



HAL
open science

Pancake Flipping is Hard

Laurent Bulteau, Guillaume Fertin, Irena Rusu

► **To cite this version:**

Laurent Bulteau, Guillaume Fertin, Irena Rusu. Pancake Flipping is Hard. 37th International Symposium on Mathematical Foundations of Computer Science (MFCS 2012), Aug 2012, Bratislava, Slovakia. pp.247-258. hal-00717712

HAL Id: hal-00717712

<https://hal.science/hal-00717712v1>

Submitted on 13 Jul 2012

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.

Pancake Flipping Is Hard

Laurent Bulteau, Guillaume Fertin, and Irena Rusu

Laboratoire d'Informatique de Nantes-Atlantique (LINA), UMR CNRS 6241
Université de Nantes, 2 rue de la Houssinière, 44322 Nantes Cedex 3, France
{Laurent.Bulteau, Guillaume.Fertin, Irena.Rusu}@univ-nantes.fr

Abstract. Pancake Flipping is the problem of sorting a stack of pancakes of different sizes (that is, a permutation), when the only allowed operation is to insert a spatula anywhere in the stack and to flip the pancakes above it (that is, to perform a prefix reversal). In the burnt variant, one side of each pancake is marked as burnt, and it is required to finish with all pancakes having the burnt side down. Computing the optimal scenario for any stack of pancakes and determining the worst-case stack for any stack size have been challenges over more than three decades. Beyond being an intriguing combinatorial problem in itself, it also yields applications, e.g. in parallel computing and computational biology. In this paper, we show that the Pancake Flipping problem, in its original (unburnt) variant, is NP-hard, thus answering the long-standing question of its computational complexity.

Keywords. Pancake problem, Computational complexity, Permutations, Prefix reversals.

1 Introduction

The pancake problem was stated in [7] as follows:

The chef in our place is sloppy, and when he prepares a stack of pancakes they come out all different sizes. Therefore, when I deliver them to a customer, on the way to the table I rearrange them (so that the smallest winds up on top, and so on, down to the largest at the bottom) by grabbing several from the top and flipping them over, repeating this (varying the number I flip) as many times as necessary. If there are n pancakes, what is the maximum number of flips (as a function of n) that I will ever have to use to rearrange them?

Stacks of pancakes are represented by permutations, and a flip consists in reversing a prefix of any length. The previous puzzle yields two entangled problems:

- Designing an algorithm that sorts any permutation with a minimum number of flips (this optimization problem is called MIN-SBPR, for Sorting By Prefix Reversals).
- Computing $f(n)$, the maximum number of flips required to sort a permutation of size n (the diameter of the so-called *pancake network*).

Gates and Papadimitriou [9] introduced the *burnt* variant of the problem: the pancakes are two-sided, and an additional constraint requires the pancakes to end with the unburnt side up. The diameter of the corresponding *burnt pancake network* is denoted $g(n)$. A number of studies [4–6, 9, 11–13] have aimed at determining more precisely the values of $f(n)$ and $g(n)$, with the following results:

- $f(n)$ and $g(n)$ are known exactly for $n \leq 19$ and $n \leq 17$, respectively [5].
- $15n/14 \leq f(n) \leq 18n/11 + O(1)$ [12, 4].
- $\lfloor (3n + 3)/2 \rfloor \leq g(n) \leq 2n - 6$ [5] (upper bound for $n \geq 16$).

Considering MIN-SBPR, 2-approximation algorithms have been designed, both for the burnt and unburnt variants [6, 8]. Moreover, Labarre and Cibulka [13] have characterized a subclass of signed permutations, called *simple permutations*, that can be sorted in polynomial time.

The pancake problems have various applications. For instance, the pancake network, having both a small degree and diameter, is of interest in parallel computing. The algorithmic aspect, i.e. the sorting problem, has applications in comparative genomics, since prefix reversals are possible elementary modifications that can affect a genome during evolution. A related problem is Sorting By Reversals [1] where any subsequence can be flipped at any step, not only prefixes. This problem is now well-known, with a polynomial-time exact algorithm [10] for the signed case, and a 1.375-approximation [2] for the APX-hard unsigned case [3].

In this paper, we prove that the MIN-SBPR problem is NP-hard (in its unburnt variant), thus answering a question which has remained open for several decades. We in fact prove a stronger result: it is known that the number of breakpoints of a permutation (that is, the number of pairs of consecutive elements that are not consecutive in the identity permutation) is a lower bound on the number of flips necessary to sort a permutation. We show that deciding whether this bound is tight is already NP-hard.

2 Notations

We denote by $\llbracket a ; b \rrbracket$ the interval $\{a, a + 1, \dots, b\}$. Let n be an integer. Input sequences are permutations of $\llbracket 1 ; n \rrbracket$, hence we consider only sequences where all elements are unsigned, and there cannot be duplicates. We use upper case for sequences, and lower case for elements.

