



HAL
open science

Playing with the Tower of Hanoi Formally

Laurent Théry

► **To cite this version:**

| Laurent Théry. Playing with the Tower of Hanoi Formally. 2020. hal-02903548

HAL Id: hal-02903548

<https://inria.hal.science/hal-02903548>

Preprint submitted on 21 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Playing with the Tower of Hanoi Formally

Laurent Théry

Stamp Project INRIA, France

Laurent.Theiry@sophia.inria.fr

Abstract

The Tower of Hanoi is a typical example that illustrates all the power of recursion in programming. In this paper, we show that it is also a very nice example for inductive proofs and formal verification. We present some non-trivial results about it that have been formalised in the COQ proof assistant.

1 Introduction

The Tower of Hanoi is often used in computer science course to teach recursion. The puzzle is composed of three pegs and some disks of different sizes. Here is a drawing of the initial configuration for five disks ¹:



Initially, all the disks are piled-up in decreasing order of size on the left peg. The goal is to move them to the right peg. There are two rules. First, only one disk can be moved at a time. Second, a larger disk can never be put on top of a smaller one.

A program that solves this puzzle can easily be written using recursion : one builds the program $P_{3:r:n+1}$ that solves the puzzle for $n + 1$ disks using the program $P_{3:r:n}$ that solves the puzzle for n disks. The algorithm proceeds

¹We use macros designed by Martin Hofmann and Berteun Damman for our drawings.

as follows. We first call $P_{3:r:n}$ to move the top- n disks to the intermediate peg.



Then we move the largest disk to its destination.



Finally, we use $P_{3:r:n}$ to move the n disks on the intermediate peg to their destination.



This simple recursive algorithm is also optimal: it produces the minimal numbers of moves.

In this paper, we consider some variants of this puzzle (with three or four pegs, with some constraints on the moves one can perform) and explain how these puzzles and their optimal solution have been formalised.

2 General setting

We first present the general settings of our formalisation that has been done in COQ using the SSREFLECT extension [1]. Then, we explain more precisely the variants of the puzzle that we have taken into consideration and how they have been formalised. We have tried as much as possible to be precise and present exactly what has been formalised. For this, we adopt the syntax of the Mathematical Component library [9] that we have been using. We believe that most of it is close enough to usual mathematical notations so the reader can figure out what is going on. If this not the case, the reader can refer to [9]. We use the following convention. Notions that are present in the library are written using a typewriter font while our own definitions are written using an italic font.

Disks

A disk is represented by its size. We use the type \mathbb{I}_n of natural numbers strictly smaller than n for this purpose.

Definition $disk\ n := \mathbb{I}_n$.

As there is an implicit conversion from \mathbb{I}_n to natural number, the comparison of the respective size of two disks is simply written as $d_1 < d_2$. Two elements (`ord0` and `ord_max`) are defined for \mathbb{I}_n when n is not zero. We use them to represent the smallest disk and the largest one.

Definition $sdisk : disk\ n.+1 := ord0$.
Definition $ldisk : disk\ n.+1 := ord_max$.

In particular, *ldisk* is the main actor of our proofs by simple induction, it represents the largest disk.

Pegs

We also use \mathbb{I}_n for pegs.

Definition $peg\ k := \mathbb{I}_k$.

We mostly use elements of *peg 3* or *peg 4* but some generic properties hold for *peg k*. An operation associated to pegs is the one that picks a peg that differs from an initial peg p_i and a destination peg p_j . It is written as $p[p_i, p_j]$. Generic and specific properties are derived from it. For example, we have:

Lemma *opeg_sym* $(p_1\ p_2 : peg\ k) : p[p_1, p_2] = p[p_2, p_1]$.
Lemma *opegDl* $(p_1\ p_2 : peg\ k.+3) : p_1 \neq p_2 \rightarrow p[p_1, p_2] \neq p_1$.
Lemma *opeg3Kl* $(p_1\ p_2 : peg\ 3) : p_1 \neq p_2 \rightarrow p[p[p_1, p_2], p_1] = p_2$.

The symmetry is valid for every number of pegs. The property of being distinct is only valid when we have more than three pegs. Finally, there is a version of the pigeon-hole principle for three pegs.

