Journal of Discrete Algorithms 36 (2016) 18–26

Contents lists available at ScienceDirect

Journal of Discrete Algorithms

www.elsevier.com/locate/jda

Folding a paper strip to minimize thickness ☆

Erik D. Demaine [a], David Eppstein [b], Adam Hesterberg [c], Hiro Ito [d], Anna Lubiw [e], Ryuhei Uehara [f,*], Yushi Uno [g]

[a] Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, USA
[b] Computer Science Department, University of California, Irvine, USA
[c] Department of Mathematics, Massachusetts Institute of Technology, USA
[d] School of Informatics and Engineering, University of Electro-Communications, Tokyo, Japan
[e] David R. Cheriton School of Computer Science, University of Waterloo, Ontario, Canada
[f] School of Information Science, Japan Advanced Institute of Science and Technology, Ishikawa, Japan
[g] Graduate School of Science, Osaka Prefecture University, Osaka, Japan

A R T I C L E I N F O

Article history:
Available online 28 October 2015

Keywords:
Linkage
NP-complete
Optimization problem
Rigid origami

A B S T R A C T

In this paper, we study how to fold a specified origami crease pattern in order to minimize the impact of paper thickness. Specifically, origami designs are often expressed by a mountain–valley pattern (plane graph of creases with relative fold orientations), but in general this specification is consistent with exponentially many possible folded states. We analyze the complexity of finding the best consistent folded state according to two metrics: minimizing the total number of layers in the folded state (so that a "flat folding" is indeed close to flat), and minimizing the total amount of paper required to execute the folding (where "thicker" creases consume more paper). We prove both problems strongly NP-complete even for 1D folding. On the other hand, we prove both problems fixed-parameter tractable in 1D with respect to the number of layers.

1. Introduction

Most results in computational origami design assume an idealized, zero-thickness piece of paper. This approach has been highly successful, revolutionizing artistic origami over the past few decades. Surprisingly complex origami designs are possible to fold with real paper thanks in part to thin and strong paper (such as made by Origamido Studio) and perhaps also to some unstated and unproved properties of existing design algorithms.

This paper is one of the few attempting to model and optimize the effect of positive paper thickness. Specifically, we consider an origami design specified by a *mountain–valley pattern* (a crease pattern plus a mountain-or-valley assignment for each crease), which in practice is a common specification for complex origami designs. Such patterns only partly specify a folded state, which also consists of an *overlap order* among regions of paper. In general, there can be exponentially many overlap orders consistent with a given mountain–valley pattern [7]. Furthermore, it is NP-hard to decide flat foldability of a mountain–valley pattern, or to find a valid flat folded state (overlap order) given the promise of flat foldability [2]. But for 1D pieces of paper, the same problems are polynomially solvable [1,3], opening the door to optimizing the effects of paper thickness among the exponentially many possible flat folded states — the topic of this paper.

☆ A preliminary version of this draft was presented at the 9th International Workshop on Algorithms and Computation (WALCOM 2015).

* Corresponding author.
E-mail addresses: edemaine@mit.edu (E.D. Demaine), eppstein@uci.edu (D. Eppstein), achester@mit.edu (A. Hesterberg), itohiro@uec.ac.jp (H. Ito), alubiw@uwaterloo.ca (A. Lubiw), uehara@jaist.ac.jp (R. Uehara), uno@mi.s.osakafu-u.ac.jp (Y. Uno).

http://dx.doi.org/10.1016/j.jda.2015.09.003
1570-8667/

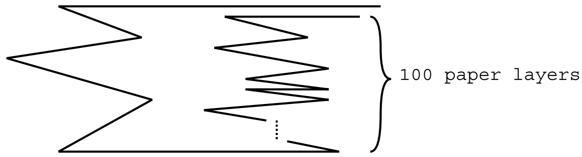

Fig. 1. How can we count the paper layers?