Consider a sequence S of length n , $S = \langle x_1, x_2, \dots, x_n \rangle$. Element x_1 is said to be the *head element* of S . Sequence S has a *breakpoint* at position r , $1 \leq r < n$ if $x_r \neq x_{r+1} - 1$ and $x_r \neq x_{r+1} + 1$. It has a *breakpoint* at position n if $x_n \neq n$. We write $d_b(S)$ the number of breakpoints of S . Note that having $x_1 \neq 1$ does not directly count as a breakpoint, and that $d_b(S) \leq n$ for any sequence of length n . For any $p \leq q \in \mathbb{N}$, we write \mathcal{I}_q^p the sequence $\langle p, p+1, p+2, \dots, q \rangle$. \mathcal{I}_n^1 is the *identity*. For a sequence of any length $S = \langle x_1, x_2, \dots, x_k \rangle$, we write *S the sequence obtained by reversing S : $^*S = \langle x_k, x_{k-1}, \dots, x_1 \rangle$. Given an integer p , we write $p + S = \langle p + x_1, p + x_2, \dots, p + x_k \rangle$.

The *flip* of length r is the operation that consists in reversing the r first elements of the sequence. It transforms

$$\begin{aligned} S &= \langle x_1, x_2, \dots, x_r, x_{r+1}, \dots, x_n \rangle \\ \text{into } S' &= \langle x_r, x_{r-1}, \dots, x_1, x_{r+1}, \dots, x_n \rangle. \end{aligned}$$

Property 1. Given a sequence S' obtained from a sequence S by performing one flip, we have $d_b(S') - d_b(S) \in \{-1, 0, 1\}$.




Fig. 1. Examples of efficient flips. Sequence $\langle 5, 2, 3, 1, 4 \rangle$ is efficiently sortable (in four flips), but $\langle 5, 2, 3, 4, 1 \rangle$ is not.

A flip from S to S' is said to be *efficient* if $d_b(S') = d_b(S) - 1$, and we reserve the notation $S \rightarrow S'$ for such flips. A sequence of size n , different from the identity, is a *deadlock* if it yields no efficient flip, and we write $S \rightarrow \perp$. By convention, we underline in a sequence the positions corresponding to possible efficient flips: there are at most two of them, and at least one if the sequence is neither a deadlock nor the identity.

A *path* is a series of flips, it is *efficient* if each flip is efficient in the series. A sequence S is *efficiently sortable* if there exists an efficient path from S to the identity permutation (equivalently, if it can be sorted in $d_b(S)$ flips). See for example Figure 1.

Let S be a sequence different from the identity, and \mathbb{T} be a set of sequences. We write $S \Longrightarrow \mathbb{T}$ if both following conditions are satisfied:

1. for each $T \in \mathbb{T}$, there exists an efficient path from S to T .
2. for each efficient path from S to the identity, there exists a sequence $T \in \mathbb{T}$ such that the path goes through T .

If \mathbb{T} consists of a single element ($\mathbb{T} = \{T\}$), we may write $S \Longrightarrow T$ instead of $S \Longrightarrow \{T\}$. Note that condition 1. is trivial if $\mathbb{T} = \emptyset$, and condition 2. is trivial if there is no efficient path from S to \mathcal{I}_n^1 . Given a sequence S , there can be several different sets \mathbb{T} such that $S \Longrightarrow \mathbb{T}$. The following properties are easily deduced from the definition of \Longrightarrow .

Property 2. Given any sequence $S \neq \mathcal{I}_n^1$,

$$\begin{aligned}
S \Longrightarrow \mathcal{I}_n^1 &\Leftrightarrow S \text{ is efficiently sortable.} \\
S \Longrightarrow \emptyset &\Leftrightarrow S \text{ is not efficiently sortable.}
\end{aligned}$$

Property 3. If $S \Longrightarrow \{S_1, S_2\}$, $S_1 \Longrightarrow \mathbb{T}_1$ and $S_2 \Longrightarrow \mathbb{T}_2$, then $S \Longrightarrow \mathbb{T}_1 \cup \mathbb{T}_2$.

3 Reduction from 3-SAT

The reduction uses a number of gadget sequences in order to simulate boolean variables and clauses with subsequences. They are organized in two levels (where level-1 gadgets are directly defined by sequences of integers, and level-2 gadgets are defined using a pattern of level-1 gadgets). For each gadget we define, we derive a property characterizing the efficient paths that can be followed if some part of the gadget appears at the head of a sequence. The proofs for all these properties follow the same pattern, with no obstacle apart from the increasing complexity of the sequences, and only the one for the Dock gadget is given in this extended abstract.

3.1 Level-1 gadgets

Dock. The dock gadget is the simplest we define. Its only goal is to store sequences of the kind ${}^*\mathcal{I}_q^{p+1}$ (with $p < q$) out of the head of the sequence, without “disturbing” any other part.

Definition 1. Given two integers p and q with $p < q$, the dock for ${}^*\mathcal{I}_q^{p+1}$ is the sequence $Dock(p, q) = D$, where

$$D = \langle p - 1, p, q + 1, q + 2 \rangle.$$

It has the following property:

Property 4. Let p and q be any integers with $p < q$, $D = Dock(p, q)$, and X and Y be any sequences. We have

$$\langle {}^*\mathcal{I}_q^{p+1}, X, D, Y \rangle \implies \langle X, \mathcal{I}_{q+2}^{p-1}, Y \rangle$$

