

Functional Pearl: La Tour D’Hanoi

Ralf Hinze

Computing Laboratory, University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, England

ralf.hinze@comlab.ox.ac.uk

Abstract

This pearl aims to demonstrate the ideas of wholemeal and projective programming using the Towers of Hanoi puzzle as a running example. The puzzle has its own beauty, which we hope to expose along the way.

Categories and Subject Descriptors D.1.1 [*Programming Techniques*]: Applicative (Functional) Programming; D.3.2 [*Programming Languages*]: Language Classifications—Applicative (functional) languages, Haskell; D.3.3 [*Programming Languages*]: Language Constructs and Features—Frameworks, patterns, recursion

General Terms Algorithms, Design, Languages

Keywords Towers of Hanoi, wholemeal programming, projective programming, Hanoi graph, Sierpiński graph, Sierpiński gasket graph, Gray code

1. Introduction

Functional languages excel at *wholemeal programming*, a term coined by Geraint Jones. Wholemeal programming means to think big: work with an entire list, rather than a sequence of elements; develop a solution space, rather than an individual solution; imagine a graph, rather than a single path. The wholemeal approach often offers new insights or provides new perspectives on a given problem. It is nicely complemented by the idea of *projective programming*: first solve a more general problem, then extract the interesting bits and pieces by transforming the general program into more specialised ones. This pearl aims to demonstrate the techniques using the popular Towers of Hanoi puzzle as a running example. This puzzle has its own beauty, which we hope to expose along the way.

2. The Hanoi graph

The Towers of Hanoi puzzle was invented by the French number theorist Édouard Lucas more than a hundred years ago. It consists of three vertical pegs, on which discs of mutually different diameters can be stacked. For reference, we call the pegs A, B and C and let a , b and c range over pegs.

$\text{data Peg} = \text{A} \mid \text{B} \mid \text{C}$

I own a version of the puzzle where the pegs are arranged in a row. However, the mathematical structure of the puzzle becomes clearer,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP’09, August 31–September 2, 2009, Edinburgh, Scotland, UK.
Copyright © 2009 ACM 978-1-60558-332-7/09/08...\$10.00

if we arrange them in a circle. Initially, the discs are placed on one peg in decreasing order of diameter. The task is then to move the disks, one at a time, to another peg subject to the rule that a larger disk must not be placed on a smaller one.

This restriction implies that a configuration can be represented by a list of pegs: the first element determines the position of the largest disc, the second element the position of the second largest disc, and so forth. Consequently, there are 3^n possible configurations where n is the total number of discs. Lucas’ original puzzle contained 8 discs. The instructions of the puzzle refer to an old Indian legend, attributed to the French mathematician De Parville, according to which monks were given the task of moving a total of 64 discs. Since the day of the world’s creation, they transfer the discs, one per day. According to the legend, once they complete their sacred task, the world will come to an end.

Now, taking a wholemeal approach, let us first develop the big picture. The set of all configurations together with the set of legal moves defines a graph, which turns out to enjoy a nice inductive definition. If there are no discs to move around, the graph consists of a singleton node: \circ . For the inductive step we reason as follows: the largest disc can only be moved, if all the smaller discs reside on *one* other peg. The smaller discs, however, can be moved independent of the largest one. As the largest disc may rest on one of three pegs, the graph for $n + 1$ discs consequently incorporates three copies of the graph for n discs linked by three edges. The diagram in Fig. 1 illustrates the construction. The graph has the shape of a triangle; the dashed lines indicate the sub-graphs (for $n = 0$ the sub-graphs collapse to singleton nodes); the three solid lines connect the sub-graphs. The inductive construction shows that




Figure 1. Inductive construction of the Hanoi graph.

the graph is planar: it can be drawn so that no edges intersect. Fig. 2 displays the graph for 4 discs. To reduce clutter, the peg of the largest disc is always written in the centre of the respective sub-graph, with the size of the font indicating the size of the disc. Can you find the configuration $[\mathbf{B}, \text{A}, \text{B}, \text{c}]$? The corners of the triangle correspond to *perfect configurations*: all the discs reside on one peg. The example graph shows that every configuration permits three different moves, except for the three perfect configurations, where only two moves are possible.




Figure 2. The Hanoi graph for 4 discs.