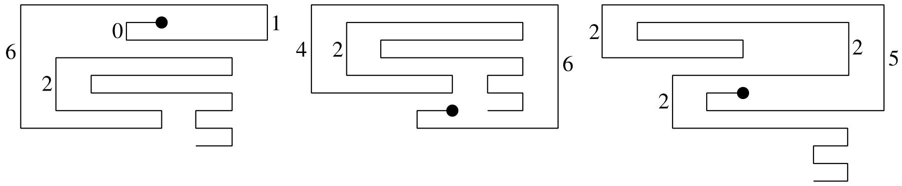

Fig. 2. Three different folded states of the crease pattern $VMVMVVMMMM$ (ending at the dot). The positive crease width of each crease is given beside its corresponding vertical segment. Each folding is better than the other two in one of the three measures, where h is the height, m is the maximum crease width, and t is the total crease width: (1) $h = 11, \underline{m = 5}, t = 11$, (2) $\underline{h = 8}, m = 6, t = 12$, and (3) $h = 9, m = 6, \underline{t = 9}$.

One of the first mathematical studies about paper thickness is also primarily about 1D paper. Britney Gallivan [4], as a high school student, modeled and analyzed the effect of repeatedly folding a positive-thickness piece of paper in half. Specifically, she observed that creases consume a length of paper proportional to the number of layers they must "wrap around", and thereby computed the total length of paper (relative to the paper thickness) required to fold in half n times. She then set the world record by folding a 4000-foot-long piece of (toilet) paper in half twelve times, experimentally confirming her model and analysis.

Motivated by Gallivan's model, Uehara [6] defined the *stretch* at a crease to be the number of layers of paper in the folded state that lie between the two paper segments hinged at the crease. We will follow the terminology of Umesato et al. [8] who later replaced the term "stretch" with *crease width*, which we adopt here. Both papers considered the case of a strip of paper with *equally spaced* creases but an arbitrary mountain–valley assignment. When the mountain–valley assignment is uniformly random, its expected number of consistent folded states is $\Theta(1.65^n)$ [7]. Uehara [6] asked whether it is NP-hard, for a given mountain–valley assignment, to minimize the maximum crease width or to minimize the total crease width (summed over all creases). Umesato et al. [8] showed that the first problem is indeed NP-hard, while the second problem is fixed-parameter tractable.

We consider the problem of minimizing crease width in the more general situation where the creases are not equally spaced along the strip of paper. This more general case has some significant differences with the equally spaced case. For one thing, if the creases are equally spaced, all mountain–valley patterns can be folded flat by repeatedly folding from the rightmost end; in contrast, in the general case, some mountain–valley patterns (and even some crease patterns) have no consistent flat folded state that avoids self-intersection. Flat foldability of a mountain–valley pattern can be checked in linear time [1], [3, Sec. 12.1], but it requires a nontrivial algorithm.

For creases that are not equally spaced, the notion of crease width must also be defined more precisely, because it is not so clear how to count the layers of paper between two segments at a crease. For example, in Fig. 1, although no layers of paper come all the way to touch the three creases on the left, we want the sum of their crease widths to be 100.

We consider a folded state to be an assignment of the segments to horizontal *levels* at integer y coordinates, with the creases becoming vertical segments of variable lengths. See Fig. 2 and the formal definition below. Then the *crease width* at a crease is simply the number of levels in between the levels of the two segments of paper joined by the crease. That is, it is one less than the length of the vertical segment assigned to the crease. In the case of equally spaced creases, this is the number of layers of paper between the two horizontal segments at the crease, so we have generalized the previous definition. Analogous to Uehara's open problems [6], we will study the problems of minimizing the maximum crease width and minimizing the total crease width for a given mountain–valley pattern. The total crease width corresponds to the extra length of paper needed to fold the paper strip using paper of positive thickness, naturally generalizing Gallivan's work [4].[1]

In the setting where creases need not be equally spaced, there is another sensible measure of thickness: the *height* of the folded state is the total number of levels. The height is always $n + 1$ for n equally spaced creases, but in our setting different folds of the same crease pattern can have different heights. Fig. 2 shows how the three measures can differ. Of course, the maximum crease width is always less than the height.

Our main results (Section 3) are NP-hardness of the problem of minimizing height and the problem of minimizing the total crease width. See Table 1. In addition, we show in Section 4 that the problem of minimizing height is fixed-parameter tractable, by giving a dynamic programming algorithm that runs in $O(2^{O(h \log h)} n + n \log n)$ time, where h is the minimum

[1] In this figure, we assume orthogonal bends to make the notions clear. On the other hand, Gallivan measures turns as circular arcs, this changes the length by only a constant factor. Gallivan's model seems to correspond better to practice.