Proof. An efficient path from $\langle {}^*\mathcal{I}_q^{p+1}, X, D, Y \rangle$ to $\langle X, \mathcal{I}_{q+2}^{p-1}, Y \rangle$ is given by:

$$\begin{aligned} \langle {}^*\mathcal{I}_q^{p+1}, X, D, Y \rangle &= \langle q, q - 1, \dots, p + 2, p + 1, X, p - 1, \underline{p}, q + 1, q + 2, Y \rangle \\ &\rightarrow \langle p, p - 1, \underline{*X}, p + 1, p + 2, \dots, q - 1, q, q + 1, q + 2, Y \rangle \\ &\rightarrow \langle X, p - 1, p, p + 1, p + 2, \dots, q - 1, q, q + 1, q + 2, Y \rangle \\ &= \langle X, \mathcal{I}_{q+2}^{p-1}, Y \rangle \end{aligned}$$

For each sequence in the path, we apply the only possible efficient flip, hence every efficient path between $\langle {}^*\mathcal{I}_q^{p+1}, X, D, Y \rangle$ and \mathcal{I}_n^1 (if such a path exists) begins with these two flips, and goes through $\langle X, \mathcal{I}_{q+2}^{p-1}, Y \rangle$.

Lock. A lock gadget contains three parts: a sequence which is the lock itself, a key element that “opens” the lock, and a test element that checks whether the lock is open.

Definition 2. For any integer p , $Lock(p)$ is defined by $Lock(p) = (key, test, L)$, where

$$\begin{aligned} key &= p + 10 & test &= p + 7 \\ L &= p + \langle 1, 2, 9, 8, 5, 6, 4, 3, 11, 12 \rangle \end{aligned}$$

Given a lock $(key, test, L) = Lock(p)$, we write

$$L^o = p + \langle 1, 2, 3, 4, 6, 5, 8, 9, 10, 11, 12 \rangle.$$

Sequences L and L^o represent the lock when it is closed and open, respectively. If a sequence containing a closed lock has key for head element, then efficient flips put the lock in open position. If it has $test$ for head element, then it is a deadlock if and only if the lock is closed.

Property 5. Let p be any integer, $(key, test, L) = Lock(p)$, and X and Y be any sequences. We have

- a. $\langle key, X, L, Y \rangle \implies \langle X, L^o, Y \rangle$
- b. $\langle test, X, L^o, Y \rangle \implies \langle X, \mathcal{I}_{p+12}^{p+1}, Y \rangle$
- c. $\langle test, X, L, Y \rangle \rightarrow \perp$

We use locks to emulate literals of a boolean formula: variables “hold the keys”, and in a first time open the locks corresponding to true literals. Each clause holds three test elements, corresponding to its three literals, and the clause is true if the lock is open for at least one of the test elements.

Hook. A hook gadget contains four parts: two sequences used as delimiters, a *take* element that takes the interval between the delimiters and places it in head, and a *put* element that does the reverse operation. Thus, the sequence between the delimiters can be stored anywhere until it is called by *take*, and then can be stored back using *put*.

Definition 3. For any integer p , *Hook*(p) is defined by $Hook(p) = (take, put, G, H)$, where

$$\begin{aligned} take &= p + 10 & put &= p + 7 \\ G &= p + \langle 3, 4 \rangle & H &= p + \langle 12, 11, 6, 5, 9, 8, 2, 1 \rangle. \end{aligned}$$

Given a hook $(take, put, G, H) = Hook(p)$, we write

$$\begin{aligned} G' &= p + \langle 12, 11, 6, 5, 4, 3 \rangle & H' &= p + \langle 10, 9, 8, 2, 1 \rangle \\ G'' &= p + \langle 3, 4, 5, 6, 7 \rangle & H'' &= p + \langle 12, 11, 10, 9, 8, 2, 1 \rangle. \end{aligned}$$

Property 6. Let p be an integer, $(take, put, G, H) = Hook(p)$, and X, Y and Z be any sequences. We have

- a. $\langle take, X, G, Y, H, Z \rangle \implies \langle Y, G', *X, H', Z \rangle$
- b. $\langle put, X, G', *Y, H', Z \rangle \implies \langle Y, G'', X, H'', Z \rangle$
- c. $\langle G'', X, H'', Y \rangle \implies \langle X, *\mathcal{I}_{p+12}^{p+1}, Y \rangle$

Fork. A fork gadget implements choices. It contains two parts delimiting a sequence X . Any efficient path encountering a fork gadget follows one of two tracks, where either X or $*X$ appears at the head of the sequence at some point. Sequence X would typically contain a series of triggers for various gadgets (*key*, *take*, etc.), so that X and $*X$ differ in the order in which the gadgets are triggered.

Definition 4. For any integer p , *Fork*(p) is defined by $Fork(p) = (E, F)$, where

$$E = p + \langle 11, 8, 7, 3 \rangle \quad F = p + \langle 10, 9, 6, 12, 13, 4, 5, 15, 14, 2, 1 \rangle.$$

Given a fork $(E, F) = Fork(p)$, we write

$$\begin{aligned} F^1 &= p + \langle 10, 9, 6, 7, 8, 11, 12, 13, 14, 15, 5, 4, 3, 2, 1 \rangle \\ F^2 &= p + \langle 3, 7, 8, 11, 10, 9, 6, 12, 13, 4, 5, 15, 14, 2, 1 \rangle \end{aligned}$$

Property 7. Let p be an integer, $(E, F) = \text{Fork}(p)$, and X, Y be any sequences. We have

- a. $\langle E, X, F, Y \rangle \implies \{ \langle X, F^1, Y \rangle, \langle *X, F^2, Y \rangle \}$
- b. $\langle F^1, Y \rangle \implies \langle *T_{p+15}^{p+1}, Y \rangle$
- c. $\langle F^2, Y \rangle \implies \langle *T_{p+15}^{p+1}, Y \rangle$

3.2 Level-2 gadgets

Literals. The following gadget is used only once in the reduction. It contains the locks corresponding to all literals of the formula.