Let us turn our attention to the layout of the sub-graphs. The following notation proves to be useful: the *arrangement* $x^y z$ denotes a permutation of x , y and z , which by assumption are pairwise different. Using $a^b c$ to indicate the position of the largest disc—the largest disc in the left triangle resides on a and so forth—we observe that if the corners of the graph are $a^b c$, then the corners of the sub-graphs are $a^c b$, $c^b a$ and $b^a c$, respectively. Using this observation, we can capture the informal description of the construction as a pseudo-Haskell program.

$$\begin{aligned} \text{graph}_0 \ (a^b c) &= \circ \\ \text{graph}_{n+1} \ (a^b c) &= b \triangleleft \text{graph}_n \ (c^b a) \\ &\quad \swarrow \quad \searrow \\ & a \triangleleft \text{graph}_n \ (a^c b) \quad - \quad c \triangleleft \text{graph}_n \ (b^a c) \end{aligned}$$

The function graph_n maps an arrangement $a^b c$ to an undirected graph, whose vertices are lists of pegs of length n . The notation $a \triangleleft g$ means prepend a to all the vertices of g ; \circ denotes a singleton node labelled with the empty list. We leave the type of graphs unspecified. If the type were a functor, then $a \triangleleft g$ would be $\text{fmap} \ (a :) \ g$. The call $\text{graph}_4 \ (A^B C)$ yields the graph in Fig. 2.

A few observations are in order. The definition of graph_n implies that the graph has 3^n nodes ($a_0 = 1$, $a_{n+1} = 3 \cdot a_n$) and $(3^{n+1} - 3)/2$ edges ($a_0 = 0$, $a_{n+1} = 3 \cdot a_n + 3$). Furthermore, the length of a side of the triangle is $2^n - 1$ ($a_0 = 0$, $a_{n+1} = 2 \cdot a_n + 1$). Since there are only $3! = 6$ permutations of three different items, the graph contains at most six different mini-triangles: $A^B C$, $C^A B$, $B^C A$, $A^C B$, $C^B A$ and $B^A C$. Note that the first three arrange the pegs clockwise and the last three counterclockwise. Inspecting the definition of graph_n , we see that the direction changes with every recursive step. The diagram in Fig. 3 illustrates the change of direction. This observation implies that the recursion pattern is quite regular: at any depth there are only three different recursive calls. Fig. 4 visualises the call structure using three different colours.

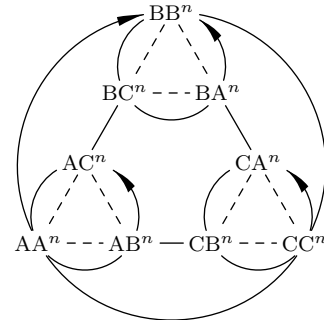


Figure 3. Change of direction.

3. Towers of Hanoi, recursively

Solving the Towers of Hanoi puzzle amounts to finding a shortest path in the corresponding graph. Clearly, the shortest path between two corners is along the side of the triangle. Projecting onto the lower side, we transform graph_n to

$$\begin{aligned} \text{hanoi}'_0 \ (a^b c) &= [[]] \\ \text{hanoi}'_{n+1} \ (a^b c) &= a \triangleleft \text{hanoi}'_n \ (a^c b) \# c \triangleleft \text{hanoi}'_n \ (b^a c) . \end{aligned}$$

Now, $a \triangleleft x$ is shorthand for $\text{fmap} \ (a :) \ x$. The call $\text{hanoi}'_n \ (a^b c)$ returns a list of configurations solving the problem of ‘moving n discs from a to c using b ’.

Instead of configurations (vertices of the graph) we can alternatively return a list of moves (corresponding to edges).

$$\begin{aligned} \text{hanoi}_0 \ (a^b c) &= [] \\ \text{hanoi}_{n+1} \ (a^b c) &= \text{hanoi}_n \ (a^c b) \# [(a, c)] \# \text{hanoi}_n \ (b^a c) \end{aligned}$$




Figure 4. Call structure of *hanoi* at recursion depth 4.

The pair (a, c) represents the instruction ‘move the topmost disc from a to c ’. Here is the sequence of moves for $n = 4$:

$\gg hanoi_4 (A^B C)$
 $[(A, B), (A, C), (B, C), (A, B), (C, A), (C, B), (A, B), (A, C),$
 $(B, C), (B, A), (C, A), (B, C), (A, B), (A, C), (B, C)]$.

Note that the lower part of the list can be obtained from the upper part via a clockwise rotation of the pegs: $a^c b$ becomes $b^a c$.

There is at least one further variation: instead of an arrangement one can pass a source and a target peg.

$hanoi_0 \quad a \perp c = []$
 $hanoi_{n+1} \ a \perp c$
 $= hanoi_n \ a \ (a \perp c) \uplus [(a, c)] \uplus hanoi_n \ (a \perp c) \ c$

The function \perp , which determines ‘the other peg’, is given by

$A \perp A = A; \quad A \perp B = C; \quad A \perp C = B$
 $B \perp A = C; \quad B \perp B = B; \quad B \perp C = A$
 $C \perp A = B; \quad C \perp B = A; \quad C \perp C = C$.

