On The Average-Case Complexity of the Bottleneck Tower of Hanoi Problem

Noam Solomon[1] and Shay Solomon[2⋆]

[1] Department of Mathematics
The Weizmann Institute of Science, Rehovot 76100, Israel
Noam.Solomon@weizmann.ac.il
[2] Department of Computer Science
Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel
shayso@cs.bgu.ac.il

Abstract. The Bottleneck Tower of Hanoi (BTH) problem, posed in 1981 by Wood [29], is a natural generalization of the classic Tower of Hanoi (TH) problem. There, a generalized placement rule allows a larger disk to be placed higher than a smaller one if their size difference is less than a given parameter $k \geq 1$. The objective is to compute a shortest move-sequence transferring a legal (under the above rule) configuration of n disks on three pegs to another legal configuration.
In SOFSEM'07, Dinitz and Solomon [7] established tight asymptotic bounds for the worst-case complexity of the BTH problem, for all values of n and k. Moreover, they proved that the average-case complexity is asymptotically the same as the worst-case complexity, for all values of $n > 3k$ and $n \leq k$, and conjectured that the same phenomenon also occurs in the complementary range $k < n \leq 3k$.
In this paper we settle the conjecture of Dinitz and Solomon [7] in the affirmative, and show that the average-case complexity of the BTH problem is asymptotically the same as the worst-case complexity, for all values of n and k. We also discuss some connections between the BTH problem, the problem of sorting with complete networks of stacks using a forklift [1, 19], and the pancake problem [11].

1 Introduction

1.1 The (Classic) TH Problem and Configuration Graph. It is fascinating that the Tower of Hanoi (TH) problem still attracts the interest of mathematicians almost 130 years after its invention by the French number theorist Edouard Lucas (1842-1891). This stems from the rich inherent mathematical structure of the problem, which can be described in the following way. We are given n disks of sizes $1, 2, \ldots, n$ that are stacked on three vertical pegs, subject to the "divine rule": never to have a larger disk above a smaller one (on the same peg). A (legal) move is to pop the top-most disk from one of the pegs and to push it to the top of one of the other two pegs, subject to the divine rule.

⋆ This research has been supported by the Clore Fellowship grant No. 81265410, by the BSF grant No. 2008430, and by the Lynn and William Frankel Center for CS.

Consider the (unweighted and undirected) graph in which the nodes are all the (legal) configurations of the problem, and there is an edge between a pair of configurations if they are reachable via a single move; we refer to this graph as the *TH configuration graph* (or shortly, *TH graph*). The problem of interest is to find *shortest paths* in this graph, i.e., shortest move-sequences transferring a given (initial) configuration to a given (final) configuration. In the most well-known case, both the initial and final configurations are perfect; in a *perfect* configuration all the disks are placed on a single peg in decreasing order of size, from disk n at the bottom to disk 1 at the top. (The name of a disk is identified with its size.) It is easy to show that there is a unique shortest path in the graph between any pair of (distinct) perfect configurations, and the length of this path is $2^n - 1$. It is also long known that the length of the shortest path between any pair of (arbitrary) configurations is at most $2^n - 1$ [30, 12]; thus, the diameter of the TH graph is realized by a pair of (distinct) perfect configurations. Hinz [13] devised an algorithm for computing a shortest path between any pair of configurations; a more efficient algorithm was given later by Romik [25].

The connection between the TH problem and the Sierpiński gasket was first observed by Stewart [27]; in particular, the TH graph is isomorphic to the discrete Sierpiński gasket (see also [17, 25]). This connection was employed by Hinz and Schief [14] to conclude that the average distance on the Sierpiński gasket is 466/885. Thus, the average distance between nodes in the TH graph is asymptotically $(1 + o(1))466/885 \cdot 2^n$; this result was also proved by Chan [4].

The TH graph was showed to be planar and (2-)connected in [28]; it was proved to be Hamiltonian in [21], but it cannot be Eulerian, as there are always nodes with odd degree in the TH graph. For more detailed discussions on properties of the TH graph, see [24, 22], and the references therein.

1.2 The Bottleneck TH Problem and Configuration Graph. In 1981, D. Wood [29] suggested a natural generalization of the TH problem, characterized by the *k-relaxed placement rule*, $k \geq 1$: *If disk j is placed higher than disk i on the same peg (not necessarily neighboring it), then their size difference $j - i$ is less than k.* A move need no longer be subject to the (strict) divine rule, but rather to the k-relaxed placement rule. The objective remains unchanged, i.e., to find shortest paths in the induced configuration graph; when $k = 1$ we arrive at the classic problem. We refer to this problem as the *Bottleneck Tower of Hanoi* (BTH) problem (following Poole [23]), and denote it by $BTH_{n,k}$; also, we refer to the induced configuration graph as the *BTH configuration graph* (or shortly, *BTH graph*), and denote it by $G^{BTH}_{n,k}$. (See Figure 1 for an illustration.)

We remark that the number of all legal (under the k-relaxed placement rule) configurations increases with k. For example, for $k = 1$, the perfect configuration of the n disks is the only legal configuration where all n disks lie on a specific peg, whereas for $k \geq n$, all $n!$ permutations of the n disks on that peg are legal.

Poole [23] suggested a natural algorithm for computing a shortest path in the BTH graph between any pair of perfect configurations, for all values of n and k, but the question whether this algorithm is optimal was left open. Beneditkis, Berend, and Safro [2] proved Poole's algorithm to be optimal for the first non-trivial case $k = 2$ only. Optimality of Poole's algorithm in the general case was proved independently by Dinitz and Solomon [6, 8] and by Chen et al. [5]. It was proved in [7] that there is more than one shortest path in the BTH graph between any pair of perfect configurations, for all $k \geq 2$; also, a complete characterization of the set of all such shortest paths was given therein, complemented with a closed formula, depending on n and k, for the cardinality of this set.