Configuration

Disks are always ordered from the largest to the smallest on a peg. This means that a configuration just needs to record which disk is on which peg. It is then just defined as a finite function from disks to pegs.

Definition *configuration* $k\ n := \{\text{ffun } \text{disk } n \rightarrow \text{peg } k\}$.

From a technical point of view, using finite functions gives for free functional extensionality (which is not valid for usual functions in COQ). As a consequence, we can use the boolean equality `==` to test equality between two configurations.

A configuration is called *perfect* if all its disks are on the same peg. So it is encoded as the constant function:

Definition *perfect* $p := [\text{ffun } d \Rightarrow p]$.

It is written as $c[p]$ in the following, or $c[p, n]$ when the number of disks needs to be given specifically.

Note that this encoding of configurations has the merit of avoiding invalid configurations. The price to pay is that we have to recover some natural notions. One of these notions is the predicate *on_top* $d\ c$ that indicates that the disk d is on top of its peg in the configuration c . It is defined as follows:

Definition *on_top* $(d : \text{disk } n) (c : \text{configuration } k\ n) :=$
 $[\text{forall } d_1 : \text{disk } n, c\ d == c\ d_1 ==> d \leq d_1]$.

For every disk on the same peg as d , its size must be larger than the one of d .

When manipulating configurations and particularly in inductive proofs, one needs to take some part of a configuration or build new configuration. The most common operation that is used in proofs by simple induction is to remove the largest disk or to add a new disk.

Definition *cliftr* ($c : \text{configuration } k \ n$) ($p : \text{peg } k$) : *configuration* $k \ n.+1$.
Definition *cunliftr* ($c : \text{configuration } k \ n.+1$) : *configuration* $k \ n$.

A dedicated notation $\downarrow [c]$ is associated with *cunliftr* c and a corresponding one, $\uparrow [c]_p$, for *cliftr* $c \ p$. A set of basic properties is derived for these operations. For example, we have:

Lemma *cunliftrK* ($c : \text{configuration } k \ n.+1$) : $\uparrow [\downarrow [c]]_c \text{ldisk} = c$.
Lemma *perfect_unliftr* ($p : \text{peg } k$) : $\downarrow [c[p]] = c[p]$.

Similarly, in proofs by strong induction, one may need to take bigger chunk of configuration. One way to do this is to be directed by the type (here the addition).

Definition *emerge* ($c_1 : \text{configuration } k \ m$) ($c_2 : \text{configuration } k \ n$) : *configuration* $k \ (m + n)$.
Definition *clshift* ($c : \text{configuration } k \ (m + n)$) : *configuration* $k \ m$.
Definition *crshift* ($c : \text{configuration } k \ (m + n)$) : *configuration* $k \ n$.

As a matter of fact, $\downarrow [c]$ and $\uparrow [c]_p$ are just defined as a special case of these operators for $m = 1$.

Other kinds of surgery have been defined but they have been so far of a less frequent use:

Definition *ccut* ($C : c \leq n$) ($c : \text{configuration } k \ n$) : *configuration* $k \ c$.
Definition *ctuc* ($C : c \leq n$) ($c : \text{configuration } k \ n$) : *configuration* $k \ (n - c)$.
Definition *cset* ($s : \{\text{set } (\text{disk } n)\}$) ($c : \text{configuration } k \ n$) : *configuration* $k \ \#|s|$.
Definition *cset2* ($sp : \{\text{set } (\text{peg } k)\}$) ($sd : \{\text{set } (\text{disk } n)\}$)
($c : \text{configuration } k \ n$) : *configuration* $\#|sp| \ \#|sd|$.

ccut and *ctuc* are the equivalent of *crshift* and *clshift* but directed by the proof C rather than the type. *cset* considers a subset of disks but with the same number of peg ($\#|s|$ is the cardinal of s). *cset2* is the most general and considers both a subset of disks and a subset of pegs.

Move

A move is defined as a relation between configurations. As we want to possibly add constraints on moves, it is parameterised by a relation r on the pegs: $r\ p_1\ p_2$ indicates that it is possible to go from peg p_1 to peg p_2 . Assumptions are usually added on the relation r (such as irreflexivity or symmetry) in order to prove basic properties of moves.