We will find some use for \perp later on. For the moment, we just note that the operation is commutative and idempotent, but not associative.

4. Towers of Hanoi, parallelly

Imagine that the monastery always accommodates as many monks as there are discs. The tallest monk is responsible for moving the largest disc, the second tallest monk for moving the second largest disc, and so forth. Can we set up a work schedule for the monastery?

Inspecting Fig. 2, we notice that, somewhat unfairly, the smallest monk is the busiest. Since the smallest triangles correspond to moves of the smallest disc, he is active every other day. We can ex-

tract his work plan from $hanoi_n$ by omitting all the moves, except when n equals 1.

$cycle_0 \quad (a^b c) = []$
 $cycle_1 \quad (a^b c) = [(a, c)]$
 $cycle_{n+1} \ (a^b c) = cycle_n \ (a^c b) \uplus cycle_n \ (b^a c)$

The function is called *cycle*, because the smallest disc cycles around the pegs: in the recursive call it is moved from a to b and then from b to c . Whether it cycles clockwise or counterclockwise depends on the parity of n —the direction changes with every recursive call of *cycle*.

Of course, the smallest disc is by no means special: all the discs cycle around the pegs, albeit at a slower pace and in alternating directions. In fact, *hanoi* satisfies the following ‘fractal’ property:

$$hanoi_{n+1} \ (a^b c) = cycle_{n+1} \ (a^b c) \ \Upsilon \ hanoi_n \ (a^b c) \ , \quad (1)$$

where Υ denotes the interleaving of two lists.

$$\begin{aligned}
 [] \quad \Upsilon \ bs &= bs \\
 (a : as) \ \Upsilon \ bs &= a : (bs \ \Upsilon \ as)
 \end{aligned}$$

The fractal property suggests the following alternative definition of *hanoi*, which has a strong parallel flavour.

$phanoi_0 \quad (a^b c) = []$
 $phanoi_{n+1} \ (a^b c) = cycle_{n+1} \ (a^b c) \ \Upsilon \ phanoi_n \ (a^b c)$

In words, the smallest monk starts to work on the first day; he is active every second day and moves his disc, say, clockwise around the pegs. The second smallest monk starts on the second day; he is active every fourth day and moves his disc counterclockwise. And so forth. Actually, only the smallest monk must remember the direction; for the larger discs there is no choice, as one of the other two pegs is blocked by the smallest disc.

There is an intriguing cross-connection to binary numbers: the activity diagram of the monks (Which monk has to work on a given day?)

$$\begin{aligned} discs_0 &= [] \\ discs_{n+1} &= discs_n \# [n] \# discs_n \end{aligned}$$

yields the *binary carry sequence* or *ruler function* as n goes to infinity (Hinze, 2008). This sequence gives the number of trailing zeros in the binary representations of the positive numbers, most significant bit first. Or put differently, it specifies the running time of the binary increment. We will make use of this observation later on.

The fractal property (1) enjoys a simple inductive proof, which makes essential use of the following *abide law*. If x_1 and x_2 are of the same length, then

$$(x_1 \# y_1) \curlywedge (x_2 \# y_2) = (x_1 \curlywedge x_2) \# (y_1 \curlywedge y_2) .$$

The basis of the induction is straightforward. Here is the inductive step.

$$\begin{aligned} & cycle_{n+2} (a^b c) \curlywedge hanoi_{n+1} (a^b c) \\ = & \{ \text{definition of } cycle \text{ and definition of } hanoi \} \\ & (cycle_{n+1} (a^c b) \# cycle_{n+1} (b^a c)) \\ & \curlywedge (hanoi_n (a^c b) \# [(a, c)] \# hanoi_n (b^a c)) \\ = & \{ \text{abide law} \} \\ & (cycle_{n+1} (a^c b) \curlywedge hanoi_n (a^c b)) \# [(a, c)] \\ & \# (cycle_{n+1} (b^a c) \curlywedge hanoi_n (b^a c)) \\ = & \{ \text{ex hypothesis} \} \\ & hanoi_{n+1} (a^c b) \# [(a, c)] \# hanoi_{n+1} (b^a c) \\ = & \{ \text{definition of } hanoi \} \\ & hanoi_{n+2} (a^b c) \end{aligned}$$

The abide law is applicable in the second step, because the two lists, $cycle_{n+1} (a^c b)$ and $hanoi_n (a^c b) \# [(a, c)]$, have the same length, namely, 2^n . Speaking of the length of lists, note that $2^{n+1} - 1 = \sum_{i=0}^n 2^i$ is a simple consequence of the fractal property.

