any $c > 1$; global communication (Example 3)
  - time $D(1 + \frac{2}{c-1} + o(1))$ using Dn^c for any $c > 1$; local communication (Example 4)
  - time $O(D \log n)$ using $k = (2 + \varepsilon)nD$ agents and local communication in general graph, for any $\varepsilon > 0$

Example 3: Fast tree exploration (global communication)

- The goal: exploration in time $O(D)$ with polynomial number of agents

Example 3: Fast tree exploration (global communication)

- The goal: exploration in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

Example 3: Fast tree exploration (global communication)

- The goal: exploration in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

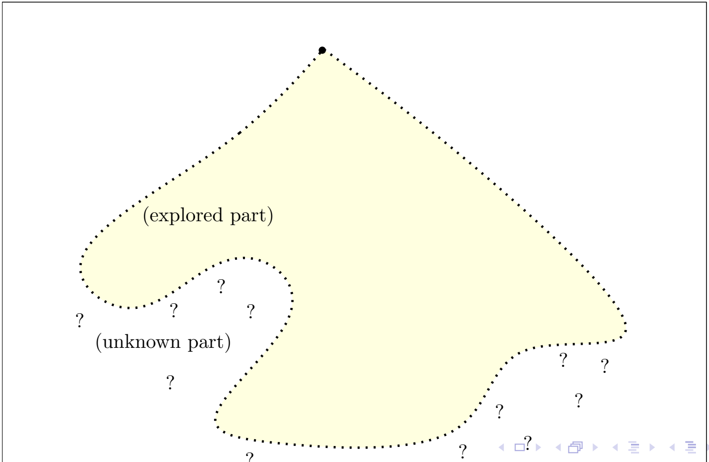

Example 3: Fast tree exploration (global communication)

- The goal: exploration in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

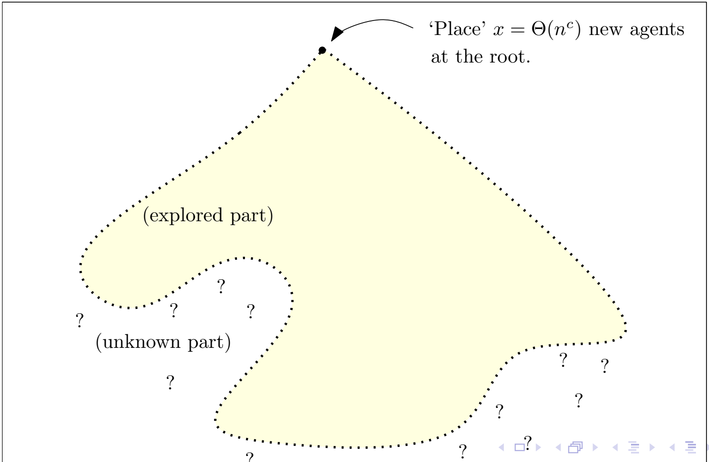

Example 3: Fast tree exploration (global communication)