Here is the formal definition of a move:

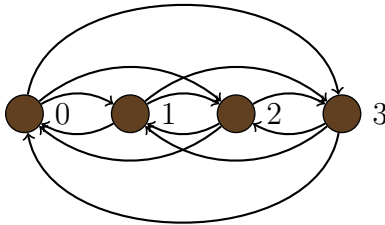
```

Definition move : rel (configuration k n) :=
  [rel c1 c2 | [exists d1 : disk n,
    [&& r (c1 d1) (c2 d1), on_top d1 c1, on_top d1 c2 &
    [forall d2, d1 != d2 ==> c1 d2 == c2 d2]]]
  
```

It simply states that there is a disk d_1 that fulfills 4 conditions:

- the move of d_1 from its peg in c_1 to its peg in c_2 is compatible with r ;
- the disk d_1 is on top of its peg in c_1 ;
- the disk d_1 is on top of its peg in c_2 ;
- it is the unique disk that moves.

The standard puzzle has no restriction on the moves between pegs. If we draw the possible moves as arrows between pegs, the picture for four pegs gives the following complete graph:



We call this version *regular*. It is denoted with the r infix. For example, $P_4:r:5$ corresponds to the puzzle with four pegs and five disks with no restriction on the moves. Its associated relation $rrel$ only enforces irreflexivity:

Definition $rrel : rel (peg\ k) := [\mathbf{rel}\ x\ y \mid x \neq y]$.
Definition $rmove : rel (configuration\ k\ n) := move\ rrel$.

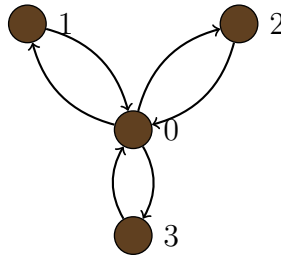
The first variant we consider is where one can only move from one peg to its neighbour. The picture for four pegs is the following:



This version is called *linear* and its associated infix is \mathbf{l} . For example, the puzzle with four pegs and five disks with linear moves is written $P_{4:\mathbf{l}:5}$. The corresponding relation $lrel$ just uses arithmetic to check neighbourhood :

Definition $lrel : rel (peg\ k) := [\mathbf{rel}\ x\ y \mid (x.+1 == y) \mid\mid (y.+1 == x)]$.
Definition $lmove : rel (configuration\ k\ n) := move\ lrel$.

Finally the last variant we consider is the one where one central peg is the only one that can communicate with its outer pegs. A picture for four pegs gives



This version is called *star* and its associated infix is \mathbf{s} . For example, the puzzle with four pegs and five disks with star moves is written $P_{4:\mathbf{s}:5}$. The corresponding relation $srel$ just uses multiplication to put the peg 0 in the center :

Definition $srel : rel (peg\ k) := [\mathbf{rel}\ x\ y \mid (x \neq y) \ \&\& \ (x * y == 0)]$.
Definition $smove : rel (configuration\ k\ n) := move\ srel$.

Note that these categories may overlap when there are few pegs. For example, $P_{3:\mathbf{l}:n}$ and $P_{3:\mathbf{s}:n}$ are identical.

2.1 Path and distance

Moves are defined as relation over configurations. So, we can see sequence of moves as paths on a graph whose nodes are the configurations. As configurations belong to a finite type, we can benefit from the elements of graph theory that are present in the Mathematical Component library. For example, (*rgraph move c*) returns the set of all the configurations that are reachable from *c* in one move, (*connect move c₁ c₂*) indicates that *c₁* and *c₂* can be connected through moves, or (*path move c cs*) gives that the sequence *cs* of configurations is a path that connects *c* with the last element of *cs*.

Now, all the surgery operations on configurations need to be lifted to path. As distinct configurations may become identical when taking sub-part, we first need to define an operation on sequences that removes repetition.

```
Fixpoint rm_rep (A : eqType) (a : A) (s : seq A) :=
  if n is b :: s1 then
    if a == b then rm_rep b s1 else b :: rm_rep b s1
  else [::]
```

It is then possible to derive properties on path. For example, we have:

```
Lemma path_clshift (c : configuration (m + n)) cs :
  path move c cs →
  path move (clshift c) (rm_rep (clshift c) [seq (clshift i) | i <- cs]).
```

As we want to show that some algorithms are optimal, the last ingredient we need is a notion of distance between configurations. Unfortunately, there is no built-in notion of distance in the Mathematical Component library, so we have to define one. For this, we first build recursively the function *connectn* that computes the set of elements that are connected with exactly *n* moves. Then, we can define the distance between two points *x* and *y* as the smallest *n* such as (*connectn r n x y*) holds. It is defined as (*gdist r x y*) and written as *d[x, y]_r* in the following.