5. When will the world come to an end?

Many visitors come to the monastery. Looking at the configuration of discs, they often wonder how many days have passed since the creation of the world. Or, when will the world come to an end?

We can answer these questions by locating the current configuration in the Hanoi graph. If we use as positions the three digits 0^2_1 —that is, 0 for the left triangle, 2 for the upper triangle and 1 for the right triangle—then we obtain the answer to the first question in binary. Fig. 5 displays the Hanoi graph for 4 discs suitably re-labelled—this graph is also known as the Sierpiński graph. The 2^4 nodes on the lower side of the triangle, and only those, are labelled with binary numbers. Consequently, if the current position contains a 2, we know that the monks have lost track. (To answer the second question, we use as positions 1^2_0 instead of 0^2_1 . If we are only interested in the distance to the final configuration, we simply replace the digit 2 by a 1.)

$$\begin{aligned} pos (a^b c) [] &= [] \\ pos (a^b c) (p : ps) & \\ | p = a &= 0 : pos (a^c b) ps \\ | p = b &= 2 : pos (c^b a) ps \\ | p = c &= 1 : pos (b^a c) ps \end{aligned}$$

For instance, $pos (A^B C) [A, B, C, A]$ yields $[0, 1, 0, 1]$, the binary representation of 5, most significant bit first. The function $pos (a^b c)$ defines a bijection between $\{A, B, C\}^n$ and $\{0, 2, 1\}^n$

for any given initial arrangement $a^b c$. This arrangement can be seen as representing the bijective mapping $\{a \mapsto 0, b \mapsto 2, c \mapsto 1\}$. Alternatively, we can use an arrangement, say, l^t_r as a representation of the ‘inverse’ mapping $\{A \mapsto l, B \mapsto t, C \mapsto r\}$ obtaining the following slightly more succinct variant of pos .

$$\begin{aligned} pos' (l^t_r) [] &= [] \\ pos' (l^t_r) (A : ps) &= l : pos' (l^r_t) ps \\ pos' (l^t_r) (B : ps) &= t : pos' (r^t_l) ps \\ pos' (l^t_r) (C : ps) &= r : pos' (l^l_r) ps \end{aligned}$$

The two variants are related by $pos (A^B C) = pos' (0^2_1)$. The latter definition is particularly easy to invert.

$$\begin{aligned} conf (a^b c) [] &= [] \\ conf (a^b c) (0 : ps) &= a : conf (a^c b) ps \\ conf (a^b c) (2 : ps) &= b : conf (c^b a) ps \\ conf (a^b c) (1 : ps) &= c : conf (b^a c) ps \end{aligned}$$

The call $conf (A^B C) [0, 1, 0, 1]$ yields $[A, B, C, A]$, the configuration we obtain after 5 days. The functions pos' and $conf$ are actually identical, if we identify A with 0, B with 2 and C with 1. Then pos' is an involutive graph isomorphism between Hanoi graphs and Sierpiński graphs of the same order.

If the monks have lost track—a 2 appears in the answer list—then we can use the idea of locating a configuration in the Hanoi graph to determine the shortest path to the final configuration c^n . This leads to the following generalised version of *hanoi*.

$$\begin{aligned} ghanoi_0 [] (a^b c) &= [] \\ ghanoi_{n+1} (p : ps) (a^b c) & \\ | p = a &= ghanoi_n ps (a^c b) \# [(a, c)] \# hanoi_n (b^a c) \\ | p = b &= ghanoi_n ps (b^c a) \# [(b, c)] \# hanoi_n (a^b c) \\ | p = c &= ghanoi_n ps (a^b c) \end{aligned}$$

(The index n is actually redundant: it always equals the length of the peg list.) Depending on the location of p , we either walk within the left triangle and then along the lower side, or within the upper triangle and then along the right side, or within the right triangle. The reader should convince herself that $ghanoi_n ps (a^b c)$ indeed yields the *shortest path* between ps and c^n . Note that the first two equations are perfectly symmetric, in fact, $ghanoi_n ps (a^b c) = ghanoi_n ps (b^a c)$. As an aside, if we fuse *ghanoi* with *length*, then we obtain a function that yields the distance to the final configuration.

6. Towers of Hanoi, iteratively

The generalised version of the puzzle—moving from an arbitrary configuration to a perfect configuration—serves as an excellent starting point for the derivation of an iterative solution, that is, a function that maps a configuration to the next move or to the next configuration.