Definition 5. Let p and m be two integers, $\text{Literals}(p, m)$ is defined by

$$\begin{aligned} \text{Literals}(p, m) &= (key_1, \dots, key_m, test_1, \dots, test_m, \Lambda) \\ \text{where } \forall i \in \llbracket 1; m \rrbracket, (key_i, test_i, L_i) &= \text{Lock}(p + 12(i - 1)) \\ \Lambda &= \langle L_1, L_2, \dots, L_m \rangle \end{aligned}$$

Let O and I be two disjoint subsets of $\llbracket 1; m \rrbracket$. We use Λ_I^O for the sequence obtained from Λ by replacing L_i by L_i^O for all $i \in O$ and by $T_{p+12i}^{p+12i-11}$ for all $i \in I$.

Elements of O correspond to open locks in Λ_I^O , while elements of I correspond to open locks which have moreover been tested. Note that $\Lambda_\emptyset^\emptyset = \Lambda$, and that $\Lambda_{\llbracket 1; m \rrbracket}^\emptyset = T_{p+12m}^{p+1}$.

Property 8. Let p and m be two integers, O and I be two disjoint subsets of $\llbracket 1; m \rrbracket$, $(key_1, \dots, key_m, test_1, \dots, test_m, \Lambda) = \text{Literals}(p, m)$, and X be any sequence. We have

- a. $\forall i \in \llbracket 1; m \rrbracket - O - I, \langle key_i, X, \Lambda_I^O \rangle \implies \langle X, \Lambda_I^{O \cup \{i\}} \rangle$
- b. $\forall i \in O, \langle test_i, X, \Lambda_I^O \rangle \implies \langle X, \Lambda_{I \cup \{i\}}^{O - \{i\}} \rangle$
- c. $\forall i \in \llbracket 1; m \rrbracket - O, \langle test_i, X, \Lambda_I^O \rangle \rightarrow \perp$

Variable. In the rest of this section, we assume that p_Λ and m are two fixed integers, and we define the gadget $(key_1, \dots, key_m, test_1, \dots, test_m, \Lambda) = \text{Literals}(p_\Lambda, m)$. Thus, we can use elements key_i and $test_i$ for $i \in \llbracket 1; m \rrbracket$, and sequences Λ_I^O for any disjoint subsets O and I of $\llbracket 1; m \rrbracket$.

We now define a gadget simulating a boolean variable x_i . It holds two series of *key* elements: the ones with indices in P (resp. N) open the locks corresponding to literals of the form x_i (resp. $\neg x_i$). When the triggering element, ν , is brought to the head, a choice has to be made between P and N , and the locks associated with the chosen set (and only them) are opened.

Definition 6. Let P, N be two disjoint subsets of $\llbracket 1; m \rrbracket$ ($P = \{p_1, p_2, \dots, p_q\}$, $N = \{n_1, n_2, \dots, n_{q'}\}$) and p be an integer, $\text{Variable}(P, N, p)$ is defined by

$$\text{Variable}(P, N, p) = (\nu, V, D)$$

$$\text{where } (take, put, G, H) = \text{Hook}(p + 2), \quad (E, F) = \text{Fork}(p + 14),$$

in $\nu = take$

$$V = \langle G, E, key_{p_1}, \dots, key_{p_q}, put, key_{n_1}, \dots, key_{n_{q'}}, F, H \rangle$$

$$D = \text{Dock}(p + 2, p + 29)$$

Given a variable gadget $(\nu, V, D) = \text{Variable}(P, N, p)$, we write

$$V^1 = \langle G'', key_{n_1}, \dots, key_{n_{q'}}, F^1, H'' \rangle$$

$$V^2 = \langle G'', key_{p_q}, \dots, key_{p_1}, F^2, H'' \rangle$$

where G'', H'', F^1, F^2 , come from the definitions of *Hook* (Definition 3) and *Fork* (Definition 4).

The following property determines the possible behavior of a variable gadget.

Property 9. Let P, N be two disjoint subsets of $\llbracket 1; m \rrbracket$, p be an integer, X and Y be two sequences, O, I be two disjoint subsets of $\llbracket 1; m \rrbracket$, and $(\nu, V, D) = \text{Variable}(P, N, p)$. For sub-property (a.) we require that $(P \cup N) \cap (O \cup I) = \emptyset$, for (b.) that $N \cap (O \cup I) = \emptyset$, and for (c.) that $P \cap (O \cup I) = \emptyset$ (these conditions are in fact necessarily satisfied by construction since all sequences considered are permutations). We have

$$\begin{aligned} \text{a. } & \langle \nu, X, V, Y, A_I^O \rangle \implies \left\{ \begin{array}{l} \langle X, V^1, Y, A_I^{O \cup P} \rangle, \\ \langle X, V^2, Y, A_I^{O \cup N} \rangle \end{array} \right\} \\ \text{b. } & \langle V^1, X, D, Y, A_I^O \rangle \implies \langle X, \mathcal{I}_{p+31}^{p+1}, Y, A_I^{O \cup N} \rangle \\ \text{c. } & \langle V^2, X, D, Y, A_I^O \rangle \implies \langle X, \mathcal{I}_{p+31}^{p+1}, Y, A_I^{O \cup P} \rangle \end{aligned}$$

Clause. The following gadget simulates a 3-clause in a boolean formula. It holds the *test* elements for three locks, corresponding to three literals. When the triggering element, γ , is at the head of a sequence, three distinct efficient paths may be followed. In each such path, one of the three locks is tested: in other words, any efficient path leading to the identity requires one of the locks to be open.