Finally, we introduce the notion of geodesic path: a path that realises the distance.

```

Definition gpath r x y p :=
  [&& path r x p, last x p == y & d[x, y]r == size p].

```

Companion theorems are derived for these basic notions. For example, the following lemma shows that concatenation behaves well with respect to distances.

```

Lemma gdist_cat r x y p1 p2 :
  gpath r x y (p1 ++ p2) → d[x, y]r = d[x, last x p1]r + d[last x p1, y]r.

```

3 Puzzles with three pegs

The proofs associated with the puzzles with three pegs are straightforward. They are done by induction on the number of disks inspecting the moves of the largest disk. What makes the simple induction work well with three pegs is that when the largest disk moves from p_i to p_j , all the smaller disks are necessarily piled-up on peg $p[p_i, p_j]$, so they make a perfect configuration on which we can often apply the inductive hypothesis directly.

3.1 Regular puzzle

It is easy to translate in COQ the algorithm described in the introduction. We write it as a recursive function that works on n disks and generate the sequence of configurations that goes from the configuration $c[p_1]$ to the configuration $c[p_2]$:

```

Fixpoint ppeg n p1 p2 :=
  if n is n1.+1 then
    let p3 := p[p1, p2] in
      [seq ↑ [i]p1 | i ← ppeg n1 p1 p3] ++ [seq ↑ [i]p2 | i ← c[p3]] :: ppeg n1 p3 p2
    else [::].

```

It is easy to prove the basic properties of this function

Lemma *size_ppeg* $n\ p_1\ p_2$: $\text{size } (\text{ppeg } n\ p_1\ p_2) = 2^n - 1$
Lemma *last_ppeg* $n\ p_1\ p_2\ c$: $\text{last } c\ (\text{ppeg } n\ p_1\ p_2) = c[p_2]$.
Lemma *path_ppeg* $n\ p_1\ p_2$: $p_1 \neq p_2 \rightarrow \text{path } \text{remove } c[p_1]\ (\text{ppeg } n\ p_1\ p_2)$.

The key property is that *ppeg* builds a path of minimal size. As a matter of fact, we have proved that this minimal path is unique.

Lemma *ppeg_min* $n\ p_1\ p_2\ cs$:
 $p_1 \neq p_2 \rightarrow \text{path } \text{remove } c[p_1]\ cs \rightarrow \text{last } c[p_1]\ cs = c[p_2] \rightarrow$
 $2^n - 1 \leq \text{size } cs \text{ ?= iff } (cs == \text{ppeg } n\ p_1\ p_2)$.

The expression $(e_1 \leq e_2 \text{ ?= iff } C)$ indicates not only that e_1 is smaller than e_2 but also the condition C tells exactly when the comparison between e_1 and e_2 is an equality. The proof is done by double induction (one on the size of cs and one on n) inspecting the moves of the largest disk in the sequence cs . As here p_1 differs from p_2 , we know that the largest disk must move at least one time in cs . If it moves exactly once, the inductive hypothesis on n let us conclude directly. If it moves at least twice, the equality never holds. If the first two moves are on different pegs, adding the inductive hypothesis on $n - 1$ twice plus the two moves of the largest disk gives us a path of size at least $2 \times (2^{n-1} - 1) + 2 = 2^n$. If it moves on one peg and then returns to the peg p_1 , the path has some repetition, so the inductive hypothesis on the size of cs let us conclude.

From this theorem, we easily derive the corollary on the distance.

Lemma *gdist_rhanoi3p* $n\ (p_1\ p_2 : \text{peg } 3)$:
 $d[c[p_1, n], c[p_2, n]]_{\text{remove}} = (2^n - 1) \times (p_1 \neq p_2)$

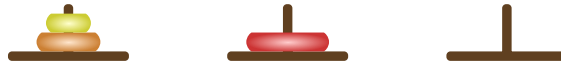
Note that we have used the automatic conversion from boolean to integer (**true** is 1 and **false** is 0) to include the case where p_1 and p_2 are the same peg.