The next move is easy to determine: we fuse *ghanoi* with the natural transformation

$$\begin{aligned} first &:: [\alpha] \rightarrow \text{Maybe } \alpha \\ first [] &= \text{Nothing} \\ first (a : as) &= \text{Just } a \end{aligned}$$

that maps a list to its first element. We obtain

$$\begin{aligned} move_0 [] (a^b c) &= \text{Nothing} \\ move_{n+1} (p : ps) (a^b c) & \\ | p = a &= move_n ps (a^c b) \# \text{Just}(a, c) \# first(hanoi_n(a^b c)) \\ | p = b &= move_n ps (b^c a) \# \text{Just}(b, c) \# first(hanoi_n(a^b c)) \\ | p = c &= move_n ps (a^b c) . \end{aligned}$$




Figure 5. The Sierpiński graph of order 4.

The operator $\#$ is overloaded to also denote concatenation of *Maybe* values ($\#$ is really the *mplus* method of *MonadPlus*).

```
(#)      :: Maybe α → Maybe α → Maybe α
Nothing # m = m
Just a  # m = Just a
```

We can simplify the definition of *move* drastically: since *Just a # m = Just a*, the calls to *hanoi* can be eliminated; because of that the index *n* is no longer needed. Furthermore, the first two equations can be unified through the use of \perp , the operator that determines ‘the other peg’. Finally, the second argument, $a^b c$, can be simplified to *c*, the target peg.

```
move []      c = Nothing
move (p : ps) c
  | p ≠ c    = move ps (p ⊥ c) # Just (p, c)
  | p == c   = move ps c
```

Here is a short interactive session that illustrates the use of *move*.

```
>>> move [A, B, B, A, C, C, C, B] C
Just (B, C)
>>> move [A, B, B, A, C, C, C, C] C
Just (A, B)
>>> move [C, C, C, C, C, C, C, C] C
Nothing
```

Since the pair (B, C) means ‘move the topmost disc from peg B to peg C, the configuration following [A, B, B, A, C, C, C, B] is [A, B, B, A, C, C, C, C]. Incidentally, if we start with the initial peg list [A, A, A, A, A, A, A, A] and target C, we obtain these configurations after 110 and 111 steps.

The function *move* implements a two-way algorithm: on the way down the recursion it calculates the target peg for each disc (using \perp); on the way up the recursion it determines the smallest disc that is not yet in place (using $\#$). The following table makes

the target pegs explicit. The rows labelled c_i lists the target peg for each disc: the first, user-supplied target is $c_0 = C$, the next target is $c_1 = p_0 \perp c_0$, and so forth.

	<i>i</i>	0	1	2	3	4	5	6	7	
110	p_i	A	B	B	A	C	C	C	B	
	c_i	C	B	B	B	C	C	C	C	A
111	p_i	A	B	B	A	C	C	C	C	
	c_i	C	B	B	B	C	C	C	C	C

The smallest disc is not in place, so it is moved from B to C. In the next round, disc 3 has to be moved from A to B.

The smaller discs are moved more frequently, so it is actually prudent to reverse the list of pegs so that the peg on which the smallest disc is located comes first. The situation is similar to the binary increment: visiting the digits from least to most significant bit is more efficient than the other way round. Let us assume for the moment that we know the ‘last target’, the value of *c* that is discarded in the first equation of *move* (the pegs that stick out in the example above). Since \perp is reversible, $p \perp c = c'$ iff $c = p \perp c'$, we can reconstruct the previous target *c* from the next target c' . The following variant of *move* makes use of this fact—the suffix ‘i’ indicates that the input list is now arranged in increasing order of diameter.

```
movei []      c' = Nothing
movei (p : ps) c'
  | p ≠ (p ⊥ c') = Just (p, p ⊥ c') # movei ps (p ⊥ c')
  | p == (p ⊥ c') = movei ps c'
```

We consistently changed the second argument of *move* to reflect the fact that we compute the previous from the next target and additionally replaced the remaining occurrences of *c* by $p \perp c'$. Again, we can simplify the code: $p \neq (p \perp c')$ is the same as $p \neq c'$. Applying *Just a # m = Just a* once more, we can eliminate the first recursive call to *movei* so that the function

stops as soon as the smallest displaced disc is found—this was the purpose of the whole exercise. We obtain

$$\begin{aligned} \text{movei } [] \quad c' &= \text{Nothing} \\ \text{movei } (p : ps) \ c' & \\ \quad | p \neq c' &= \text{Just } (p, p \perp c') \\ \quad | p == c' &= \text{movei } ps \ c' . \end{aligned}$$

In a nutshell, *movei* determines the smallest disc that does not reside on the last target.

