Almost-Optimal Deterministic Treasure Hunt in Arbitrary Graphs

Sébastien Bouchard¹, Yoann Dieudonné², Arnaud Labourel³ and Andrzej Pelc⁴

¹ LIP6, Paris, France
² MIS Lab., Université de Picardie Jules Verne, Amiens, France
³ Aix Marseille Univ, CNRS, LIS, Marseille, France
⁴ Département d'informatique, Université du Québec en Outaouais, Gatineau, Canada.



Arnaud Labourel (AMU, CNRS, LIS)

Almost-Optimal Treasure Hunt in Graphs

















Treasure hunt in graphs

The problem

A mobile agent has to find a treasure inside a network.

- Ports corresponding to edges incident to a node of degree δ are numbered $0, 1, \ldots, \delta 1$.
- The agent sees the degree and identifier of its current node.
- At each step the agent can choose a port number, move through the corresponding edge and see the incoming port number.

Benchmarking treasure hunt algorithms

Cost of a treasure hunt algorithm : number of moves performed by the agent before finding the treasure in the worst case.

d: initial distance between the agent and the treasure e(r): number of edges whose at least one extremity is at distance $\leq r - 1$ from the starting position of the agent

Lower bound for treasure hunt

The cost of any treasure hunt algorithm is at least e(d).

Idea of the lower bound



Idea of the lower bound



Idea of the lower bound



Arnaud Labourel (AMU, CNRS, LIS)

Related works on general graphs

e(r) : number of edges whose at least one extremity is at distance $\leq \textit{r}-1$

n(r): number of nodes at distance $\leq r$

Awerbuch, Betke, Rivest and Singh 1999

Treasure hunt with cost $O(e(d + o(d)) + n(d + o(d))^{1+o(1)})$

Duncan, Kobourov and Kuma 2006 Treasure hunt with cost $O(e((1 + \epsilon)d) + n((1 + \epsilon)d)/\epsilon)$ for $0 < \epsilon < 1$.

Almost optimal treasure hunt

Our result (ICALP 2021)

Treasure hunt on general graphs with cost $O(e(d) \log d)$.

Since $d \le e(d)$ the cost only differs from the best possible cost by a logarithmic factor. Our result refutes the following conjecture :

Conjecture Awerbuch, Betke, Rivest, Singh 1999 Is it possible (we conjecture not) to find a treasure in time nearly linear in the number of those vertices and edges whose distance to the source is less than or equal to that of the treasure?

Restricted models

Rope-restricted model

- The agent is attached to its starting position by a rope of length *L*
- The rope unwinds by a length 1 with every forward edge traversal and rewinds by a length of 1 with every backward edge traversal.

Fuel-restricted model

- The agent has a fuel tank of capacity *B*.
- Each move consumes one unit of fuel.
- The agent can fully replenish its tank at its starting position.

r: radius of the graph

For any constant ϵ , our treasure hunt algorithm can be transformed into:

- an algorithm for rope-restricted model with a rope of length $(1+\epsilon)\mathbf{r}$
- an algorithm for fuel-restricted model with a fuel tank of capacity $2(1+\epsilon)\mathbf{r}$

Both algorithms have cost $O(e(d) \log d)$.

Easy case: trees with knowledge

For now, we assume that:

- the graph is a tree,
- we know the values e(i) for every *i*.

Idea of the algorithm

Iteratively explore balls that grows exponentially.

Algorithm

• Compute values *h_i* such that:

• $h_0 = 1$

- $h_{i+1} = \min\{h \mid e(h) \ge 2e(h_i)\}$ for all $i \ge 0$
- Iteratively explore balls of radius h_i for i = 0, 1, 2, ...











Trees with knowledge (second attempt)

Problem with the first attempt

If the balls grow too fast $e(h_{i+1}) >> e(h_i)$ and the difference of their radius is ≥ 2 then the treasure can be hidden at distance $h_{i+1} - 1$ and the cost is $\Omega(e(h_{i+1})) >> e(h_{i+1} - 1)$.

 \Rightarrow Add an exploration at depth $h_{i+1} - 1$ in this case.

