Searching a tree with signals: routing mobile sensors for targets emitting radiation, chemicals or scents

Steve Alpern (Warwick Business School, UK) and Thomas Lidbetter (Rutgers Business School, US)

Thursday 26th October 2023

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1935826.
Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Background: Isaac’s problem and Gal’s solution: Hide and seek on a network

- Rufus Isaacs (1965) Differential Games
- Shmuel Gal (1979) Search Games with Mobile and Immobile Hider
- Shmuel Gal (2000) On the Optimality of a Simple Strategy for Searching Graphs
- Steve Alpern (2011) A new approach to Gal's Theory of Search Games on Weakly Eulerian networks

Search for Immobile Hider on a Network

- Every arc of a network Q has a length
- Total length of Q is μ
- Distance function d on Q is the "shortest path" metric

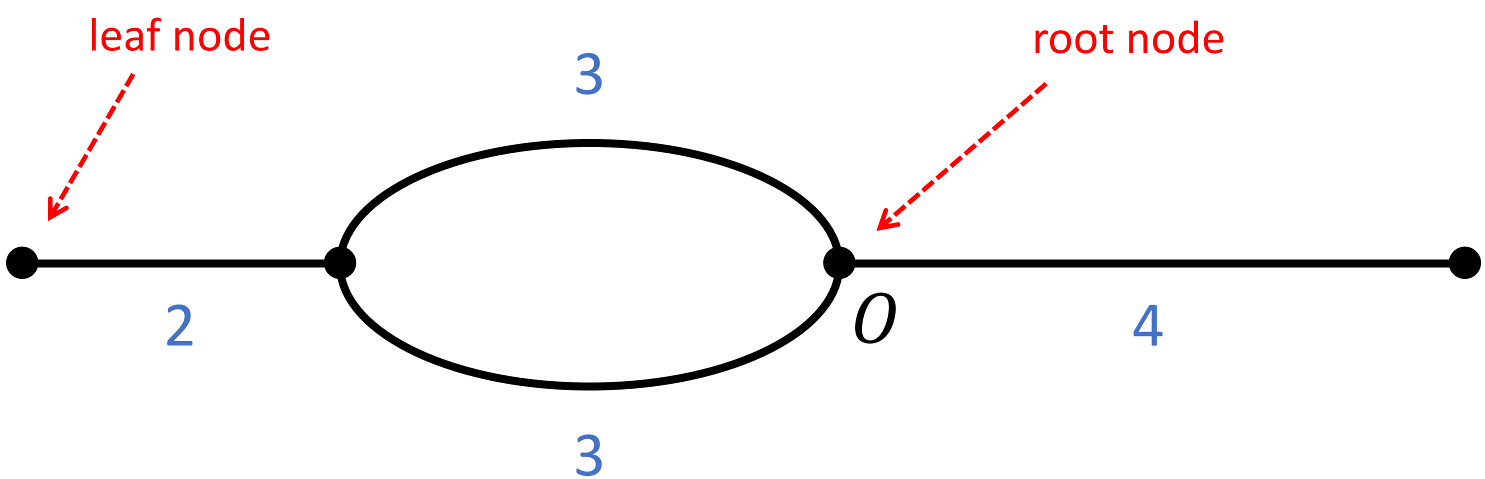

The game

- A strategy for Hider (maximizer): a point in Q (not necessarily a node)
- Mixed strategy h for Hider is a distribution over Q
- A strategy for Searcher (minimizer) is a unit speed path $S(t), t \geq 0$ which covers Q. (Unit speed $\Leftrightarrow$ $d\big(S(t_1), S(t_2)\big) \leq t_2 - t_1$ for $t_2 \geq t_1 \geq 0$.)
- Mixed strategy for the Searcher is a probability distribution over such paths
- The payoff is the *search time* $T = T(S, H) = \min\{t: S(t) = H\}$. For mixed strategies s and h, write $T(s, h)$ for the *expected search time*.
- The game has a value $V = V(Q, O)$, optimal (min-max) mixed Searcher strategies and ε-optimal (max-min) mixed Hider strategies.
- I.e. there is a number V such that the Searcher has a mixed strategy that guarantees the expected payoff is at most V whatever the Hider does and the Hider has a mixed strategy that guarantees the expected payoff is at least V whatever the Searcher does.