So, for an iterative version of *hanoi* we have to maintain two pieces of information: the current configuration and the current ‘last target’. It remains to determine the initial last target and how the last target changes after each move. If the list of pegs is given in the original decreasing order, then we can transform *move* to

$$\begin{aligned} \text{last } [] \quad c &= c \\ \text{last } (p : ps) \ c &= \text{last } ps \ (p \perp c) , \end{aligned}$$

which yields the last target. It is not hard to see that *last* is an instance of the famous *foldl*: $\text{last } ps \ c = \text{foldl } (\perp) \ c \ ps$. If we reverse the list, we simply have to replace *foldl* by *foldr* additionally using the fact that \perp is commutative: $\text{foldr } (\perp) \ c \ ps = \text{foldl } (\perp) \ c \ (\text{reverse } ps)$.

Next, we augment *movei* so that it returns the next configuration instead of the next move and additionally the next ‘last target’.

$$\begin{aligned} \text{step } ([], \quad c') &= \text{Nothing} \\ \text{step } (p : ps, \ c') & \\ \quad | p \neq c' &= \text{Just } ((p \perp c') : ps, p \perp c') \\ \quad | p == c' &= \mathbf{do} \ \{ (ps', c) \leftarrow \text{step } (ps, c'); \\ &\quad \text{return } (p : ps', p \perp c) \} \end{aligned}$$

Consider the second equation: p is moved to $p \perp c'$. After the move, the disc resides on the target peg. Consequently, the next target is also $p \perp c'$ —recall that \perp is idempotent. This target is then updated in the third equation ‘on the way back’ mimicking *move*’s mode of operation. The function *step* runs in constant amortised time, since it performs the same number of steps as the binary increment—recall that the activity diagram of the monks coincides with the binary carry sequence.

The next last target can, in fact, be easily calculated by hand. Consider the following two subsequent moves (as usual, a , b and c are pairwise different).

$$\begin{array}{cccccc} & n & n-1 & n-2 & & 1 & 0 \\ \cdots & a & c & c & \cdots & c & c \\ \cdots & b & c & c & \cdots & c & c & c \\ \cdots & b & c & c & \cdots & c & c \\ \cdots & b & b & a & \cdots & b & a & b \end{array}$$

Assume that the first configuration ends with an even number of cs . The topmost disc of a is then moved to b . The new succession of target pegs consequently alternates between b and a : since the number of cs is even, the new last target is b ; for an odd number, it is a . So, the monks can be instructed as follows: determine the smallest disc that is not on c . Transfer it from a to b . If the disc’s number is even, the new last target is b ; otherwise, it is a .

If we solve the original puzzle, that is, if the configurations lie on one of the sides of the triangle, then the next last target is even easier to determine: like the discs, it cycles around the pegs. If the pegs are arranged A^B_C and we move the discs from A to C , then the last target moves counterclockwise for an even number of discs and clockwise for an odd number.

Summing up, we obtain the following iterative implementation for solving the *generalised* Hanoi puzzle.

$$\begin{aligned} \text{ihanoi } ps \ c &= \text{map fst } (\text{iterate step } (ps, \text{foldr } (\perp) \ c \ ps)) \\ \text{iterate} &:: (a \rightarrow \text{Maybe } a) \rightarrow (a \rightarrow [a]) \\ \text{iterate } f \ x &= x : \mathbf{case} \ f \ x \ \mathbf{of} \ \{ \text{Nothing} \rightarrow []; \\ &\quad \text{Just } x' \rightarrow \text{iterate } f \ x' \} \end{aligned}$$

7. Longest paths and Sierpiński’s triangle

So far we have considered shortest paths in the Hanoi graph. Since the destruction of the world hangs in the balance, as a gift to future generations, we might want to look for the *longest path*. In the following variant of *hanoi* the largest disc makes a detour. (According to the Indian legend, the temple is actually in Bramah rather than in Hanoi, hence the name of the function.)

$$\begin{aligned} \text{bramah}_0 \quad (a^b_c) &= [] \\ \text{bramah}_{n+1} \quad (a^b_c) &= \text{bramah}_n \ (a^b_c) \ \# \ [(a, b)] \\ &\quad \# \ \text{bramah}_n \ (c^b_a) \ \# \ [(b, c)] \\ &\quad \# \ \text{bramah}_n \ (a^b_c) \end{aligned}$$

The largest disc is first moved from a to b , and then from b to c . Since bramah_n returns $3^n - 1$ moves ($a_0 = 0$, $a_{n+1} = 3 \cdot a_n + 2$), we have actually found a *Hamiltonian path*. The path has another interesting property: if the pegs are arranged in a row, $a \ b \ c$, then discs are only moved between adjacent pegs.