Definition 7. Let $a, b, c \in \llbracket 1; m \rrbracket$ be pairwise distinct integers and p be an integer, $\text{Clause}(a, b, c, p)$ is defined by

$$\text{Clause}(a, b, c, p) = (\gamma, \Gamma, \Delta)$$

$$\text{where } (E_1, F_1) = \text{Fork}(p + 2), \quad (take_1, put_1, G_1, H_1) = \text{Hook}(p + 21), \\ (E_2, F_2) = \text{Fork}(p + 45), \quad (take_2, put_2, G_2, H_2) = \text{Hook}(p + 33),$$

in $\gamma = take_1$

$$\Gamma = \langle G_1, E_1, take_2, put_1, test_c, F_1, G_2, E_2, test_a, put_2, test_b, F_2, H_2, H_1 \rangle$$

$$\Delta = \langle \text{Dock}(p + 2, p + 17), \text{Dock}(p + 21, p + 60) \rangle$$

Given a clause gadget $(\gamma, \Gamma, \Delta) = \text{Clause}(a, b, c, p)$, we write

$$\begin{aligned}\Gamma^1 &= \langle G_1'', \text{test}_c, F_1^1, G_2'', \text{test}_b, F_2^1, H_2'', H_1'' \rangle \\ \Gamma^2 &= \langle G_1'', \text{test}_c, F_1^1, G_2'', \text{test}_a, F_2^2, H_2'', H_1'' \rangle \\ \Gamma^3 &= \langle G_1'', \text{take}_2, F_1^2, G_2, E_2, \text{test}_a, \text{put}_2, \text{test}_b, F_2, H_2, H_1'' \rangle\end{aligned}$$

The following two properties determine the possible behavior of a clause gadget. The main point is that, starting from a sequence $\langle \gamma, X, \Gamma, Y, \Lambda_I^O \rangle$, there is one efficient path for each true literal in the clause (ie. each literal with index in O).

Property 10. Let X and Y be any sequences, and O, I be two disjoint subsets of $\llbracket 1; m \rrbracket$. We have

$$\langle \gamma, X, \Gamma, Y, \Lambda_I^O \rangle \implies \mathbb{T},$$

where \mathbb{T} contains from 0 to 3 sequences, and is defined by:

$$\begin{aligned}\langle X, \Gamma^1, Y, \Lambda_{I \cup \{a\}}^{O - \{a\}} \rangle &\in \mathbb{T} \text{ iff } a \in O \\ \langle X, \Gamma^2, Y, \Lambda_{I \cup \{b\}}^{O - \{b\}} \rangle &\in \mathbb{T} \text{ iff } b \in O \\ \langle X, \Gamma^3, Y, \Lambda_{I \cup \{c\}}^{O - \{c\}} \rangle &\in \mathbb{T} \text{ iff } c \in O\end{aligned}$$

Property 11. Let Y and Z be any sequences, and O, I be two disjoint subsets of $\llbracket 1; m \rrbracket$. We have

- a. If $b, c \in O$, then $\langle \Gamma^1, Y, \Delta, Z, \Lambda_I^O \rangle \implies \langle Y, \mathcal{I}_{p+62}^{p+1}, Z, \Lambda_{I \cup \{b, c\}}^{O - \{b, c\}} \rangle$
- b. If $a, c \in O$, then $\langle \Gamma^2, Y, \Delta, Z, \Lambda_I^O \rangle \implies \langle Y, \mathcal{I}_{p+62}^{p+1}, Z, \Lambda_{I \cup \{a, c\}}^{O - \{a, c\}} \rangle$
- c. If $a, b \in O$, then $\langle \Gamma^3, Y, \Delta, Z, \Lambda_I^O \rangle \implies \langle Y, \mathcal{I}_{p+62}^{p+1}, Z, \Lambda_{I \cup \{a, b\}}^{O - \{a, b\}} \rangle$

3.3 Reduction

Let ϕ be a boolean formula over l variables in conjunctive normal form, such that each clause contains exactly three literals. We write k the number of clauses, $m = 3k$ the total number of literals, and $\{\lambda_1, \dots, \lambda_m\}$ the set of literals. Let $n = 31l + 62k + 12m$.