Denote the diameter of the BTH graph $G^{BTH}_{n,k}$ by $Diam(n, k)$. Tight asymptotic bounds for $Diam(n, k)$ were established in [7].

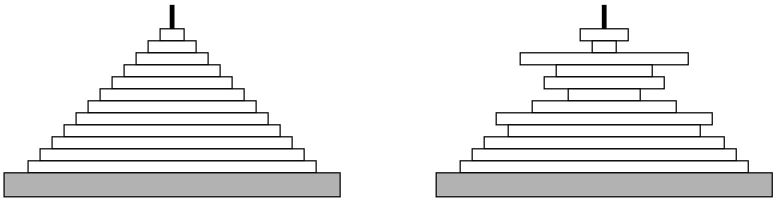

Fig. 1. On the left, the perfect configuration of the disk set $\{1, 2, \ldots, 12\}$ is depicted. Another configuration of this disk set is depicted on the right, where the disks are placed from bottom to top in the following order: 12, 11, 10, 8, 9, 6, 3, 5, 4, 7, 1, 2; it is legal for $BTH_{12,k}$, if $k \geq 5$, but it is illegal if $k \leq 4$, since disk 7 is above disk 3.

Theorem 1 (Theorem 2 in [7]).

$$Diam(n,k) = \begin{cases} \Theta(n \cdot \log n) & \textit{if } n \leq k, \\ \Theta(k \cdot \log k + (n-k)^2) & \textit{if } k < n \leq 2k, \\ \Theta(k^2 \cdot 2^{\frac{n}{k}}) & \textit{if } n > 2k. \end{cases}$$

We remark that the upper bound proof of Theorem 1 is constructive. That is, given an arbitrary pair C, C' of configurations in $G_{n,k}^{BTH}$, a move-sequence of length at most $Diam(n,k)$ transferring C to C' is provided. Refer to [26] for further detail; see Theorem 3.2.1 therein, and the corresponding proof.

We denote by $Avg(n,k)$ the average distance between nodes in $G_{n,k}^{BTH}$.

Notice that $Avg(n,k) \leq Diam(n,k)$, for all values of n and k.

The following theorem from [7] asserts that $Avg(n,k)$ and $Diam(n,k)$ are asymptotically the same, for all values of $n \leq k$ and $n > 3k$.

Theorem 2 (Theorem 3 in [7]). *For all values of $n \leq k$ and $n > 3k$, $Avg(n,k) = \Theta(Diam(n,k))$.*

Finally, it was conjectured in [7] that $Avg(n,k)$ and $Diam(n,k)$ are asymptotically the same also in the complementary range $k < n \leq 3k$. (See Conjecture 1 in [7].) In this paper we settle the conjecture of [7] in the affirmative.

Theorem 3. *For all $k < n \leq 3k$, $Avg(n,k) = \Theta(Diam(n,k))$.*

1.3 Stack Sorting, Fork Stacks, and the Pancake Problem. Consider an arbitrary configuration in the BTH graph $G_{n,k}^{BTH}$, and some path between this configuration and a perfect one on some fixed peg. The move-sequence corresponding to this path can be viewed as a "sorting sequence" that rearranges the n disks in the "correct order" on the fixed peg. If we can find a sorting sequence for any configuration, then we have at hand a "sorting algorithm". Note that by applying the sorting algorithm twice, we can transfer any configuration in the BTH graph to any other (not necessarily perfect) configuration.

In the case $n \leq k$, the pegs behave as stacks. The problem of stack sorting was introduced by Knuth [18], and has received much attention in the literature (see [3, 16, 19, 10], and the references therein). Variants of this problem include imposing restrictions on the legal moves (e.g., a disk that has been popped from a stack may never be pushed back to it again) and considering more than three stacks; in fact, another variant of stack sorting that has been studied in [20, 19] is a natural generalization of the BTH problem, where the placement of disks on a single stack is subject to constraints that are modeled by a conflict graph.

In the case $n > k$ things become more complicated. We need to introduce some notation. Consider a move M of some disk from peg X to peg Y; the pegs

X and Y are the source and destination pegs of M, respectively, whereas the third peg $Z \neq X, Y$ is called the *spare peg* of M. A move M in a move-sequence S is called *switched* (with respect to S) if (i) it is not the first move of S, and (ii) the spare peg of M is different from the spare peg of the preceding move in S. Consider a move-sequence S of the disk set $\{1, 2, \ldots, n\}$, and restrict attention to the k' largest disks $\{n-k'+1, n-k'+2, \ldots, n\}$, for some (small) parameter k'. Notice that any disk in this set cannot be placed higher than any disk in $\{1, 2, \ldots, n-k'+1-k\}$ on the same peg. Thus, it is easy to see that if S contains ℓ switched moves of the k' largest disks, then at least ℓ packet-moves of the disk set $\{1, 2, \ldots, n-k'+1-k\}$ are required; a *packet-move* of a disk set D is a move-sequence transferring the entire disk set D from one peg to another. Finally, the *k'-switched distance* between a pair C, C' of configurations in the BTH graph $G^{BTH}_{n,k}$ is defined as the minimum number of switched moves of the k' largest disks in any move-sequence transferring C to C'. Dinitz and Solomon [7] made a critical use of this notion of switched distance to obtain both upper and lower bounds for $Diam(n, k)$ and $Avg(n, k)$. In particular, they proved (see Lemma 2 in [7]) that the average $2k$-switched distance, taken over all nodes of $G^{BTH}_{n,k}$, is $\Omega(k)$. This lemma of [7] implies that $\Omega(k)$ packet-moves of the disk set $\{1, 2, \ldots, n-3k+1\}$ are required on the average; consequently, a tight asymptotic lower bound on $Avg(n, k)$ was derived in [7], for the range $n > 3k$ only. To prove a tight lower bound on $Avg(n, k)$ in the entire range of $n > k$, we strengthen Lemma 2 of [7] significantly by showing (see Theorem 5 in Section 3) that *for any parameter* $k' = 1, 2, \ldots, O(\log k)$ (rather than just for the specific case $k' = 2k$), the average k'-switched distance, taken over all nodes of $G^{BTH}_{n,k}$, is $\Omega(k')$.