Optimal Searcher strategy for trees

Lemma: Let S be any depth-first tour of a tree Q with root O and let S_r be the reverse tour. Let s be the search that chooses S and S_r with equal probability. Then for any $H \in Q, T(s, H) \leq \mu$. Hence $V \leq \mu$.

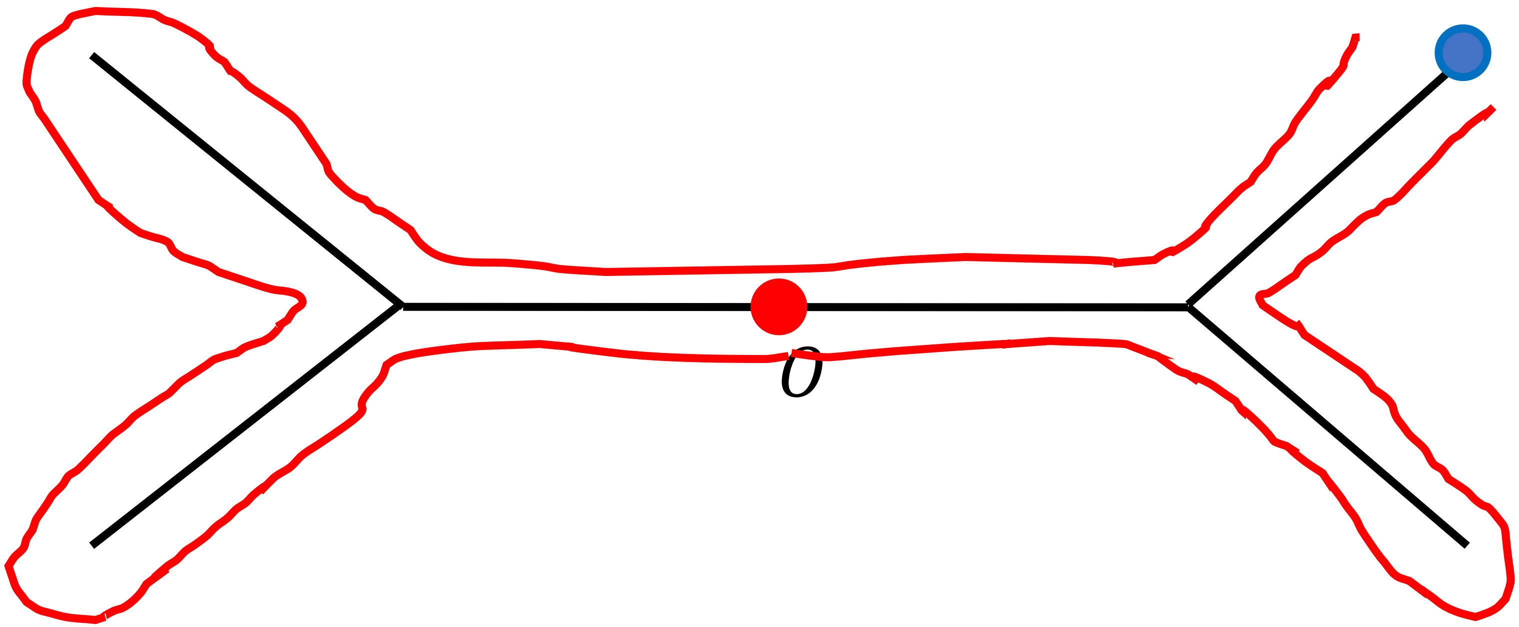

Optimal Searcher strategy for trees

Lemma: Let S be any depth-first tour of a tree Q with root O and let S_r be the reverse tour. Let s be the tour that chooses S and S_r with equal probability. Then for any $H \in Q, T(s, H) \leq \mu$. Hence $V \leq \mu$.

