Theoretical Computer Science ••• (••••) •••–•••

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

The worst parallel Hanoi graphs ☆

Ernst L. Leiss [a,*], Isaac Mackey [b]

[a] University of Houston, United States
[b] University of Virginia, United States

A R T I C L E   I N F O

Article history:
Received 26 October 2016
Received in revised form 14 September 2017
Accepted 17 September 2017
Available online xxxx
Communicated by A. Fink

Keywords:
Towers of Hanoi
Parallel moves
Hanoi graphs
Maximal number of moves

A B S T R A C T

The "Towers of Hanoi" is a problem that has been extensively studied and frequently generalized. In particular, it has been generalized to be played on arbitrary directed graphs and using parallel moves of two types. We ask what is the largest number of parallel moves, in either of the two models, that is required to move n disks from the starting node to the destination node. Not all directed graphs allow solving this problem; we will call those graphs that do Hanoi graphs. In previous work, we settled the question of what are the worst sequential Hanoi graphs, that is, those graphs that require the largest number of sequential moves. We also demonstrated that the characterization of sequential Hanoi graphs carries over the parallel Hanoi graphs. Here, we determine the worst Hanoi graphs provided parallel moves are allowed. It turns out that for one of the two models of parallel moves, the worst graphs are quite different from the worst sequential graphs, while in the other model of parallelism, there is little difference with the sequential situation.

1. Background

The Towers of Hanoi game is a classical problem that is frequently (ab)used in data structures and algorithms classes to illustrate the power of recursion. For a comprehensive review of the literature on this problem, see [12] and [5]. In [6], the game was generalized to be played on graphs; specifically, there is a finite directed graph $G = (V, E)$ with two distinguished nodes S and D, there are n disks of n different sizes on node S such that no larger disk may lie on top of a smaller disk, and the objective is to move the n disks from S to D subject to the following rules:

1. Only one disk may be moved at a time and only along an edge in G.
2. A disk is always placed on top of all the disks on the node where it is moved and no larger disk may ever be placed on top of a smaller disk.

If the problem can be solved for a given graph G for all $n \geq 1$, the Hanoi problem is called solvable. If for a given graph the associated Hanoi problem is not solvable, the Hanoi problem is called finite. For finite Hanoi problems see [7] and [2].

There is a rather elegant characterization of all those graphs with solvable Hanoi problems.

Theorem 0 *([6]). A Hanoi problem with graph $G = (V, E)$ is solvable if and only if there exist three different nodes v_1, v_2, and v_3 in V such that:*

☆ Support under NSF grants IIS-1359199 (REU Site Program) and DGE-1433817 is acknowledged.
* Corresponding author.
E-mail addresses: coscel@cs.uh.edu (E.L. Leiss), iam6ep@virginia.edu (I. Mackey).

https://doi.org/10.1016/j.tcs.2017.09.012
0304-3975/© 2017 Elsevier B.V. All rights reserved.

1. *There exists a path in G from S to one of the three nodes* v_1, v_2, v_3.
2. *There exists a path in G from one of the three nodes* v_1, v_2, v_3 *to D*.
3. *There exist paths from* v_1 *to* v_2, *from* v_2 *to* v_3, *and from* v_3 *to* v_1 *in G*.

Thus, whether a sequential Hanoi problem is solvable depends exclusively on the graph on which it is played and on the choice of nodes S and D. We will call a directed graph $G = (V, E)$ with distinguished nodes S and D in V a sequential Hanoi graph (G, S, D) if it satisfies the conditions of Theorem 0.

The transition from the original problem (on three pegs) to one involving arbitrary directed graphs is a major generalization. Here, we examine another generalization, namely one that involves parallelism. More specifically, we want to retain the rules outlined above but permit moving more than one disk at a time. Below, we describe in detail the two models that arise naturally in this context. Informally, we may think of agents that move disks from one node (peg) to another along a directed edge. In the sequential problem, there is exactly one agent. In the parallel models, we assume there are several agents which can move disks. These movements for each agent must obey the standard rules for Hanoi games but two or more agents might move disks at the same time. This would correspond to processors carrying out computations in parallel (see [8]). The two models we assume here differ in the way agents may interact; both models arise naturally when considering parallelism, but in one we assume that the two (or more) agents must operate entirely independently of each other (the one-step model) while in the other model an agent may move a disk to a node from which another agent removed a disk (provided all the other requirements are satisfied).

Specifically, in [10], we introduced the two models of parallel moves. In the one-step model, a move consists of a single step; therefore two moves may be done in parallel under this model only if both destination nodes of the disks are not occupied by smaller disks. If there are more than two agents, any two of them must satisfy these requirements. Thus, given two sequential moves

$$v_1{-}d_1 \to v_2 \quad \text{and} \quad v_3{-}d_2 \to v_4,$$

they can be done in parallel under the one-step model iff

(a) $d_1 \neq d_2$,
(b) d_1 is the top disk on v_1,
(c) d_2 is the top disk on v_3,
(d) v_2 does not contain a disk smaller than d_1,
(e) v_4 does not contain a disk smaller than d_2,
(f) $v_2 \neq v_3$,
(g) $v_2 \neq v_4$, and
(h) $v_4 \neq v_1$.

Note that the first three requirements imply the necessary condition that $v_1 \neq v_3$. Therefore, the nodes v_1, v_2, v_3, v_4 must be four distinct nodes.

In the two-step model, a move consists of the lifting of the disk from the node on which it resides and then the placing of that disk on the destination node. It follows that we may have the following situation: disk d_1 resides on node v_1 and disk d_2 resides on node v_2; d_1 is to move from v_1 to v_2 and d_2 is to move from v_2 to node v_3. Assuming all other requirements are satisfied, in the two-step model these two moves may be carried out in parallel, since first both disks are lifted from their respective nodes thereby removing these disks, and in the second step the two disks are moved to their respective destination nodes. It should be clear that in the one-step model, this would not be possible. Therefore, not all of the eight conditions for the one-step mode apply to the two-step model. Thus, two sequential moves

$$v_1{-}d_1 \to v_2 \quad \text{and} \quad v_3{-}d_2 \to v_4$$

can be done in parallel under the two-step model iff

(a) $d_1 \neq d_2$,
(b) d_1 is the top disk on v_1,
(c) d_2 is the top disk on v_3,
(d) if $v_2 \neq v_3$, then v_2 does not contain a disk smaller than d_1, if $v_2 = v_3$, then v_2 does not contain a disk smaller than d_1 once d_2 is removed,
(e) if $v_4 \neq v_1$, then v_4 does not contain a disk smaller than d_2, if $v_4 = v_1$, then v_4 does not contain a disk smaller than d_2 once d_1 is removed, and
(f) $v_2 \neq v_4$.

An important move that can be carried out in the two-step model but not in the one-step model is the swap move. Briefly, if two nodes i and j contain p and q as their top disks respectively and there are the edges (i, j) and (j, i) in the graph, then the disks p and q can be interchanged in one move in the two-step model (provided this does not place a larger