There is a close connection between the notion of switched distance and the notion of *fork stack* [1]. Rather than moving a single element from (the top of) one stack to (the top of) another, a fork stack is equipped with a *forklift* that can be used for moving multiple elements from one peg to another in a single step. It it easy to see that in the case $n \leq k$, the n-switched distance between a pair of configurations in $G^{BTH}_{n,k}$ is (essentially) equal to the minimum number of steps needed to get between these configurations using three fork stacks. In addition, the notions of switched distance and fork stack are closely related to the well-studied pancake problem, where the objective is to sort permutations by prefix reversal (see [11, 15], and the references therein). We remark that an upper bound on the number $f(n)$ of prefix reversals required to sort a permutation of n elements provides the same (up to a constant factor) upper bound on the n-switched distance between a pair of gathered configurations in $G^{BTH}_{n,k}$, $n \leq k$; a configuration is called *gathered* if all n disks lie on the same peg. Also, a lower bound on the n-switched distance between a pair of gathered configurations in $G^{BTH}_{n,k}$ provides the same (up to a constant factor) lower bound on $f(n)$.

1.4 Structure of the Paper. In Section 2 we present the notation that is used throughout the paper. Section 3 is devoted to the proof of Theorem 3. We start (Section 3.1) with proving a statement that is central in the proof of Theorem 3, and then employ this statement (Section 3.2) to complete the proof of the theorem. Finally, in Section 4 we outline some directions for future work.

2 Definitions and Notation

For any disk set D, any configuration C of D, and any $D' \subseteq D$, the *restriction* $C|_{D'}$ is C with all disks not in D' removed. Note that if C is legal (under the k-relaxed placement rule), then its restriction $C|_{D'}$ is legal as well.

For a pair of positive integers w, z, such that $w \leq z$, we denote the sets $\{w, w+1, \ldots, z\}$ and $\{1, 2, \ldots, z\}$ by $[w, z]$ and $[z]$, respectively. The entire disk set $[n]$ is divided into $\lceil \frac{n}{k} \rceil$ *blocks* $B_i = B_i(n)$: $B_1 = [(n-k+1), n]$, $B_2 = [(n-2k+1), (n-k)]$, $\ldots$, $B_{\lceil \frac{n}{k} \rceil} = [1, (n - (\lceil \frac{n}{k} \rceil - 1) \cdot k)]$. Note that the set of disks in any block is allowed to be placed on the same peg in an arbitrary order. For any $n \geq 1$, let $Small(n)$ denote the set $[n] \setminus B_1(n)$.

We will use the following result of [6, 8] in the sequel.

Theorem 4 (Theorem 2 in [6], Theorem 3.2 in [8]). *Let $n = sk + r$, where $0 \leq r < k$, and define $b_{n,k} = (k+r) \cdot 2^s - k$. Under the rules of $BTH_{n,k}$, the length of any packet-move of $[n]$ is at least $b_{n,k}$.*

Remark: It is easy to see that $b_{n,k} = \Theta(k \cdot 2^{\frac{n}{k}})$.

3 Proof of Theorem 3

In this section we prove that $Avg(n, k) = \Omega(Diam(n, k))$, for all $k < n \leq 3k$. Since $Avg(n, k) \leq Diam(n, k)$, for all values of n and k, Theorem 3 would follow.

The following statement is central in our proof.

Theorem 5. *Let $k' \leq k$. For at least half of all pairs of legal configurations of $[n]$, the minimum number of switched moves of disks in $[n - k' + 1, n]$ required to get from one configuration to another is $\frac{k'}{24} - 1$.*

In Section 3.1 we prove Theorem 5. We then employ Theorem 5 in Section 3.2 to complete the proof of Theorem 3.

3.1 Proof of Theorem 5

To prove Theorem 5, we show that for any legal configurations C_{init}, the minimum number of switched moves of disks in $[n - k' + 1, n]$ required to get from C_{init} to at least half the legal configurations of $[n]$ is $\frac{k'}{24} - 1$.

Let $B_1^{k'}(n), B_2^{k'}(n)$ denote the k' largest disks of $B_1(n), B_2(n)$ respectively. In case $k' = k$, these sets coincide with $B_1(n), B_2(n)$ respectively.

For $a, b, c \in B_1^{k'}(n)$, we call the ordered triple (a, b, c) a *triad*.