- The goal: exploration in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

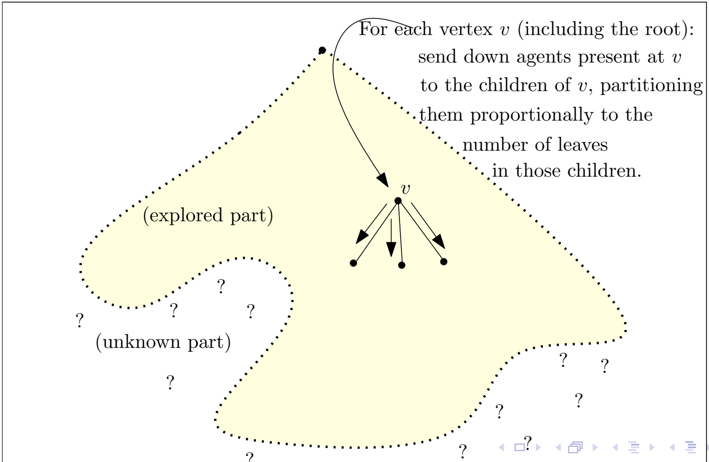

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

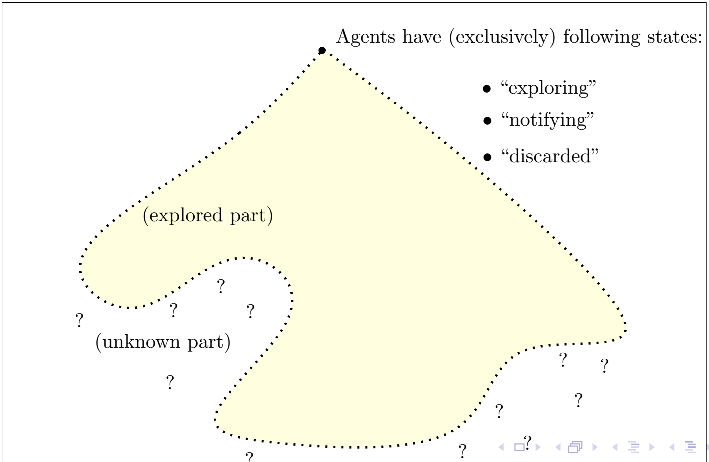

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

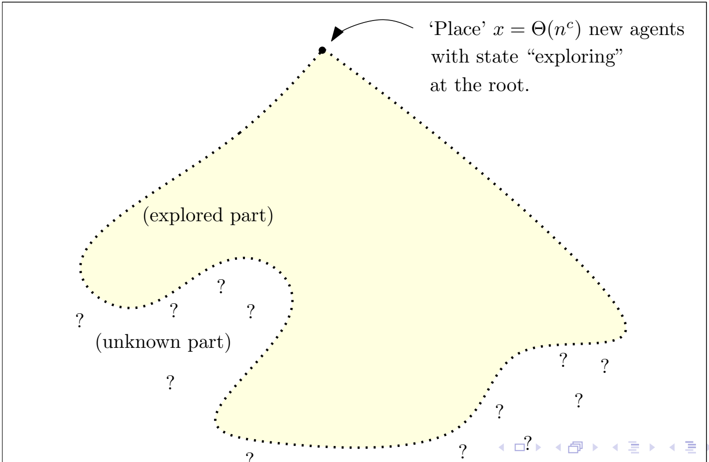

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

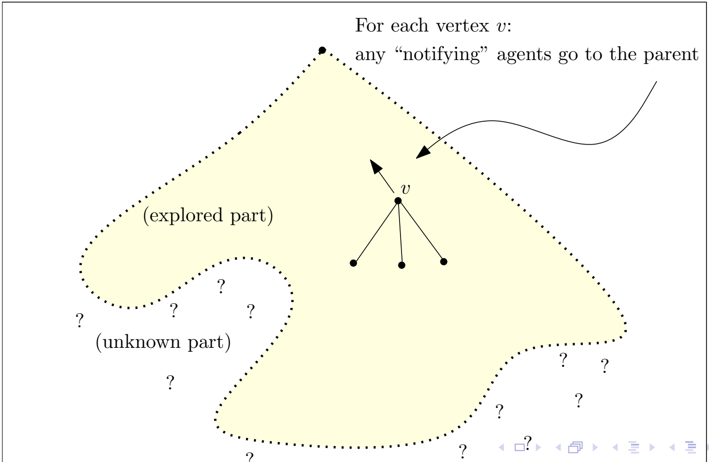

