The 14th Annual Meeting of the Asian Association for Algorithms and Computation

October 22-24, 2021, Tainan, Taiwan

AAAC 2021

http://aaac2021.ie.nthu.edu.tw/

Contents

I. Agenda

Friday, October 22, 2021

16:00 – 17:00	**Tutorial:** Voronoi and Voronoi-like Diagrams Evanthia Papadopoulou Chair: Chung-Shou Liao
17:00 – 19:00	Welcome Reception

Saturday, October 23, 2021

08:30 – 08:50	Registration
08:50 – 09:00	Opening: D. T. Lee
09:00 – 10:00	**Keynote Speech:** Bounded Hanoi Kazuo Iwama Chair: D. T. Lee
10:00 – 10:20	Coffee/Tea Break
10:20 – 12:00	Session A1 Online & Approximation Algorithms Session Chair: Chung-Shou Liao Session B1 Geometric Computing Session Chair: Hee-Kap Ahn
12:00 – 14:00	Lunch
14:00 – 15:40	Session A2 Miscellaneous Topics Session Chair: Po-An Chen Session B2 Online & Approximation Algorithms Session Chair: Hyung-Chan An
15:40 – 16:00	Coffee/Tea Break
16:00 – 17:00	**Keynote Speech:** Broadcast and Epidemics on Random Networks Luca Trevisan Chair: Kai-Min Chung
17:00 – 19:00	Conference Banquet

Sunday, October 24, 2021

14:00 – 15:40	Session A3 Graph Theory & Graph Algorithms Session Chair: Wing-Kai Hon Session B3 Miscellaneous Topics Session Chair: Hirotaka Ono
15:40 – 16:00	Coffee/Tea Break
16:00 – 17:00	**Keynote Speech:** Scheduling to Optimize Energy and Electricity Cost Prudence Wong Chair: Siu-Wing Cheng
17:00 – 17:10	Closing: Siu-Wing Cheng Best Student Presentation Award
18:00 – 21:00	COCOON 2021 Reception

Session A1 Online & Approximation Algorithms (Session Chair: Chung-Shou Liao)

- 10:20-10:40 Ya-Chun Liang, Kuan-Yun Lai, Ho-Lin Chen and Kazuo Iwama. Tight Competitive Analyses of Online Car-sharing Problems
- 10:40-11:00 Hao-Ping Yeh, Wei Lu, Li-Hsuan Chen, Ling-Ju Hung, Ralf Klasing and Sun-Yuan Hsieh. Approximation Algorithms for the Star p-Hub Center Routing Problem
- 11:00-11:20 Shi-Chun Tsai, Meng-Tsung Tsai and Tsung-Ta Wu. An Empirical Study of Finding a Most Frequent Fraction and Its Applications
- 11:20-11:40 Yi-Chang Liang and Hung-Lung Wang. The colorful knapsack center problem
- 11:40-12:00 Sheng-Yen Ko, Ho-Lin Chen, Siu-Wing Cheng, Wing-Kai Hon and Chung-Shou Liao. General Max-Min Fair Allocation

Session B1 Geometric Computing (Session Chair: Hee-Kap Ahn)

- 10:20-10:40 Andrew Bloch-Hansen, Roberto Solis-Oba and Andy Yu. High Multiplicity Strip Packing with Three Rectangle Types
- 10:40-11:00 Siu-Wing Cheng and Man Ting Wong. Self-Improving Voronoi Construction for a Hidden Mixture of Product Distributions
- 11:00-11:20 Hwi Kim, Jaegun Lee and Hee-Kap Ahn. Rectangular Partitions of a Rectilinear Polygon
- 11:20-11:40 Jongmin Choi, Dahye Jeong and Hee-Kap Ahn. Covering Convex Polygons by Two Congruent Disks

- 11:40-12:00 Mincheol Kim and Hee-Kap Ahn. Minimum-Link Shortest Paths for Polygons amidst Rectilinear Obstacles

Session A2 Miscellaneous Topics (Session Chair: Po-An Chen)

- 14:00-14:20 Ching-Hsiang Lin and Wing-Kai Hon. Duality Theorem in Kontsevich's Pebble Game and Its Generalization
- 14:20-14:40 Po-An Chen, Yiling Chen, Chi-Jen Lu and Chuang-Chieh Lin. Profitable Prediction Market Making
- 14:40-15:00 Yu-Lun Wu and Hung-Lung Wang. Correcting matrix products over the ring of integers
- 15:00-15:20 Kaito Suzuki, Diptarama Hendrian, Ryo Yoshinaka and Ayumi Shinohara. Query Learning of Symbolic Weighted Finite Automata
- 15:20-15:40 Yu-Hsuan Huang, Yao-Ching Hsieh and Mi-Ying Huang. Accountable Ring Signature From Isogeny Group Action

Session B2 Online & Approximation Algorithms (Session Chair: Hyung-Chan An)