Proof: Let t be such that $S(t) = H$. Then $T(S_r, H) \leq 2\mu - t$. So

$$T(s, H) = \frac{1}{2} T(S, H) + \frac{1}{2} T(S^r, H) \leq \frac{1}{2} t + \frac{1}{2} (2\mu - t) = \mu.$$

Equal Branch Density (EBD) Hider Distribution for Trees

Definition: The EBD Hider distribution is concentrated on the leaf nodes and at every branch node the ratio of hiding probability to branch length (*density*) of each branch is equal.

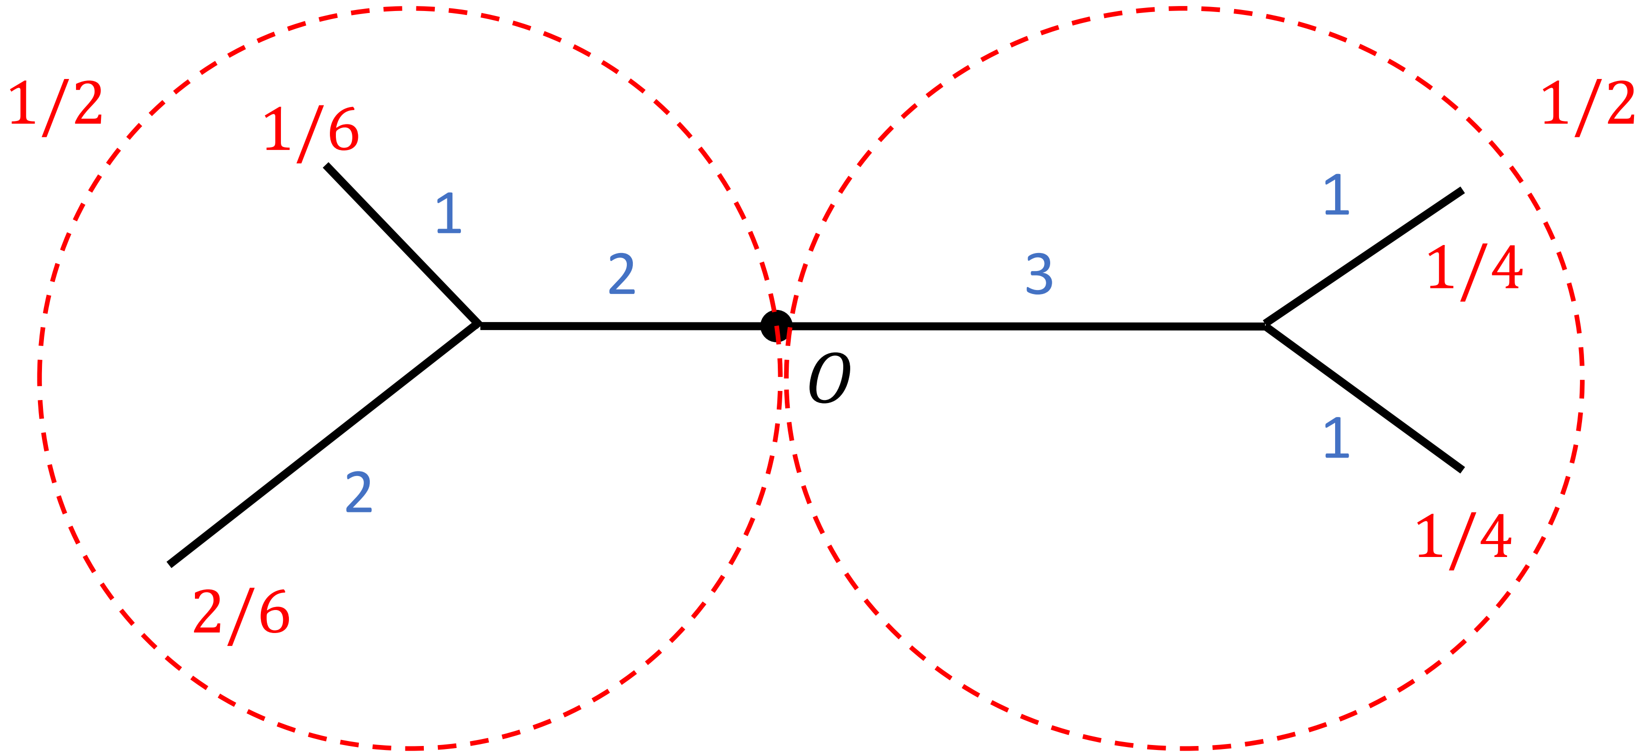

Depth-first search is a best response against the EBD

Lemma: Any depth-first search S is a best response against the EBD distribution, h and has expected search time $T(S,h) = \mu$.

Proof sketch

(i) Any two depth-first searches S_1 and S_2 have the same expected search time because S_1 can be transformed into S_2 by successively swapping the order of search of equal density subtrees that share a root.

(ii) If S is any depth-first search and S^r is its time reverse search then for any leaf node v,

$$T(S,v) + T(S^r,v) = 2\mu,$$

so

$$T(S,h) + T(S^r,h) = 2\mu,$$

so

$$T(S,h) = T(S^r,h) = \mu.$$

(iii) Proof by contradiction that any depth-first search is a best response.

$V = \mu$ for trees