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

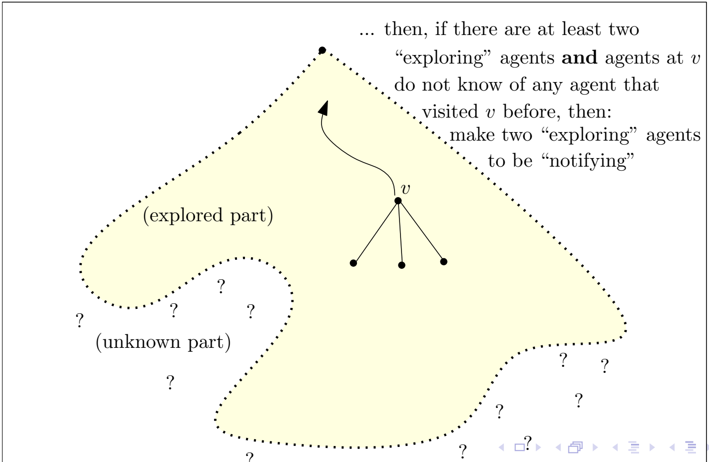

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

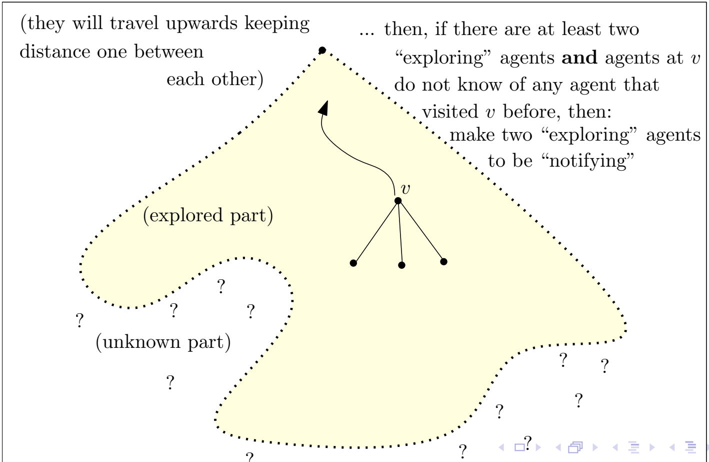

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

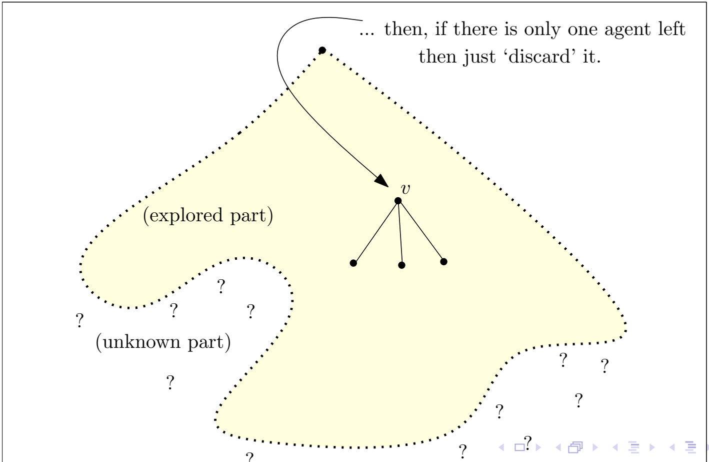

Example 4: Fast tree exploration (local communication)

- The goal: exploration also in time $O(D)$ with polynomial number of agents
- Single *step* is as follows:

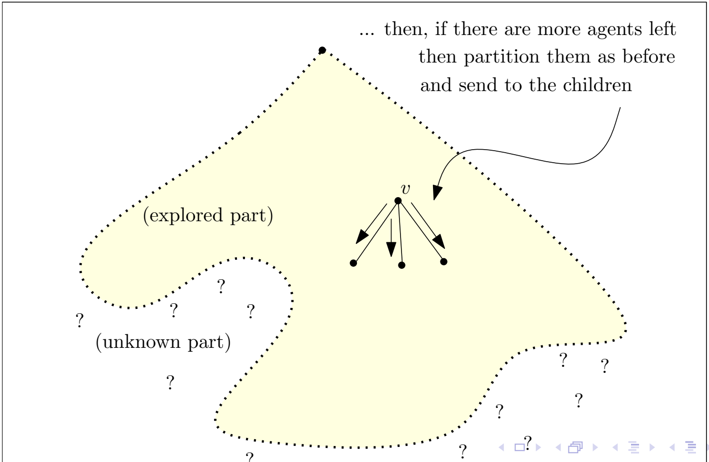

Fast graph exploration

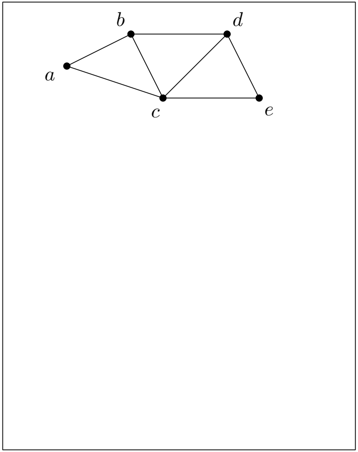

Fast graph exploration

- we simulate exploration of G by exploring a 'virtual' tree $\mathcal{T}$

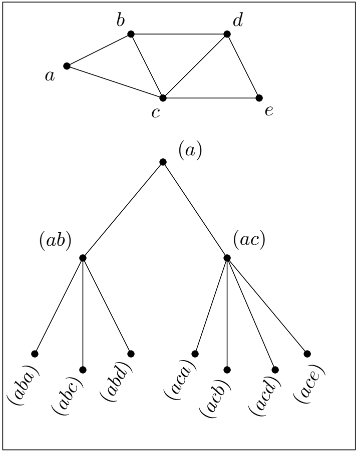

Fast graph exploration

- we simulate exploration of G by exploring a 'virtual' tree $\mathcal{T}$
- one virtual move in $\mathcal{T}$ gives one step in G

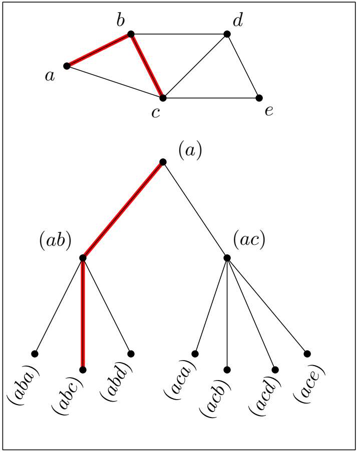

Fast graph exploration

- we simulate exploration of G by exploring a ‘virtual’ tree $\mathcal{T}$
- one virtual move in $\mathcal{T}$ gives one step in G
- agent placed on P in $\mathcal{T}$ is present at the end vertex of P in G

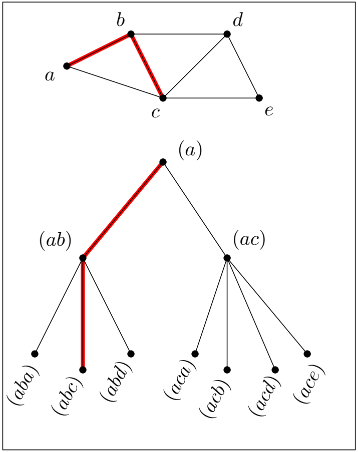

Fast graph exploration

- we simulate exploration of G by exploring a 'virtual' tree $\mathcal{T}$
- one virtual move in $\mathcal{T}$ gives one step in G
- agent placed on P in $\mathcal{T}$ is present at the end vertex of P in G
- the size of $\mathcal{T}$ is exponential but it is enough to explore a polynomial-size subtree

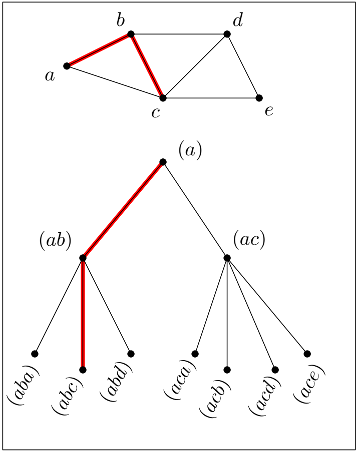

Thank you!