New algorithm: pick in order for h_{i+1}

) Pick
$$h_i + 1$$
 if $e(h_i + 1) \ge 2e(h_i)$

Pick (if it exists) the minimum radius h such that $2e(h_i) \le e(h) \le 4e(h_i)$

2 Pick the maximum radius h such that $e(h) < 2e(h_i)$

Trees with knowledge (second attempt)



Cost of the new algorithm

The treasure is found at cost 16e(d).

Key ideas of the proof :

- The cost of exploring the last ball (the one in which the treasure is found) is at most four times e(d).
- The size of the explored balls grows exponentially (either $\times 2$ in one iteration of the loop or $\times 4$ in two iterations), hence the total cost is proportional to the cost of the exploration of the last ball.

Idea of the algorithm

Assuming that the ball of radius h_i has been explored. Search the next radius h_{i+1} by trying to explore balls of radius between $h_i + 1$ and $2h_i$ with a binary search.

Exploration will be done with a Budgeted DFS that tries to explore the ball of radius h but will abort and backtrack if the number of explored edges is $> 4e(h_i)$ and $h > h_i + 1$.







Arnaud Labourel (AMU, CNRS, LIS)





Cost of treasure hunt for trees

Cost of the new algorithm The treasure is found at cost $O(\log(d).e(d))$.

- We count the cost of explorations with doubling radius separately: there are at most log d such explorations and the cost of each one is O(e(d)) and so their total cost is O(log(d).e(d))
- Other *e*(*h_i*) grows exponentially, hence the total cost is proportional to the cost of the last iteration.
- The cost of the last iteration is O(log(d).e(d)) because there are at most log(d) tries of explorations at each iteration and each one has a cost in O(e(d)).

General case: arbitrary graph

We assume by induction that we already have explored the ball of radius h_i and we want to explore all nodes at distance $\leq h_i + l$.

We maintain:

A list of incomplete nodes, i.e., node at distance
h_i + *l* with incident unexplored edges

• A tree containing all the incomplete nodes The agent will visit each incomplete node u in some DFS order of the tree and explore all unexplored nodes at distance $\leq l$ of u.

Like with the treasure-hunt on trees, the exploration will have a budget and abort if there are too many edges.

Spanning tree and incomplete nodes















Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Unlike the case of trees, it is difficult to explore nodes at distance *d* since the exact distance is not known by the agent.



Solution: pruning algorithm

We use the pruning technique of Duncan et al. 2006.

Rough idea of the algorithm

Repeat the following steps until there are no more incomplete node:

- Move to an incomplete node u
- Prune the subtree rooted in u by creating subtrees
- Explore at depth I/2 from every node of the subtree containing u
- Opdate the list of incomplete nodes
- Update the subtrees by removing completed ones and merging those sharing a node.

backtrack to u

Pruning algorithm

Properties of the pruning

- the maximum depth on the original pruned is I/2
- ullet no pruned subtree has depth less than I/4









locally explore from the first incomplete node all nodes at distance l





prune the tree, remove any complete tree





Rough idea of the analysis

Three kinds of edge traversal :

- moves between pruned subtrees
- moves inside pruned subtrees (in O(e(h_i + I)) since the pruned subtrees are disjoint)
- new edge exploration (≤ 2e(h_i + I) since each of these edge is traversed once in each direction)

One of the key idea is that each pruned subtree has a size $\geq I/4$ and thus the moves between subtrees (O(I) in order to move from one subtree to another one) is proportional to the number of moves inside the subtrees.