Theorem: Let Q be a tree with root O. Then $V = \mu$.

Proof:
(i) $V \leq \mu$ (Searcher uses equiprobable mixture of a DF tour and its reverse)
(ii) $V \geq \mu$ (Hider uses EBD distribution)

Our work: searching a tree with signals

Modifications to original game:

- Now the Searcher is *assumed* to perform a depth-first search.
- The tree is assumed to be binary.
- Every time the Searcher reaches a branch node, she receives a signal indicating which branch the Hider is in.
- If the Hider is in one of the branches, the signal is correct with probability $p \geq 1/2$.
- If the Hider isn't in either branch, the signal can be anything.
- If $p = 1/2$, this is equivalent to the original game.

Optimal Hider strategy

Definition: We recursively define a Hider strategy λ_Q = a probability distribution on the leaf nodes of Q. (If Q is one arc, only one choice is possible.)

Case 1:

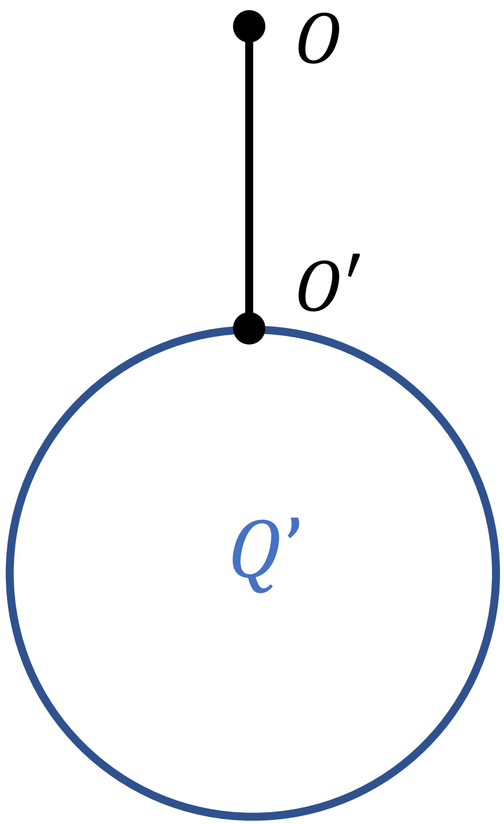

$\boxed{\text{Let } \lambda_Q = \lambda_{Q'}}$

Optimal (worst case) Hider strategy

Case 2: Let $D(Q, O, \lambda_Q)$ be the average distance from O to leaf nodes of Q, weighted according to λ_Q.

WLOG, $D(Q_1, O_1, \lambda_{Q_1}) \geq D(Q_2, O_2, \lambda_{Q_2})$.