The last theorem only talks about going from a perfect configuration to another one. What about the distance between two arbitrary configurations c_1 and c_2 ? It seems natural to apply the same greedy strategy than for the perfect configuration. We proceed recursively. If c_1 *ldisk* is equal to c_2 *ldisk*, we just perform the recursive call for smaller disks. If they are different, we

perform a recursive call to move the smaller disks in c_1 to an intermediate peg, $p[c_1 \text{ldisk}, c_2 \text{ldisk}]$, then move the largest disk to its position in c_2 , and finally perform another recursive call to move the smaller disks to their position in c_2 . Unfortunately, this strategy is not optimal anymore. We can illustrate this with an example with 3 disks. Let us suppose that we try to go from the initial configuration:



to the final position:



The greedy strategy would move the largest disk only once and would have size five, one for the largest disk, plus two times two moves for the two small disks. The most effective strategy, instead, is to move the largest disk first to the right peg:



and now the greedy strategy gives a solution of size three. So, we have a solution of size four.

We have formalised exactly what the optimal solutions are :

- between an arbitrary configuration and a perfect configuration, the greedy strategy is always optimal;
- between two arbitrary configurations c_1 and c_2 , the optimal strategy just needs to compare the one-jump solution with two-jumps solution only for the largest disk d such that $c_1 d$ differs from $c_2 d$.

3.2 Linear puzzle

Implementing the greedy strategy for the linear puzzle is slightly more complicated. As we can move only between pegs that are neighbour, the largest

disk may need to jump twice. But from the optimality point of view the situation is much simpler, the greedy strategy is always optimal. This is formally proved and we get the expected theorem about distance between perfect configurations:

Lemma *gdist_lhanoi3p* n (p_1 p_2 : peg 3) :
 $d[c[p_1, n], c[p_2, n]]_{lmove} =$
 if *lrel* p_1 p_2 then $(3^n - 1)/2$ else $(3^n - 1) \times (p_1 \neq p_2)$

Note that as there are n disks and 3 pegs, there are 3^n possible configurations. This last theorem tells us than the solution that goes from the perfect configuration where all the disks are on the left peg to the perfect configuration where they are on the right peg visit all the configurations!

4 Puzzles with four pegs

Adding a peg changes completely the situation. If the previous simple recursive algorithm still works, it does not give anymore an optimal solution. The new strategy is implemented by the so-called Frame-Stewart algorithm. We explain it using the regular puzzle with four pegs. Then, we explain how the proofs about the distances of $P_{4:r:n}$ and $P_{4:s:n}$ have been formalised in Coq.

4.1 The Frame-Stewart Algorithm

Let us build, $P_{4:r:n}$, an algorithm that moves the disks from the leftmost peg to the rightmost one for the regular puzzles with four pegs.



We choose an arbitrary m smaller than n and use $P_{4:r:m}$ to move the top- m disks to an intermediate peg.



The remaining $n - m$ disks can now freely move except on this intermediate peg, so we can use $P_{3:r:n-m}$ to move them to their destination.



and reuse $P_{4:r:k}$ to move the top- m disks to their destination.



We can even choose the m parameter as to minimize the number of moves, so we have $|P_{4:r:n}| = \min_{m < n} 2|P_{4:r:m}| + (2^{n-m} - 1)$. This strategy can be generalised to an arbitrary numbers k of pegs, leading to the recurrence relation:

$$|P_{k:r:n}| = \min_{m < n} 2|P_{k:r:m}| + |P_{k-1:r:n-m}|.$$

Knowing if this general program is optimal is an open question but it has been shown optimal for $P_{4:r:n}$ and $P_{4:s:n}$. It is what we have formalized. Note that, from the proving point of view, the new strategy just seems to move from simple induction to strong induction with this new parameter m . As a matter of fact, this new strategy is just a generalisation of the previous one. If we apply it to three pegs, the only way we can move $n - m$ pegs for the $P_{2:r:n-m}$ puzzle is by taking $m = n - 1$.