Consider some configuration C' of $[n]$, and let $C'|_{B_1^{k'}(n)}$ be its restriction to $B_1^{k'}(n)$. For each peg X, let us divide the set of disks in $C'|_{B_1^{k'}(n)}$ on X into triads and a residue of size at most six, according to their placement at $C'|_{B_1^{k'}(n)}$, as follows. Each triad consists of three consecutive disks of $B_1^{k'}(n)$ placed on the same peg, X, whereas there may remain a residue of size at most two, at the bottom of each peg. Let $l_{C'}$ denote the number of such triads; these triads are referred to as the triads of C'. Clearly,

$$l_{C'} \geq \frac{k'}{3} - 2. \tag{1}$$

A disk is called switched w.r.t. a move-sequence S if it participates in at least one switched move. (See Section 1.3 for the definition of switched move.) We say that a triad is *switched* w.r.t. a move-sequence S from C_{init} to C', if at least one disk in that triad is switched w.r.t. S. A triad is called *cheap*, if the disks in it at C_{init} are consecutive on some peg, and either preserve their order at C', or reverse it; otherwise, it is called *expensive*.

The following claim is proved in [7].

Claim. Any expensive triad in C' is switched w.r.t. any move-sequence from C_{init} to C'.

Thus, it suffices to show that for at least half the legal configurations, C', there are at least $\frac{k'}{24} - 1$ expansive triads in C' (w.r.t the initial configuration C_{init}).

Having partitioned the disks of $C'|_{B_1^{k'}(n)}$ into triads, usually there will be other disks in C' placed in the "spaces" between the disks of these triads. For each triad (a, b, c) (not necessarily corresponding to a specific configuration), the collection of three spaces: above a, between a and b, and between b and c is referred to, as the *envelope* of the triad (a, b, c).
We say that a triad is *sparse* (w.r.t C') if the number of disks belonging to $B_2^{k'}(n)$ which are placed in its envelope is at most 3.

Let $l^s_{C'}$ denote the number of sparse triads in C'.

Claim. We have $l^s_{C'} \geq \frac{k'}{12} - 2$.

Proof. For each non-sparse triad, there are at least 4 disks from $B_2^{k'}(n)$ placed in its envelope. As $|B_2^{k'}(n)| = k'$, there can be at most $\frac{k'}{4}$ non-sparse triads. By Equation (1) there are at least $\frac{k'}{3} - 2$ triads, and subtraction from it at most $\frac{k'}{4}$ non-sparse triads completes the proof of the claim. □

For each triad (a, b, c), a triad together with disks from $B_2^{k'}(n)$ placed in its envelope, is called a *completed triad.*

We show next, that given any configuration C', there are less cheap completed triads than expansive ones (w.r.t the initial configuration C_{init}). In fact, we show something slightly stronger, which does not depend on the given configuration C'.

For a collection of p disks $a_1, \dots, a_p \in B_2^{k'}(n)$, and an *ordered* set $Y \subseteq B_1^{k'}(n)$, let $N(a_1, \dots, a_p, Y)$ denote the number of legal configurations of the disks in the union $\{a_1, \dots, a_p\} \cup Y$, whose restrictions to the disks of Y is Y. The following lemma will play a critical role in the proof of Proposition 1.

Lemma 1. *Given $a_1, \dots, a_p, Y$ as above, with $a_1 < a_2 < \ldots < a_p$, let t_i be the number of legal ways to place disk a_i in one of the spaces between the disks of Y. Then $N(a_1, \dots, a_p, Y) = \prod_{i=1}^{p}(t_i + i - 1)$.*

Proof. The proof is via induction on p. The basis of induction $p = 1$ is clear. For the induction step, we assume the claim holds for $p - 1$ and prove it for $p \leq k'$. Assume first that $a_1, \dots, a_{p-1}$ were legally inserted. To count the number of legal ways to insert an additional disk a_p, we first notice that there are t_p legal placements of a_p when ignoring the disks $a_1, \dots, a_{p-1}$. Since a_p is bigger than $a_1, \dots, a_{p-1}$ and smaller than all disks in $B_1^{k'}(n)$, each one of $a_i,\ i = 1, \dots, p-1$ that were inserted gives an additional legal placement for a_p. This gives $p - 1$ additional legal placements for a_p, giving $t_p + p - 1$ legal placements for a_p once the $a_1, \dots, a_{p-1}$ were legally placed. By the induction hypothesis, the number of legal configurations of $\{a_1, \dots, a_{p-1}\} \cup Y$ is $\prod_{i=1}^{p-1}(t_i + i - 1)$. Multiplying it by $t_p + p - 1$ completes the proof of Lemma 1. □

For any a, b, c in $B_1^{k'}(n)$, let $f(a, b, c)$, the *face* of (a, b, c), denote the union of the triads (a, b, c) and (c, b, a). The 6 distinct permutations of $\{a, b, c\}$ are divided into three faces $f(a, b, c), f(a, c, b), f(b, a, c)$.
With the above notation, when $Y = (a, b, c)$ is a triad, and $p = 3$, let

$$N^{a_1,\dots,a_p}_{f(a,b,c)} := N(a_1, \dots, a_p, (a, b, c)) + N(a_1, \dots, a_p, (c, b, a))$$

denote the number of legal ways to complete the triads of the face $f(a,b,c)$ w.r.t the disks $a_1,\dots,a_p$ being placed in its associated envelope.

Proposition 1 ("Triangle inequality"). *With the above notation*

$$N^{a_1,\dots,a_p}_{f(a,b,c)} \leq N^{a_1,\dots,a_p}_{f(a,c,b)} + N^{a_1,\dots,a_p}_{f(c,a,b)}.$$

Proof. Assume without loss of generality that $a_1 < \dots < a_p$. The number of legal ways to complete a triad (a,b,c) w.r.t p disks $a_1,\dots,a_p$ placed in its envelope is determined by an increasing monotone triplet $(s_1,\dots,s_p)$, $0 \leq s_i \leq 3$, where s_i denotes below how many of $\{a,b,c\}$ a_i can be legally placed. Using Lemma 1, we give a formula for this number; let $A < B < C$ be the ordering of (a,b,c).