The Hamiltonian path for four discs is displayed in Fig. 6. The picture is quite appealing. Actually, if we move the sub-triangles closer to each other so that the corners touch, we obtain a nice fractal image. Fig. 7 shows the result for 7 discs. The corresponding graph is known as the discrete Sierpiński gasket graph.¹ The picture has been drawn using Functional Metapost’s turtle graphics (Korittky, 1998).

$$\begin{aligned} \text{curve}_0 \quad d &= \text{forward } \& \ \text{turn } (2 * d) \ \& \ \text{forward} \\ \text{curve}_{n+1} \quad d &= \text{turn } d \ \& \ \text{curve}_n \ (-d) \ \& \ \text{turn } d \ \& \ \text{curve}_n \ d \\ &\quad \& \ \text{turn } d \ \& \ \text{curve}_n \ (-d) \ \& \ \text{turn } d \end{aligned}$$

The command *forward* moves the turtle one step forward, *turn* d turns the turtle clockwise by d degrees, and $\&$ sequences two turtle actions. Since turtle graphics is state-based—the turtle has a position and faces a direction—recursive definitions typically maintain an invariant. To draw the ‘triangle’ a^b_c , we start at a looking at b and stop at c looking away from b . The overall change of direction is twice the second argument of *curve*, which for equilateral triangles is either 60° or -60° .

The curve is closely related to Sierpiński’s arrowhead curve. In fact, both curves yield Sierpiński’s triangle as n goes to infinity. As an aside, Sierpiński’s triangle is a so-called *fractal curve*: it has the Hausdorff dimension $\log 3 / \log 2 \cong 1.58496$, as it consists of three self-similar pieces with magnification factor 2.

8. Gray code

The function *bramah* enumerates the configurations $\{A, B, C\}^n$ changing only one peg at a time. In other words, the succession of configurations corresponds to a *ternary Gray code*! To investigate the correspondence a bit further, here is a version of *bramah* that returns a list of configurations, rather than a list of moves.

$$\begin{aligned} \text{bramah}'_0 \quad (a^b_c) &= [[]] \\ \text{bramah}'_{n+1} \quad (a^b_c) &= a \triangleleft \text{bramah}'_n \ (a^b_c) \\ &\quad \# \ b \triangleleft \text{bramah}'_n \ (c^b_a) \\ &\quad \# \ c \triangleleft \text{bramah}'_n \ (a^b_c) \end{aligned}$$

¹The term gasket graph is actually not used consistently. Some authors call the counterpart of the Hanoi graph the Sierpiński gasket graph, and its contracted variant the Sierpiński graph.




Figure 6. A Hamiltonian path in the Hanoi graph for 4 discs.




Figure 7. A Hamiltonian path in the Hanoi graph for 7 discs.

There are two standard ternary Gray codes: the so-called modular code, the digits vary 012|201|120 \dots , and the reflected code, the digits vary 012|210|012 \dots . The definition above yields the latter, as the discs are only moved between adjacent pegs. In fact, we have $bramah'_n (c^b_a) = reverse (bramah'_n (a^b_c))$. Using this property on the second recursive call, we can simplify $bramah'_n (0^1_2)$ somewhat.

$$\begin{aligned} gray3_0 &= [[]] \\ gray3_{n+1} &= 0 \triangleleft gray3_n \# 1 \triangleleft reverse \ gray3_n \# 2 \triangleleft gray3_n \end{aligned}$$

For comparison, here is the definition of the Gray *binary* sequence.

$$\begin{aligned} gray2_0 &= [[]] \\ gray2_{n+1} &= 0 \triangleleft gray2_n \# 1 \triangleleft reverse \ gray2_n \end{aligned}$$

Actually, the *binary* Gray code is also hidden in the Tower of Hanoi puzzle. We only have to work with a different configuration space: instead of $\{A, B, C\}^n$ we use $\{0, 1\}^n$ keeping track whether a disc has been moved an even or an odd number of times. The initial configuration AA^n becomes 00^n , the final configuration CC^n becomes 10^n . Since the i th disc is moved 2^i times, all the discs make an even number of moves, except for the largest, which makes a single move.

We can easily adapt *ihanoi* to generate binary Gray code. We take as a starting point the first definition of *movei*, slightly modified so that the next configuration is returned instead of the next move (we also applied the *Just a # m = Just a* simplification).

$$\begin{aligned} configi [] \quad c' &= Nothing \\ configi (p : ps) c' & \\ \quad | p \neq (p \perp c') &= Just (p \perp c' : ps) \\ \quad | p == (p \perp c') &= p \triangleleft configi ps (p \perp c') \end{aligned}$$

Now, the Gray code equivalent of \perp is the Boolean operation *exclusive or*, that is, inequality of Booleans. This implies that the last target c' corresponds to a *parity bit*. Thus, *configi* becomes (the code uses Booleans rather than bits)

$$\begin{aligned} codei [] \quad p &= Nothing \\ codei (b : bs) p & \\ \quad | b \neq (b \neq p) &= Just ((b \neq p) : bs) \\ \quad | b == (b \neq p) &= b \triangleleft codei bs (b \neq p) . \end{aligned}$$