4.2 Regular puzzle

The proof given in [3] that shows that the Frame-Stewart algorithm is optimal for the regular puzzle with four pegs is rather technical. As a matter of fact

this technicality was a motivation of our formalisation. The proof is very well written and very convincing but contains several cases. A formalisation ensures that no detail has been overlooked. As an anecdote, we had started the formalisation of the proof that the Frame-Stewart is always optimal given in [6]. Stuck on the formalisation of Corollary 3, we have contacted the author that told us that there was a flaw in the proof as documented in [5].

Here, we will only highlight the overall structure of the proof. We refer to the paper proof given in [3] for the details. The first step is to relate $|P_{4:r:n}|$ with triangular numbers. Following [3], we write $|P_{4:r:n}|$ as $\Phi(n)$. We introduce the notation Δn for the sum of the n^{th} first natural numbers :

$$\Delta n = \sum_{i \leq n} i = \frac{n(n+1)}{2}.$$

A number n is triangular if it is a Δi for some i . By analogy to the square root, we introduce the triangular root ∇n :

$$\Delta(\nabla n) \leq n < \Delta(\nabla(n+1)).$$

Now, we can give explicit formula for $\Phi(n)$:

$$\Phi(n) = \sum_{i < n} 2^{\Delta i}$$

It is relatively easy to show that the function Φ verifies the recurrence relation of the Frame-Stewart algorithm: $\Phi(n) = \min_{m < n} 2\Phi(m) + (2^{n-m} - 1)$. Then, what is left to be proved is that it is the optimal solution. The key ingredient of the proof is of course to find the right inductive invariant. This is done thanks to a valuation function Ψ that takes a finite set over the natural numbers and returns a natural number:

$$\Psi E = \max_{L \in \mathbb{N}} ((1-L)2^L - 1 + \sum_{n \in E} 2^{\min(\nabla n, L)})$$

The idea is that E will contain the disks we are interested in. If we consider the set $[n]$ of all the natural numbers smaller than n (i.e. we are interested by all the disks)

$$[n] = \{\text{set } i \mid i < n\}$$

we get back the Φ function we are aiming at:

$$\Psi[n] = \frac{\Phi(n+1) - 1}{2} = \frac{\sum_{i < n} 2^{\nabla(i+1)}}{2}$$

Now we can present the central theorem. We consider two configurations u and v of n disks and the four pegs p_0, p_1, p_2 and p_3 . If v is such that there is no disk on pegs p_0 and p_1 and E is defined as

$$E = \{\text{set } i \mid \text{the disk } i \text{ is on the peg } p_0 \text{ in } u\}$$

the invariant is:

$$d[u, v]_{\text{remove}} \geq \Psi E$$

The proof proceeds by strong induction of the number of disks. It examines a geodesic path p from u to v . If p_2 is the peg where the disk $ldisk$ is in v (it cannot be p_0 nor p_1), it considers T the largest disk that was initially on the peg p_0 and visits at least one time the peg p_3 . If such a disk does not exist, the inequality easily holds. Then, it considers inside the path p the configuration x_0 before which the disk T leaves the peg p_0 for the first time and the configuration x_3 in which the disk T reaches the peg p_3 for the first time. Similarly, it considers the configuration z_0 before which the disk n leaves the peg p_0 for the first time and the configuration z_2 before which the disk n reaches the peg p_2 for the last time. Examining the respective positions of x_0, x_3, z_0 and z_2 in p and applying some surgery on the configurations of the path p in order to fit the inductive hypothesis it concludes that the inequality holds in every cases.

The six-page long proof of the main theorem 2.9 in [3] translates to a 1000-line long Coq proof.

Lemma *gdist_le_psi* ($u \ v : \text{configuration } 4 \ n$) ($p_0 \ p_2 \ p_3 : \text{peg } 4$) :
 $(\wedge \ p_3 \ != \ p_2, \ p_3 \ != \ p_0 \ \& \ p_2 \ != \ p_0) \rightarrow (\text{codom } v) \setminus \text{subset } [:: \ p_2 ; \ p_3] \rightarrow$
 $\Psi \ [\text{set } i \mid u \ i == \ p_0] \leq d[u, v]_{\text{remove}}.$

From which, we easily derive the expected theorem:

Lemma *gdist_rhanoi4* ($n : \text{nat}$) ($p_1 \ p_2 : \text{peg } 4$) :
 $p_1 \ != \ p_2 \rightarrow d[c[p_1, n], c[p_2, n]]_{\text{remove}} = \Phi \ n.$