Lemma 2. *With the above notation, the number of legal completions of the triad (a,b,c) with respect to $a_1,\dots,a_p$ placed in its associated envelope, namely $N(a_1,\dots,a_p,(a,b,c))$, is*

$\prod_{i=1}^{p}(\min(s_i,2)+i)$, *for* $(a,b,c)=(A,B,C)$.
$\prod_{i=1}^{p}(2\chi_{s_i=3}+i)$, *for* $(a,b,c)=(C,B,A)$.
$\prod_{i=1}^{p}(\min(s_i,1)+\chi_{s_i=3}+i)$, *for* $(a,b,c)=(A,C,B)$.
$\prod_{i=1}^{p}(\max(1,s_i)+i-1))$, *for* $(a,b,c)=(B,C,A)$.
$\prod_{i=1}^{p}(2\chi_{s_i=3}+i)$, *for* $(a,b,c)=(C,A,B)$ *(same as* $(a,b,c)=(C,B,A)$*).*
$\prod_{i=1}^{p}(2\chi_{s_i\geq 2}+i)$, *for* $(a,b,c)=(B,A,C)$.

Proof. Assume $a_1 < \dots < a_p$. Let t_i denote the number of legal ways to place disk a_i in the envelope associated to the triad (a,b,c). The space above a is always legal. There are two possible legal ways to insert each a_i, according to the sequence (s_1,s_2,s_3), and the ordering of (a,b,c). For instance, if $(a,b,c) = (A,B,C)$ then $t_i = \min(s_i,2)+1$.
By Lemma 1, letting $Y = \{A,B,C\}$, the number of ways to complete the triad (A,B,C) w.r.t $a_1,\dots,a_p$ placed in its envelope is $N(a_1,\dots,a_p,Y) = \prod_{i=1}^{p}(t_i + i-1)$. Substituting the value of t_i for each configuration of (A,B,C), completes the proof of Lemma 2 in this case. The other orderings of (a,b,c) are treated analogously. This completes the proof of Lemma 2. □

Next, we continue the proof of Proposition 1. We have

$$N^{a_1,\dots,a_p}_{f(C,A,B)} = N(a_1,\dots,a_p,(C,A,B)) + N(a_1,\dots,a_p,(B,A,C)) =$$

$$N(a_1,\dots,a_p,(C,B,A)) + N(a_1,\dots,a_p,(B,A,C)) \leq$$

$$N(a_1,\dots,a_p,(C,B,A)) + N(a_1,\dots,a_p,(A,B,C)) =$$

$$N^{a_1,\dots,a_p}_{f(A,B,C)} < N^{a_1,\dots,a_p}_{f(A,B,C)} + N^{a_1,\dots,a_p}_{f(A,C,B)},$$

where the second equality follows by the equality

$$N(a_1,\dots,a_p,(C,B,A)) = N(a_1,\dots,a_p,(C,A,B)),$$

and the first inequality follows by the fact that the triad (A,B,C) admits maximal $N(a_1,\dots,a_p,(a,b,c))$ among all 6 permutations of a,b,c. This implies that the face $f(C,A,B)$ satisfies the triangle inequality. By similar reasoning, we have

$$N^{a_1,\dots,a_p}_{f(A,C,B)} = N(a_1,\dots,a_p,(A,C,B)) + N(a_1,\dots,a_p,(B,C,A)) \leq$$

$$N(a_1,\dots,a_p,(A,B,C)) + N(a_1,\dots,a_p,(B,A,C)) <$$

$$N(a_1,\dots,a_p,(A,B,C)) + N(a_1,\dots,a_p,(C,B,A)) +$$

$$N(a_1,\dots,a_p,(C,A,B)) + N(a_1,\dots,a_p,(B,A,C)) =$$

$$N^{a_1,\dots,a_p}_{f(A,B,C)} + N^{a_1,\dots,a_p}_{f(C,A,B)},$$

implying the face $f(A,C,B)$ also satisfies the triangle inequality. It remains to show that the face $f(A,B,C)$ also satisfies the triangle inequality, i.e. that

$$N^{a_1,\dots,a_p}_{f(A,B,C)} \leq N^{a_1,\dots,a_p}_{f(A,C,B)} + N^{a_1,\dots,a_p}_{f(C,A,B)}.$$

As

$$N(a_1,\dots,a_p,(C,B,A)) = N(a_1,\dots,a_p,(C,A,B)),$$

it reduces to proving that $N(a_1,\dots,a_p,(A,B,C)) < N(a_1,\dots,a_p,(A,C,B)) + N(a_1,\dots,a_p,(B,A,C)) + N(a_1,\dots,a_p,(B,C,A))$.
We divide into cases according to the $(s_1,\dots,s_p)$-sequence.
If all $s_i \leq 1$ then

$$N(a_1,\dots,a_p,(A,B,C)) = N(a_1,\dots,a_p,(A,C,B)).$$

If $s_p \geq 2$, whereas $s_i \leq 1$ for $i < p$, Lemma 2 implies that

$$N(a_1,\dots,a_p,(A,B,C)) = \left(\prod_{i=1}^{p-1}(s_i+i)\right)(2+p) <$$

$$\left(\prod_{i=1}^{p-1}(s_i+i)\right)(1+p) + \left(\prod_{i=1}^{p-1} i\right)(2+p) + \left(\prod_{i=1}^{p-1} i\right)(1+p) \leq$$