Definition 8. We define the sequence S_ϕ as the permutation of $\llbracket 1; n \rrbracket$ obtained by:

$$\begin{aligned}(key_1, \dots, key_m, test_1, \dots, test_m, \Lambda) &= \text{Literals}(31l + 62k, m) \\ \forall i \in \llbracket 1; l \rrbracket, & \quad P_i = \{j \in \llbracket 1; m \rrbracket \mid \lambda_j = x_i\} \\ & \quad N_i = \{j \in \llbracket 1; m \rrbracket \mid \lambda_j = \neg x_i\} \\ & \quad (\nu_i, V_i, D_i) = \text{Variable}(P_i, N_i, 31(i - 1)), \\ \forall i \in \llbracket 1; k \rrbracket, & \quad (a_i, b_i, c_i) = \text{indices such that the } i\text{-th clause of } \phi \text{ is } \lambda_{a_i} \vee \lambda_{b_i} \vee \lambda_{c_i} \\ & \quad (\gamma_i, \Gamma_i, \Delta_i) = \text{Clause}(a_i, b_i, c_i, 31l + 62(i - 1)) \\ S_\phi &= \langle \nu_1, \dots, \nu_l, \gamma_1, \dots, \gamma_k, V_1, \dots, V_l, \Gamma_1, \dots, \Gamma_k, D_1, \dots, D_l, \Delta_1, \dots, \Delta_k, \Lambda_\emptyset^\emptyset \rangle\end{aligned}$$

Two things should be noted in this definition. First, elements key_i and $test_i$ are used in the clause and variable gadgets, although they are not explicitly stated in the parameters (cf. Definitions 6 and 7). Second, one could assume that literals are sorted in the formula ($\phi = (\lambda_1 \vee \lambda_2 \vee \lambda_3) \wedge \dots$), so that $a_i = 3i - 2$, $b_i = 3i - 1$ and $c_i = 3i$, but it is not necessary since these values are not used in the following.

We now aim at proving Theorem 1 (p. 11), which states that S_ϕ is efficiently sortable if and only if the formula ϕ is satisfiable. Several preliminary lemmas are necessary, and the overall process is illustrated in Figure 2.

Variable assignment.

Definition 9. A full assignment is a partition $\mathcal{P} = (T, F)$ of $\llbracket 1; l \rrbracket$. Using notations from Definition 8, we define the sequence $S_\phi[\mathcal{P}]$ by:

$$\begin{aligned} \text{For all } i \in \llbracket 1; l \rrbracket, \quad V'_i &= \begin{cases} V_i^1 & \text{if } i \in T \\ V_i^2 & \text{if } i \in F \end{cases} \\ O &= \bigcup_{i \in T} P_i \cup \bigcup_{i \in F} N_i \\ S_\phi[\mathcal{P}] &= \langle \gamma_1, \dots, \gamma_k, V'_1, \dots, V'_l, \Gamma_1, \dots, \Gamma_k, D_1, \dots, D_l, \Delta_1, \dots, \Delta_k, A_\emptyset^O \rangle \end{aligned}$$

With the following lemma, we ensure that any sequence of efficient flips from S_ϕ begins with a full assignment of the boolean variables, and every possible assignment can be reached using only efficient flips.

Lemma 1.

$$S_\phi \implies \{S_\phi[\mathcal{P}] \mid \mathcal{P} \text{ full assignment}\}$$

Going through clauses. Now that each variable is assigned a boolean value, we need to verify with each clause that this assignment satisfies the formula ϕ . This is done by selecting, for each clause, a literal which is true, and testing the corresponding lock. As in Definition 8, for any $i \in \llbracket 1; k \rrbracket$ we write (a_i, b_i, c_i) the indices such that the i -th clause of ϕ is $\lambda_{a_i} \vee \lambda_{b_i} \vee \lambda_{c_i}$ (thus, $a_i, b_i, c_i \in \llbracket 1; m \rrbracket$).

Definition 10. Let \mathcal{P} be a full assignment. A full selection σ is a subset of $\llbracket 1; m \rrbracket$ such that, for each $i \in \llbracket 1; k \rrbracket$, $|\{a_i, b_i, c_i\} \cap \sigma| = 1$ (hence $|\sigma| = k$). A full selection σ and a full assignment $\mathcal{P} = (T, F)$ are compatible, if, for every $i \in \sigma$, literal λ_i is true according to assignment \mathcal{P} (that is, $\lambda_i = x_j$ and $j \in T$, or $\lambda_i = \neg x_j$ and $j \in F$). Given a full selection σ and a full assignment $\mathcal{P} = (T, F)$ which are compatible, we define the sequence $S_\phi[\mathcal{P}, \sigma]$ by:

$$\begin{aligned} \forall i \in \llbracket 1; l \rrbracket, \quad V'_i &= \begin{cases} V_i^1 & \text{if } i \in T \\ V_i^2 & \text{if } i \in F \end{cases} & \forall i \in \llbracket 1; k \rrbracket, \quad \Gamma'_i &= \begin{cases} \Gamma_i^1 & \text{if } a_i \in \sigma \\ \Gamma_i^2 & \text{if } b_i \in \sigma \\ \Gamma_i^3 & \text{if } c_i \in \sigma \end{cases} \\ O &= \bigcup_{i \in T} P_i \cup \bigcup_{i \in F} N_i - \sigma & I &= \sigma \\ S_\phi[\mathcal{P}, \sigma] &= \langle V'_1, \dots, V'_l, \Gamma'_1, \dots, \Gamma'_k, D_1, \dots, D_l, \Delta_1, \dots, \Delta_k, A_I^O \rangle \end{aligned}$$

$$S_\phi = \langle \nu_1, \dots, \nu_l, \gamma_1, \dots, \gamma_k, V_1, \dots, V_l, \Gamma_1, \dots, \Gamma_k, D_1, \dots, D_l, \Delta_1, \dots, \Delta_k, A_\theta^0 \rangle$$