5 Star puzzle

We first recall how to apply the Frame-Stewart algorithm to the star puzzle. Let us build the program $P_{4:s:n}$ that generates the moves between two perfect

configurations: one on an outer peg p_i ($p_i \neq 0$) and the other on another outer peg p_j ($p_j \neq p_i \neq 0$). We first choose a parameter m and use $P_{4:s:m}$ to move the top- m disk to the third outer peg p_k ($p_k \neq p_j \neq p_i \neq 0$). Now, we use $P_{3:s:n-m}$ (which is identical to $P_{3:1:n-m}$) to move the $n - m$ from peg p_i to peg p_j . Finally, we use $P_{4:s:k}$ to move the top- m disk from peg p_k to peg p_j . This leads to the recurrence relation:

$$|P_{4:s:n}| = \min_{m < n} 2|P_{4:s:m}| + (3^{n-m} - 1).$$

Now, we have to find a mathematical object that verifies this recurrence relation. This time it is not the triangular numbers but the increasing sequence α_1 of the elements $2^i 3^j$. The first elements of this sequence are 1, 2, 3, 4, 6, 8, 9. If we define $S_1(n) = \sum_{i < n} \alpha_1(i)$, it is relatively easy to prove that

$$S_1(n) = \min_{m < n} 2S_1(m) + (3^{n-m} - 1)/2.$$

It follows that $2S_1$ verifies the recurrence relation.

Here, the proof of optimality is even more intricate, we just give an idea of the inductive invariant and how the proof proceeds. We refer to [4] for a detailed and very clear exposition of the proof. The first generalisation is to consider distance not only between two configurations but between $l + 1$ configurations (i.e. a distance between u_0 and u_l passing through u_1, \dots, u_{l-1}): $\sum_{i < l} d[u_i, u_{i+1}]$. These intermediate configurations ($0 < i < l$) are alternating. If p_1, p_2 and p_3 are the outer pegs, the configuration u_i is supposed to have its disks on pegs p_2 and $a[p_1, p_3](i)$ where alternation is defined as

$$a[p_i, p_j](0) = p_i \quad a[p_i, p_j](n + 1) = a[p_j, p_i](n)$$

Taking into account this new parameter l , we need to lift S_1 to a parametrised function S_l .

$$S_l(n) = \min_{m < n} 2S_l(m) + l(3^{n-m} - 1)/2$$

and α_1 to $\alpha_l(n) = S_l(n + 1) - S_l(n)$. Finally, we introduce penalty function β defined as

$$\beta_{n,l}(k) = \text{if } 1 < l \text{ and } k + 1 = n \text{ then } \alpha_l(k) \text{ else } 2\alpha_1(k)$$

Given these definitions, the inductive invariant looks like:

Lemma *main* $(p_1 p_2 p_3 : \text{peg } 4) n l (u : \{\text{ffun } \mathbb{I}_{l,+1} \rightarrow \text{configuration } 4 n\}) :$
 $p_1 \neq p_2 \rightarrow p_1 \neq p_3 \rightarrow p_2 \neq p_3$
 $p_1 \neq p_0 \rightarrow p_2 \neq p_0 \rightarrow p_3 \neq p_0 \rightarrow$
 $(\forall k, 0 < k < l \rightarrow \text{codom } (u k) \text{ subset } [:: p_2; a[p_1, p_3] k]) \rightarrow$
 $(S_{-}[l] n) . *2 \leq \text{sum}_{-}(i < l) d[u i, u i.+1]_{-}\text{smove} +$
 $\text{sum}_{-}(k < n) (u \text{ ord0 } k \neq p_1) * \beta_{-}[n, l] k +$
 $\text{sum}_{-}(k < n) (u \text{ ord}_{\text{max}} k \neq a[p_1, p_3] l) * \beta_{-}[n, l] k.$

The proof is done by simple induction on n . It is then split in several cases depending on the number of elements i such that $u_i(\text{ldisk}) = a[p_1, p_3](i)$. Furthermore, the inductive invariant has this alternating assumption. So, often, one needs to divide the path in a bunch of alternating sub-paths in order to apply the inductive hypothesis on each of these sub-paths. This leads to a lower-bound where various values of $S_i(m)$ appear. Key properties of $S_l(n)$ (i.e. its convexity in n and concavity in l) are then used to derive simpler lower-bound.