$$N(a_1,\dots,a_p,(A,C,B)) + N(a_1,\dots,a_p,(B,A,C)) + N(a_1,\dots,a_p,(B,C,A)),$$

where the first inequality follows by the fact that for $s_i \leq 1, i \leq p-1, 2 \leq p \leq 3$,

$$\prod_{i=1}^{p-1}(s_i+i) \leq \prod_{i=1}^{p-1} i(3+2p).$$

If $s_p, s_{p-1} \geq 2$, in case $s_1 \geq 2$ or $p = 2$ we have

$$N(a_1,\dots,a_p,(A,B,C)) = N(a_1,\dots,a_p,(B,A,C)).$$

The remaining case is $p = 3, s_1 \leq 1,\ s_2, s_3 \geq 2$, where

$$N(a_1,\dots,a_p,(A,B,C)) = (s_1+1)\cdot 4\cdot 5 < (s_1+1)\cdot 3\cdot 4 + 4\cdot 5 + 3\cdot 4 \leq$$

$$N(a_1,\dots,a_p,(A,C,B)) + N(a_1,\dots,a_p,(B,A,C)) + N(a_1,\dots,a_p,(B,C,A)),$$

completing the proof of Proposition 1. □

Given $a,b,c \in B_1^{k'}(n)$, and $a_1,\dots,a_p \in B_2^{k'}(n),\ p \leq 3$, let $T^{a_1,\dots,a_p}_{a,b,c}$, ($C^{a_1,\dots,a_p}_{a,b,c}$, $E^{a_1,\dots,a_p}_{a,b,c}$) denote the collection of possible, (cheap, expansive, respectively) completions of triads on the disks a,b,c w.r.t the disks $a_1,\dots,a_p$ placed in its envelope. (Notice that all orderings of a,b,c are allowed here) When it causes no confusion, the subscript $a,b,c,a_1,\dots,a_p$ is dropped.

The triangle inequality (Lemma 1) implies that (independently of any specific configuration)

Corollary 1. *With the above notation, there exists a bijective map*

$$f_{a,b,c}^{a_1,\dots,a_p} : T_{a,b,c}^{a_1,\dots,a_p} \to T_{a,b,c}^{a_1,\dots,a_p}$$

such that $f(C) \cap C = \emptyset, f^2 = id, f|_{T \setminus (C \cup f(C))} = id$. *In other words,* f *sends a cheap completed triad* t *to an expansive one, it sends* $f(t)$ *back to* t*, and the other triads of* T *remain untouched by* f.

Next, we observe that for an *illegal* configuration C of $[n]$, there exist two disks a and b on the same peg, such that a is located somewhere above b and $a \geq b+k$. Such an incidence is referred to as a *clash* of a and b, and we say that a *clashes* with b.

Given a legal configuration C' of $[n]$, a completed triad $t \in T_{a,b,c}^{a_1,\dots,a_p}$ belonging to C', and some other completed triad $t \in T_{a,b,c}^{a_1,\dots,a_p}$, let $(C' : t \leftrightarrow t')$ denote the configuration of $[n]$ obtained by changing the order the $3+p$ elements belonging to t according to their order in t', (both t and t' have the same elements), and leaving the other disks of C' untouched.

Lemma 3. *With the above notation,* $(C' : t \leftrightarrow t')$ *is a legal configuration of* $[n]$.

Proof. If $(C' : t \leftrightarrow t')$ is an illegal configuration of $[n]$, then there is a clash involving some disk of t. First, in each peg of C', the disks $[n-k-k']$ are placed above the highest triad on this peg. This implies that these disks cannot clash with disks of the permuted triad. Second, disks belonging to $B_1(n) \setminus B_1^{k'}(n)$ cannot clash with disks from $B_1'(n), B_2'(n)$, as the absolute value of the difference would not exceed k. Finally, there can be no inner clashes between the disks of t, as t' is a legal completed triad. Thus, there can be no clash at all, completing the proof of Lemma 3. □

With the above notation, let

$$\tilde{f}(C', t) := (C' : t \leftrightarrow f(t)).$$

Lemma 3 then implies

Corollary 2. $\tilde{f}(C', t)$ *is a legal configuration of* $[n]$.

Applying the map $\tilde{f}$ iteratively we obtain

Lemma 4. *Given a legal configuration* C_{init}*, there exists a bijection* F_{init} *(depending on* C_{init}*), on the set of legal configuration, such that for a legal configuration* C' *of* $[n]$,

1. *The configuration* $F_{init}(C')$ *is legal.*
2. *The number of expansive triads of* C' *together with the number of expansive triads of* $F_{init}(C')$ *is at least as the number of sparse triads of* C'.

Proof. 1. For the configuration C', let $t_i, i = 1, \dots, l_{C'}^s$ denote the collection of completed sparse triads in it, listed from bottom to top on the three pegs, one after the other. We define F_{init} by iterative applications of $\tilde{f}$ to C' and the completed triads, or more formally,

$$F_{init}^0(C') := C'.$$

$$F_{init}^i(C') := \tilde{f}(F_{init}^{i-1}(C'), t_i), i = 1, \dots, l_{C'}^s.$$

Let

$$F_{init}(C') := F_{init}^{l_{C'},s}(C').$$

By Corollary 2, standard induction implies that each F^i is a legal configuration of $[n]$, and therefore so is $F_{init}(C')$.
Since $f^2 = id$, it is easy to verify that $F_{init}^2 = id$, implying in particular that F_{init} is a bijection.
2. As each cheap (completed) triad is mapped under F_{init} to an expansive one, this part of the lemma is clear, and the proof of Lemma 4 is complete. □

Since for each legal configuration C' of $[n]$ there are at least $\frac{k'}{12} - 2$ sparse triads, Lemma 4 implies that either C' has at least $\frac{k'}{24} - 1$ expansive triads, or else $F_{init}(C')$ has at least $\frac{k'}{24} - 1$ expansive triads. Theorem 5 follows.