Again, we can simplify the code a bit: inequality and equality of Boolean values are associative (Backhouse and Fokkinga, 2001), so $b \neq (b \neq p)$ is simply p . Using *False # b = b* and *True # b = ¬ b*, we obtain

$$\begin{aligned} codei [] \quad p &= Nothing \\ codei (b : bs) p & \\ \quad | p &= Just (¬ b : bs) \\ \quad | ¬ p &= b \triangleleft codei bs b . \end{aligned}$$

In words: we traverse the list $p : bs$ up to the first 1; the following bit, if any, is flipped.

As for *ihanoi*, we have to maintain two pieces of information: the current Gray code and the current *parity bit*. The latter is easy to update: it is flipped in each step. Summing up, we obtain the following Gray code generator.

$$\begin{aligned} igray bs &= map fst (iterate step (bs, foldr (\neq) True bs)) \\ \text{where } step (bs, p) &= \text{do } \{ bs' \leftarrow codei bs p; \\ &\quad \text{return } (bs', \neg p) \} \end{aligned}$$

This is, in fact, the functional version of Knuth's Algorithm G (2005).

9. Conclusion and further reading

We have come to the end of the Tour D'Hanoi. In the spirit of the wholemeal approach we started with an inductive definition of the

Hanoi *graph*. From that we derived a series of programs evolving around the Tower of Hanoi theme. Knowing the big picture was jolly useful: for instance, calculating the number of moves could be reduced to the problem of locating a configuration in the graph. Projective program transformations are abundant: *hanoi* is derived from *graph* by projecting onto the lower side of the graph, *cycle* is derived from *hanoi* by mashing out the moves of the larger discs, and so forth.

The transformations could be made rigorous within the Algebra of Programming framework (Bird and de Moor, 1997). Occasionally, this comes at an additional cost. For instance, to derive *cycle* from *hanoi* we would additionally need the disc number, which is not present in *hanoi*'s output.

A lot is left to explore. There are literally hundreds of papers on the subject: Paul Stockmeyer's comprehensive bibliography has a total of 369 entries (2005). From that bibliography I learned that my definition of the Hanoi graph is not original: Er introduced it to analyse the complexity of the generalised Tower of Hanoi problem (1983). The original instructions of the game already alluded to the recursive procedure for solving the puzzle. It has been used since to illustrate the concept of recursion. The parallel version—usually classified as iterative—is due to Buneman and Levy (1980). Backhouse and Fokkinga (2001) show that each disc cycles around the pegs exploiting the associativity of equivalence. To the best of the author's knowledge the iterative, or stepwise variant is original. The connection to Gray codes has first been noticed by Gardner (1972).

Acknowledgments

Special thanks are due to Daniel James for enjoyable discussions and for suggesting a number of stylistic and presentational improvements. Thanks are also due to the anonymous referees for an effort to make the paper less dense.

References

- Backhouse, Roland, and Maarten Fokkinga. 2001. The associativity of equivalence and the Towers of Hanoi problem. *Information Processing Letters* 77:71–76.
- Bird, Richard, and Oege de Moor. 1997. *Algebra of Programming*. London: Prentice Hall Europe.
- Buneman, Peter, and Leon Levy. 1980. The Towers of Hanoi problem. *Information Processing Letters* 10(4–5):243–244.
- Er, M.C. 1983. An analysis of the generalized Towers of Hanoi problem. *BIT* 23:429–435.
- Gardner, Martin. 1972. Mathematical games: The curious properties of the Gray code and how it can be used to solve puzzles. *Scientific American* 227(2):106–109. Reprinted, with Answer, Addendum, and Bibliography, as Chapter 2 of *Knotted Doughnuts and Other Mathematical Entertainments*, W. H. Freeman and Co., New York, 1986.
- Hinze, Ralf. 2008. Functional Pearl: Streams and Unique Fixed Points. In *Proceedings of the 2008 International Conference on Functional Programming*, ed. Peter Thiemann, 189–200. ACM Press.
- Knuth, Donald E. 2005. *The Art of Computer Programming, Volume 4, Fascicle 2: Generating All Tuples and Permutations*. Addison-Wesley Publishing Company.
- Korittky, Joachim. 1998. *Functional METAPOST*. Diplomarbeit, Universität Bonn.
- Stockmeyer, Paul K. 2005. The Tower of Hanoi: A bibliography. Available from <http://www.cs.wm.edu/~pkstoc/biblio2.pdf>.