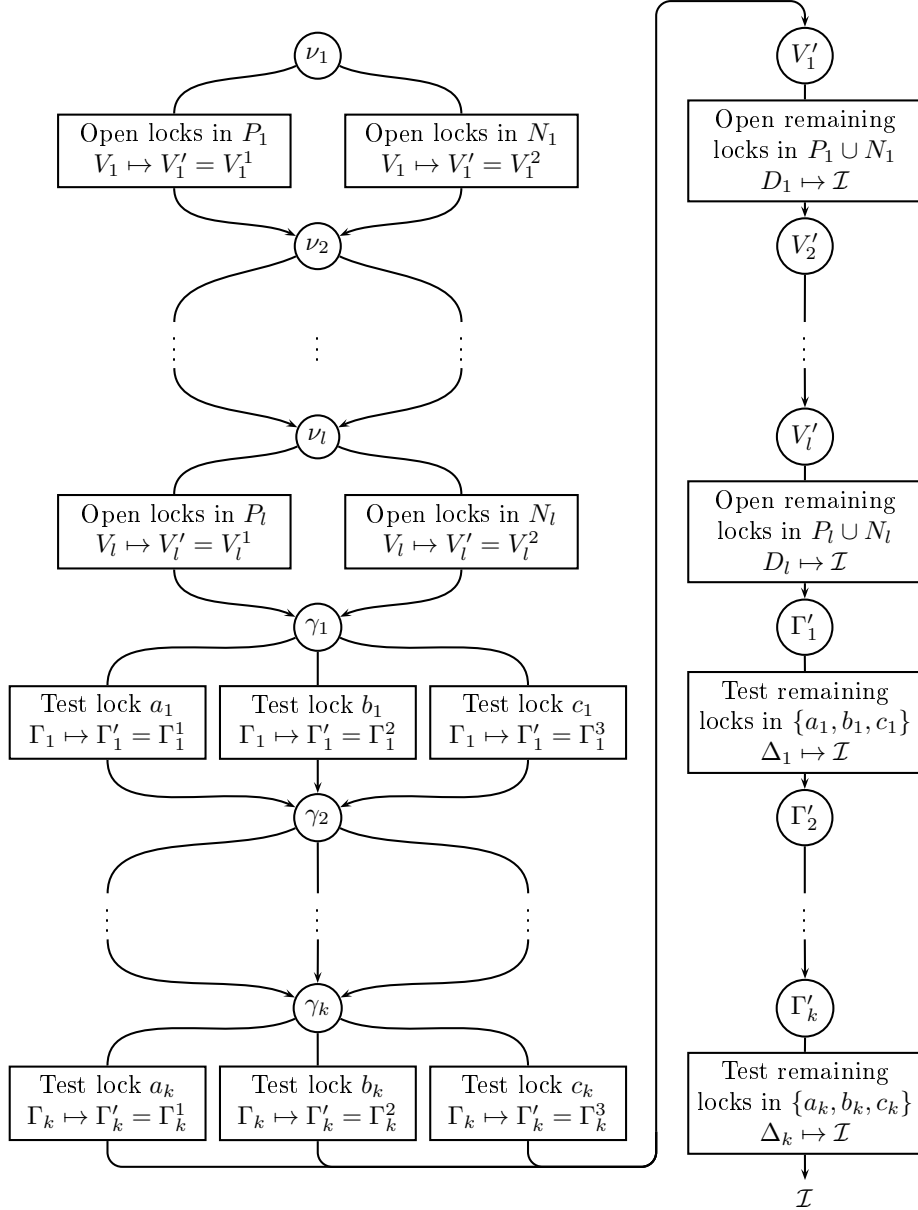


Fig. 2. Description of an efficient sorting of S_ϕ . Circular nodes correspond to head elements or sequences especially relevant (landmarks). We start with the head element of S_ϕ : ν_1 . From each landmark, one, two or three paths are possible before reaching the next landmark, each path having its own effects, stated in rectangles, on the sequence. Possible effects are: transforming a subsequence of S_ϕ (symbol \mapsto), opening a lock, testing a lock (such a path requires the lock to be open).

With the following lemma, we ensure that after the truth assignment, every efficient path starting from S_ϕ needs to select a literal in each clause, under the constraint that the selection is compatible with the assignment.

Lemma 2. *Let \mathcal{P} be a full assignment. Then*

$$S_\phi[\mathcal{P}] \implies \{S_\phi[\mathcal{P}, \sigma] \mid \sigma \text{ full selection compatible with } \mathcal{P}\}$$

Beyond clauses.

Lemma 3. *Let \mathcal{P} be a full assignment and σ be a full selection, such that \mathcal{P} and σ are compatible (provided such a pair exists for ϕ). Then*

$$S_\phi[\mathcal{P}, \sigma] \implies \mathcal{I}_n^1$$

Theorem 1.

$$S_\phi \implies \mathcal{I}_n^1 \text{ iff } \phi \text{ is satisfiable.}$$

Proof. Assume first that $S_\phi \implies \mathcal{I}_n^1$. By Lemma 1, there exists a full assignment $\mathcal{P} = (T, F)$ such that some path from S_ϕ to the identity uses $S_\phi[\mathcal{P}]$. Note that $S_\phi[\mathcal{P}] \implies \mathcal{I}_n^1$. Now, by Lemma 2, there exists a full selection σ , compatible with \mathcal{P} , such that some path from $S_\phi[\mathcal{P}]$ to the identity uses $S_\phi[\mathcal{P}, \sigma]$. Consider the truth assignment $x_i := \text{True} \Leftrightarrow i \in T$. Then each clause of ϕ contains at least one literal that is true (the literal whose index is in σ), and thus ϕ is satisfiable.