The paper proof of the *main* lemma is 15-page long and translates into 3500 lines of COQ proof script. As it contains several crossed references between cases, the formal proof of the *main* lemma is composed of 3 separate sub-lemmas plus one use of the "without loss of generality" tactic [7].

Finally, the Frame-Stewart algorithm gives an upper-bound to the distance and the *main* lemma applied with $l = 1$ gives a lower-bound. Altogether we get the expected theorem:

Lemma *gdist_shanoi4* $(n : \text{nat}) (p_1 p_2 : \text{peg } 4) :$
 $p_1 \neq p_2 \rightarrow p_1 \neq p_0 \rightarrow p_1 \neq p_0 \rightarrow d[c[p_1, n], c[p_2, n]]_{-}\text{smove} = (S_{-}[1] n) . *2.$

6 Conclusion

We have presented a formalisation of the Tower of Hanoi with three and four pegs. Of course, we have only scratched the surface of what can be proved. We refer the reader to [8] for an account of all the mathematical objects this simple puzzle can be linked to.

We started this formalisation as a mere exercise. Then, we got addicted and tried to prove more difficult results. The entire development is available

at <https://github.com/they/hanoi>. An attractive aspect of this formalisation is that it uses very elementary mathematics. So, it is heavily testing the capability to do various flavour of inductive proofs and to manipulate big operators [2]. To give only one example, in order to prove the concavity of the function S_l we had to revisit the merge sort algorithm. It is usually given as a beginner exercise as a way to merge two sorted lists. Here, it is used to merge two increasing functions f and g in an increasing function $fmerge(f, g)$ and we had to prove that

$$fmerge(f, g)(n) = \max_{i \leq n} \min(f(i), g(n - i))$$

When doing it formally, it is very easy to get lost in this kind of proofs.

Finally, the main contribution of this work is the formal proofs about the distance between two perfect configurations for the 4 pegs versions. We believe that they are a natural companion to the paper proofs. These paper proofs are very technical. We have mechanically checked all the details. As a matter of fact, we have been using very little automation, so our formal proofs follow very closely the paper proofs. The main difference is that our formalisation used natural numbers only. So, we have tried to avoid as much as possible subtraction. In our formal setting, $(m - n) + n = m$ is a valid theorem only if we add the assumption that $n \leq m$. So an expression such as $a - b \leq c - d$ in the paper proof is translated into $a + d \leq b + c$ in the formal development. These formal proofs can largely be improved. We plan to add more automation and get more concise proof scripts in the near future.

References

- [1] *The SSReflect proof language*. Section of the Coq refence manual, available at <https://coq.inria.fr/refman/proof-engine/ssreflect-proof-language.html>.
- [2] Yves Bertot, Georges Gonthier, Sidi Ould Biha, and Ioana Paşca. Canonical Big Operators. In *TPHOLS*, volume 5170 of *LNCS*, Montreal, Canada, 2008.
- [3] Thierry Bousch. La quatrième tour de Hanoi. *Bull. Belg. Math. Soc. Simon Stevin*, 21(5):895–912, 2014.

- [4] Thierry Bousch. La Tour de Stockmeyer. *Séminaire Lotharingien de Combinatoire*, 77(B77d), 2017.
- [5] Thierry Bousch, Andreas M. Hinz, Sandi Klavžar, Daniele Parisse, Ciril Petr, and Paul K. Stockmeyer. A note on the Frame-Stewart conjecture. *Discrete Math., Alg. and Appl.*, 11(4), 2019.
- [6] Roberto Demontis. What is the least number of moves needed to solve the k-peg Towers of Hanoi problem? *Discrete Math., Alg. and Appl.*, 11(1), 2019.
- [7] John Harrison. Without loss of generality. In *TPHOLs 2009*, volume 5674 of *LNCS*, pages 43–59, 2009.
- [8] Andreas M. Hinz, Sandi Klavžar, Uroš Milutinović, and Ciril Petr. *The Tower of Hanoi – Myths and Maths*. Birkhäuser, 2013.
- [9] Assia Mahboubi and Enrico Tassi. *Mathematical Components*. available at <https://math-comp.github.io/mcb/book.pdf>.