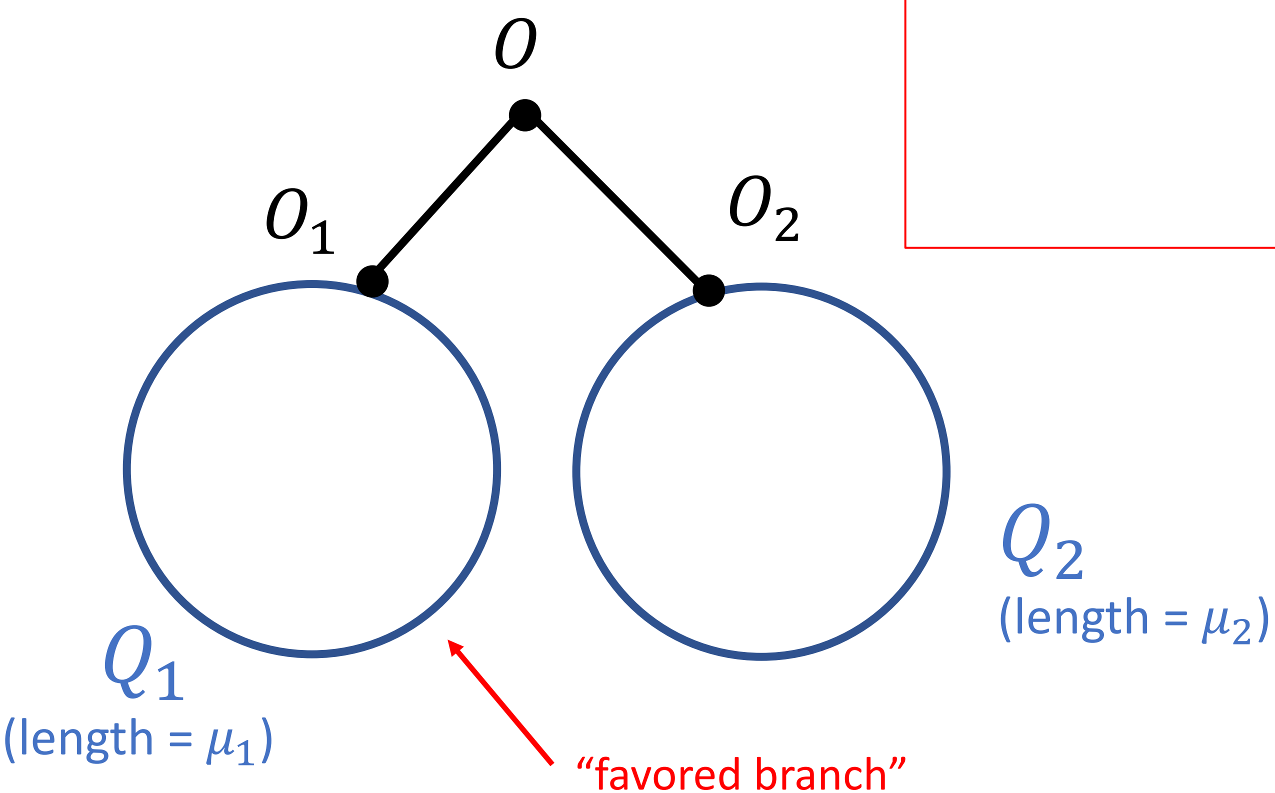

Solution of game

Theorem: For a tree Q with root O and length μ,

(i) The hiding strategy λ_Q is optimal.
(ii) When at a branch node with branches Q_1 and Q_2 of lengths μ_1 and μ_2, it is optimal for the Searcher to search the favored branch with probability

$$\beta = \frac{(2p-1)\big(D(Q_1)-D(Q_2)\big)}{2(p\mu_1+(1-p)\mu_2)},$$

where Q_1 is the favored branch. With probability $1-\beta$ the Searcher follows the signal.
(iii) The value V of the game is

$$V = 2(1-p)\mu + (2p-1)D(Q).$$

Comments

- If all leaves are at the same distance from O (eg. perfect binary trees) then the Searcher should always follow the signal.
- The value is non-increasing in p.
- What about asymmetric speed networks?
- What about variable p?
- What about the depth-first assumption?