Assume now that ϕ is satisfiable: consider any truth assignment making ϕ true, write T the set of indices such that $x_i = \text{True}$, and $F = \llbracket 1; l \rrbracket - T$. Write also σ a set containing, for each clause of ϕ , the index of one literal being true under this assignment. Then σ is a full selection, compatible with the full assignment $\mathcal{P} = (T, F)$. By Lemmas 1, 2 and 3 respectively, there exist efficient paths $S_\phi \implies S_\phi[\mathcal{P}]$, $S_\phi[\mathcal{P}] \implies S_\phi[\mathcal{P}, \sigma]$ and $S_\phi[\mathcal{P}, \sigma] \implies \mathcal{I}_n^1$. Thus sequence S_ϕ is efficiently sortable.

Using Theorem 1, we can now prove the main result of the paper.

Theorem 2. *The following problems are NP-hard:*

- *Sorting By Prefix Reversals (MIN-SBPR)*
- *deciding, given a sequence S , whether S can be sorted in $d_b(S)$ flips*

Proof. By reduction from 3-SAT. Given any formula ϕ , create S_ϕ (see Definition 8, the construction requires a linear time). By Theorem 1, the minimum number of flips necessary to sort S_ϕ is $d_b(S_\phi)$ iff ϕ is satisfiable.

4 Conclusion

In this paper, we have shown that the Pancake Flipping problem is NP-hard, thus answering a long-standing open question. We have also provided a stronger result, namely, deciding whether a permutation can be sorted with no more than one flip per breakpoint is also NP-hard. However, the approximability of MIN-SBPR is still open: it can be

seen that sequence S_ϕ can be sorted in $d_b(S_\phi) + 2$ flips, whatever the formula ϕ , hence this construction does not prove the APX-hardness of the problem.

Among related important problems, the last one having an open complexity is now the burnt variant of the Pancake Flipping problem. An interesting insight into this problem is given in a recent work from Labarre and Cibulka [13], where the authors characterize a subclass of permutations that can be sorted in polynomial time, using the breakpoint graph [1]. Another development consists in trying to improve the approximation ratio of 2 for the Pancake Flipping problem, both in its burnt and unburnt versions.

References

1. V. Bafna and P. Pevzner. Genome rearrangements and sorting by reversals. In *FOCS*, pages 148–157. IEEE, 1993.
2. P. Berman, S. Hannenhalli, and M. Karpinski. 1.375-approximation algorithm for sorting by reversals. In R. Möhring and R. Raman, editors, *ESA*, volume 2461 of *Lecture Notes in Computer Science*, pages 200–210. Springer, 2002.
3. P. Berman and M. Karpinski. On some tighter inapproximability results (extended abstract). In J. Wiedermann, P. van Emde Boas, and M. Nielsen, editors, *ICALP*, volume 1644 of *Lecture Notes in Computer Science*, pages 200–209. Springer, 1999.
4. B. Chitturi, W. Fahle, Z. Meng, L. Morales, C.O. Shields, I. Sudborough, and W. Voit. An $(18/11)n$ upper bound for sorting by prefix reversals. *Theoretical Computer Science*, 410(36):3372–3390, 2009.
5. J. Cibulka. On average and highest number of flips in pancake sorting. *Theoretical Computer Science*, 412(8-10):822–834, 2011.
6. D. Cohen and M. Blum. On the problem of sorting burnt pancakes. *Discrete Applied Mathematics*, 61(2):105–120, 1995.
7. H. Dweighter [pseudonym of J. E. Goodman]. *American Mathematics Monthly*, 82(1), 1975.
8. J. Fischer and S. Ginzinger. A 2-approximation algorithm for sorting by prefix reversals. In G. S. Brodal and S. Leonardi, editors, *ESA*, volume 3669 of *Lecture Notes in Computer Science*, pages 415–425. Springer, 2005.
9. W. Gates and C. Papadimitriou. Bounds for sorting by prefix reversal. *Discrete Mathematics*, 27(1):47–57, 1979.
10. S. Hannenhalli and P. Pevzner. Transforming cabbage into turnip: polynomial algorithm for sorting signed permutations by reversals. In *STOC*, pages 178–189. ACM, 1995.
11. M. Heydari and I. Sudborough. On sorting by prefix reversals and the diameter of pancake networks. In *Proceedings of the First Heinz Nixdorf Symposium on Parallel Architectures and Their Efficient Use*, pages 218–227, London, UK, 1993. Springer-Verlag.
12. M. Heydari and I. Sudborough. On the diameter of the pancake network. *Journal of Algorithms*, 25(1):67–94, October 1997.
13. A. Labarre and J. Cibulka. Polynomial-time sortable stacks of burnt pancakes. *Theoretical Computer Science*, 412(8-10):695–702, 2011.