3.2 Completing the Proof of Theorem 3

Having proved Theorem 5, we now turn to prove the desired lower bounds on $Avg(n, k)$ in the range $k < n \leq 3k$, thus completing the proof of Theorem 3.

Lemma 5. *For all $n > k$, $Avg(n, k) = \Omega(k \cdot \log k)$.*

Proof. The statement is trivial if $k = O(1)$. We henceforth assume that k is super-constant.

Let C be an arbitrary legal configuration of $[n]$, i.e., an arbitrary node in the BTH graph $G^{BTH}_{n,k}$. Consider a breadth-first search (BFS) tree T_C of $G^{BTH}_{n,k} = (V, E)$ rooted at C. To prove the lemma, it suffices to show that the average distance between C and all other nodes in T_C is $\Omega(k \cdot \log k)$.

It is easy to see that the number of all legal (under the k-relaxed placement rule) configurations of $[n]$ is at least $\Omega(k!)$, yielding $|V| = \Omega(k!)$. Observe that the maximum degree of G^{BTH} is at most 6, hence the maximum degree of T_C is at most 6 as well. It follows that at most $\frac{1}{5}(6^{i+1}-1)$ nodes are at distance at most i in T_C from the root C, for each index $i \geq 0$. Substituting $i = \lfloor \log_6 |V| \rfloor - 1$, we conclude that at least $\frac{4}{5} \cdot |V|$ nodes are at distance at least $\lfloor \log_6 |V| \rfloor = \Omega(\log k!) = \Omega(k \cdot \log k)$ in T_C from C. The lemma follows. □

Lemma 6. *For all $k < n \leq 3k$, $Avg(n, k) = \Omega((n - k)^2)$.*

Proof. Define $k' = \frac{n-k}{2}$, and note that $k' \leq k$. The statement is trivial if $k' = O(1)$. We henceforth assume that k' is super-constant.

By Theorem 5, for at least half of all pairs of legal configurations of $[n]$, at least $\frac{k'}{24} - 1$ switched moves of disks in $[n - k' + 1, n]$ are required to get from one configuration to another. Observe that for any integer $\ell \geq 1$, any move-sequence of $[n]$ that contains ℓ switched moves of disks in $[n - k' + 1, n]$ requires at least ℓ packet-moves of the disk set $Small(n - k' + 1) = [n - k' + 1 - k] = [\frac{n-k}{2} + 1]$. By Theorem 4, each packet-move of $[\frac{n-k}{2} + 1]$ requires at least $b_{\frac{n-k}{2}+1,k} \geq \frac{n-k}{2} + 1$ moves. It follows that at least $(\frac{k'}{24} - 1) \cdot b_{\frac{n-k}{2}+1,k} \geq (\frac{k'}{24} - 1) \cdot (\frac{n-k}{2} + 1) = \Omega((n-k)^2)$ moves are required to get between at least half of all pairs of legal configurations of $[n]$. Hence $Avg(n, k) = \Omega((n - k)^2)$, and we are done. □

Lemmas 5 and 6 imply that

$$Avg(n, k) = \begin{cases} \Omega(k \cdot \log k + (n - k)^2) & \text{if } k < n \leq 2k, \\ \Omega(k^2) & \text{if } 2k < n \leq 3k. \end{cases}$$

By Theorem 1, we conclude that $Avg(n, k) = \Omega(Diam(n, k))$, for all $k < n \leq 3k$. Theorem 3 follows.

4 Future Work

We conclude the paper by outlining some directions for future work.

1. Disregarding constant factors, Theorems 1, 2 and 3 provide tight bounds for $Diam(n,k)$ and $Avg(n,k)$, for all values of n and k. A challenging open problem is to determine the precise constant factors that are hidden within the Θ-notation of these bounds. As mentioned in Section 1, for the special case $k = 1$, the precise constant factors are known.
2. Another intriguing question is to find a pair of configurations that realize the diameter $Diam(n,k)$, for general values of n and k. For the special case $k = 1$, it is known that the diameter is realized by a pair of perfect configurations (on different pegs). However, it is not difficult to show that this phenomenon does not generalize for larger values of k. We believe that this question is of particular interest in the case $n \leq k$, where the pegs behave as stacks.
3. Another natural question in this context is to obtain closed formulae, depending on n and k, for the following two quantities: (i) The number of all legal configurations of $[n]$ on the three pegs, i.e., the number of nodes in the BTH graph $G^{BTH}_{n,k}$. (ii) The number of all (legal) gathered configurations of $[n]$. For very small values of k (say, $k \leq 4$), we can obtain closed formulae for these two quantities by solving recurrence formulae; for example, for $k = 2$, the number of all gathered configurations of $[n]$ is given by the $(n+1)$th Fibnoacci number F_{n+1}. However, obtaining closed formulae for general values of n and k should be significantly more difficult.
4. Theorem 1 provides tight asymptotic bounds for $Diam(n,k)$. As was mentioned in Section 1.2, the upper bound proof of this theorem is constructive. That is, given an arbitrary pair C, C' of configurations in $G^{BTH}_{n,k}$, a move-sequence of length at most $Diam(n,k)$ transferring C to C' is provided. It would be interesting to devise an algorithm for computing a shortest path between any pair of configurations in $G^{BTH}_{n,k}$. In Section 1.1 we mentioned that for the special case $k = 1$, such an algorithm was given in [13, 25]. If computing a shortest path in the general case of $k \geq 2$ is too difficult, one can settle for a path whose length is greater than the shortest one by a "sufficiently small" factor; this question in the particular case of $n \leq k$ coincides with an open question on stacks that was raised in [19].
5. Finally, we believe that investigating additional properties of the BTH graph $G^{BTH}_{n,k}$ is a promising direction for future work.