Complete code of our algorithm

Algorithm 2 lreasureHunt(x)	Algorithm 4 GlobalExpanzion(l, m)	Algorithm 5 CDFS(l, b)
$\label{eq:constraints} \begin{array}{l} r := the current node; \\ W := \{(r), 0\} / r & M \text{ is a global variable} \\ repeat \\ \textbf{Search(r);} \\ \hline \textbf{Bowth 3 Facul(r)} \\ we usual the current node; m:= [M]; \\ for = -i.cd(r) = oi = (1 + a); for $	$ \begin{array}{l} \mathbf{z} := \operatorname{bet} \operatorname{correct} \operatorname{and}_{\mathbb{C}} \\ \mathbf{z} := \operatorname{bet} \operatorname{args} \operatorname{correct}_{\mathbb{C}} \operatorname{d} \operatorname{d} \operatorname{borndow}_{\mathbb{C}} \operatorname{d}^{d} \operatorname{borndow}_{\mathbb{C}} \\ \begin{array}{l} \mathrm{d} \operatorname{borndow}_{\mathbb{C}} \operatorname{d}^{d} \operatorname{borndow}_{\mathbb{C}} \\ \mathrm{d} \operatorname{borndow}_{\mathbb{C}} \operatorname{d}^{d} \operatorname{borndow}_{\mathbb{C}} \\ \mathrm{d} \operatorname{borndow}_{\mathbb{C}} \left(\ldots \right) \\ \mathrm{d} \operatorname{borndow}_{\mathbb{C}} \left(\operatorname{borndow}_{\mathbb{C}} \left(\ldots \right) \\ \mathrm{d} bor$	$ \begin{array}{lll} r := her current and c T := (\{r\}, \emptyset\}, hourd := h; \\ z \ f (I > 0, hourd := h; \\ while, node v; \\ while, is deverable of the M and hourd \geq 0 do while, is the number for gost at node v in \mathcal{M}_i if I > 0 for a summary for the part by which the agent has just entered node w; \\ p_{1,i} = her part by which the agent has just entered node w; \\ p_{1,i} = her part by which the agent has just entered node w; \\ p_{1,i} = her part by which the agent has just entered node w; \\ p_{1,i} = her (v, w), ((v, w, t_{i}, v_{i}, t_{i}))); \\ p_{1,i} = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i}))); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i}))); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i}))); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i}))); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, v_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} K^{i} = ((v, w), ((v, w, t_{i}, t_{i})); \\ f = \int_{i}^{i} $
Algorithm 7 Prune(l)	Algorithm 6 LocalExpansion	(l,b)
2 T _v := the tree of T containing node v;	1 bound := b: v := the current	node:
a $T := T \setminus \{T_v\};$	2 if v is incomplete in M and	d no tree of T contains node v then
4 Root T _v at node v:	2 II o is incomplete in 541 and	a no acc of f comains node c chen

κ.

6

7

8

9

10

12

13

- 6 T_u := the subtree of T_v rooted at u
- 7 if $e_{T_u}(u) \ge \lfloor \frac{t}{4} \rfloor 1$ then
- s $T := T \cup \{T_u\};$
- 9 Remove from T_v all nodes that belong to T_u and all edges that are incident to a node of T_u;
- 10 $T := T \cup \{T_v\};$

Algorithm 8 Explore(l, b)

- 1 bound := b; i := 1; v := the current node;
- 2 T := the tree of T containing node v;
- 3 V := array containing all the nodes of T sorted in the order of the first visit through 11 the DFS traversal of T from node v;
- 4 while $i \le |V|$ and bound ≥ 0 do
- 5 MoveTo(T, V[i]);
- if node V[i] is incomplete in M then
- 7 $(bound, T') := CDFS(\lfloor \frac{1}{2} \rfloor, bound);$

s
$$T := T \cup \{T'\}$$

9 return bound;

14 return bound:

 $MoveTo(\mathcal{M}, u)$:

bound := Explore(l, bound);

 $\mathcal{T} := (\mathcal{T} \setminus \{T, T'\}) \cup \{T''\};$

Prune(l):

4 while $IncompleteNodes(v, M, l) \cap Nodes(T) \neq \emptyset$ and $bound \ge 0$ do

node having the smallest label in $T \sqcup T'$;

beginning of the current iteration of the while loop;

u := the node with the smallest label in IncompleteNodes $(v, \mathcal{M}, l) \cap Nodes(\mathcal{T})$:

T'' := the spanning tree produced by the BFS traversal of $T \sqcup T'$ from the

Execute in the reverse order all the edge traversals that have been made since the

Remove from T every tree for which all the nodes are complete in M;

while there are two trees T and T' in T having a common node do

Summary of results and open problem

Our result (ICALP 2021)

Treasure hunt on general graphs with cost $O(e(d) \log d)$.

Open problem

Is there a treasure hunt algorithm on general graphs with cost O(e(d)) (asymptotically optimal) ?



Thanks for your attention !