References

1. Albert, M. H., Atkinson, M. D.: Sorting with a Forklift. Electr. J. Comb. on(2) (2002)
2. Beneditkis, S., Safro, I.: Generalizations of the Tower of Hanoi Problem. Final Project Report, supervised by D. Berend, Dept. of Mathematics and Computer Science, Ben-Gurion University (1998)
3. Bóna, M., A Survey of Stack-Sorting Disciplines. Electron. J. Combin 9(2) (2002-2003)
4. Chan, T.-H.: A statistical analysis of the Towers of Hanoi problem. Int. J. Comput. Math. 28, 5765 (1989)
5. Chen, X., Tian, B., Wang, L.: Santa Claus' Towers of Hanoi. Graphs and Combinatorics 23(1), 153167 (2007)
6. Dinitz, Y., Solomon, S.: Optimal Algorithms for Tower of Hanoi Problems with Relaxed Placement Rules. In: Proc. of 17th ISAAC, pp. 36-47 (2006)
7. Dinitz, Y., Solomon, S.: On Optimal solutions for the Bottleneck Tower of Hanoi problem. In: Proc. of 33rd SOFSEM, pp. 248-259 (2007)

8. Dinitz, Y., Solomon, S.: Optimality of an algorithm solving the Bottleneck Tower of Hanoi problem. ACM Transactions on Algorithms 4(3), Article 25 (2008)
9. Er, M. C.: An analysis of the generalized Towers of Hanoi problem. BIT 23, 429435 (1983)
10. Felsner, S., Pergel, M.: The Complexity of Sorting with Networks of Stacks and Queues. In: Proc. of 16th ESA, pp. 417-429 (2008)
11. Gates, W. H., Papadimitriou, C. H.: Bounds for sorting by prefix reversal. Discrete Mathematics 27(1), 47-57 (1979)
12. Hinz, A. M.: The Tower of Hanoi. Enseign. Math. 35, 289321 (1989)
13. Hinz, A. M: Shortest paths between regular states of the Tower of Hanoi. Inform. Sci. 63, 173181 (1992)
14. Hinz, A. M., Schief, A.: The average distance on the Sierpiński gasket. Probab. Th. and Rel. Fields 87, 129138 (1990)
15. Heydaria, M. H., Sudborough, I. H.: On the Diameter of the Pancake Network. Journal of Algorithms 25(1), 67-94 (1997)
16. Jiang, T., Li, M., Vitányi, P M. B.: Average-Case Analysis of Algorithms Using Kolmogorov Complexity. J. Comput. Sci. Technol. 15(5): 402-408 (2000)
17. Klavzar, S., Milutinovic, U.: Graphs $S(n,k)$ and a variant of the Tower of Hanoi problem. Czechoslovak Math. J. 47 (122), 95104 (1997)
18. Knuth, D. E.: The Art of Computer Programming, Volume 2 (Fundamental Algorithms), Third Edition. Addition-Wesley (1997)
19. König, F G., Lübbecke, M. E.: Sorting with Complete Networks of Stacks. In: Proc. of 19th ISAAC, pp. 895-906 (2008)
20. König, F G., Lübbecke, M. E., Möhring, R. H., Schäfer, G., Spenke, I.: Solutions to Real-World Instances of PSPACE-Complete Stacking. In: Proc. of 15th ESA, pp. 729-740 (2007)
21. Lu, X.: Towers of Hanoi Graphs. Int. J. Comput. Math. 19, 23-38 (1986)
22. Nurlybaev, A. N.: K zadache Lyuka o bashnyakh iee obobshchenii, Izv. Nats. Akad. Nauk Respub. Kazakhstan Ser. Fiz.-Mat. 1(182), 24-30 (1995)
23. Poole, D.: The Bottleneck Towers of Hanoi Problem. J. of Recreational Math. 24(3), 203-207 (1992)
24. Poole, D. G.: The Towers and Triangles of Professor Claus (or, Pascal Knows Hanoi). Math. Mag. 67, 323-344 (1994)
25. Romik, D.: Shortest Paths in the Tower of Hanoi Graph and Finite Automata. SIAM J. Discrete Math. 20(3): 610-622 (2006)
26. Solomon, S.: Algorithms and Lower Bounds on Hanoi Tower Problems. M.Sc. thesis, Deptartment of Computer-Science, Ben-Gurion University, Israel (Sept. 2006). Available online via http://www.cs.bgu.ac.il/∼shayso/Thesis.pdf
27. Stewart, I.: Le lion, le lama et la laitue, Pour la Science 142, 102107 (1989)
28. Scorer, R. S., Grundy, P. M., Smith, C. A. B.: Some binary games, Math. Gaz. 28, 96-103 (1944)
29. Wood, D.: The Towers of Brahma and Hanoi revisited. J. of Recreational Math. 14(1), 17-24 (1981-1982)
30. Wood, D.: Adjudicating a Towers of Hanoi contest. Internat. J. Comput. Math. 14, 199207 (1983).