- 14:00-14:20 Yongho Shin and Hyung-Chan An. Making Three Out of Two: Three-Way Online Correlated Selection
- 14:20-14:40 Taehoon Ahn, Jongmin Choi, Chaeyoon Chung, Hee-Kap Ahn, Sang Won Bae and Sang Duk Yoon. Rearranging a Sequence of Points onto a Line
- 14:40-15:00 Byeonguk Kang, Jongmin Choi and Hee-Kap Ahn. Intersecting Disks using Two Congruent Disks
- 15:00-15:20 Fu-Hong Liu, Hsiang-Hsuan Liu and Prudence W.H. Wong. Greedy is Optimal for Online Restricted Assignment and Smart Grid Scheduling for Unit Size Jobs
- 15:20-15:40 Jonathan Toole-Charignon and Hsiang-Hsuan Liu. Online Independent Set with Amortized Late Accept/Reject

Session A3 Graph Theory & Graph Algorithms (Session Chair: Wing-Kai Hon)

- 14:00-14:20 Hsiao-Yu Hu, Ya-Chun Liang, Jian-Xi Shao and Chung-Shou Liao. Learning-Augmented Algorithms for Online TSP
- 14:20-14:40 Cheng-Hung Chiang and Meng-Tsung Tsai. Single-Pass Streaming Algorithms to Partition Graphs into Few Forests
- 14:40-15:00 Chun-Hsiang Chan, Cheng-Yu Shih and Ho-Lin Chen. On the Computational Power of Phosphate Transfer Reaction Networks
- 15:00-15:20 Dun-Wei Cheng, Jo-Yi Chang, Chen-Yen Lin, Limei Lin, Yanze Huang, Krishnaiyan Thulasiraman and Sun-Yuan Hsieh. Node Failure Survivability: An Efficient Logical Topology Mapping Algorithm for IP-over-WDM Optical Networks

- 15:20-15:40 Dun-Wei Cheng, Kai-Hsun Yao and Sun-Yuan Hsieh. The Construction of Multiple Independent Spanning Trees on Generalized Recursive Circulant Graphs

Session B3 Miscellaneous Topics (Session Chair: Hirotaka Ono)

- 14:00-14:20 Hiroshi Eto, Hironori Kiya and Hirotaka Ono. Hardness Results on Generalized Puyopuyo
- 14:20-14:40 Taekang Eom, Seungjun Lee and Hee-Kap Ahn. Largest similar copies of convex polygons in polygon
- 14:40-15:00 Corentin Allair and Antoine Vigneron. Pattern Matching in Doubling Spaces
- 15:00-15:20 Jihoon Hyun, Sewon Park and Martin Ziegler. Lazy Data Types
- 15:20-15:40 Koya Watanabe, Diptarama Hendrian, Ryo Yoshinaka, Takashi Horiyama and Ayumi Shinohara. Efficient Construction of Cryptarithm Catalogues over Deterministic Finite Automata

II. Keynote Speeches

- Kazuo Iwama (RIMS, Kyoto University)
- Title: Bounded Hanoi
- Abstract:

  The classic Tower of Hanoi puzzle involves moving a set of disks on three pegs. The number of moves required for a given number of disks is easy to determine, but when the number of pegs is increased to four or more this becomes more challenging. After 75 years the answer for four pegs was resolved only recently, and this \emph{time complexity} question remains open for five or more pegs. In this article the \emph{space complexity}, i.e., how many disks need to be accommodated on the pegs involved in the transfer, is considered for the first time. Suppose m disks are to be transferred from some peg L to another peg R using k intermediate \emph{work pegs} of sizes $j_1,\ldots,j_k$, then how large can m be? We denote this value by $H(j_1,\ldots,j_k)$. If $k=1$, as in the classic problem, the answer is easy: $H(j)=j+1$. We have the exact value for two work pegs, but so far only very partial results for three or more pegs. For example, $H(10!,10!)=26336386137601$ and $H(0!,1!,2!,...,10!)=16304749471397$, but we still do not know the value for $H(1,i,j)$ except for very small i and j. This is a joint work with Mike Paterson, University of Warwick and will appear in AMM.

- Luca Trevisan (Bocconi University, Italy)
- Title: Broadcast and Epidemics on Random Networks
- Abstract:
  we discuss two processes on random networks.

  First we discuss the flooding process, in which information is broadcast in a network in such a way that every informed node immediately informs all the neighbors. In dynamic networks in which nodes continually enter and exit the network, activating and deactivating random network links, we present a result showing that the process quickly converges to a state in which almost all nodes, or even all nodes, are informed, depending on whether or not dropped connections are replaced by new connections.

  Then we discuss the SIR process of epidemic spreading applied to a model of random networks similar to the Watts-Strogatz model, in which there is a mix of "local" connections and "random long distance" connections. We establish thresholds for the value of R0 that leads to large-scale epidemic spreading and show that even this very simple model is able to recover a number of realistic features, such as the way the epidemic spreads via a series of local outbreaks, and how even a small number of "super-spreader" events can have a disproportionate impact.

  We present these two results together to emphasize how they both reduce to similar questions (how large are the connected components of certain random graphs) that can be addressed with similar techniques (analyze a BFS-like exploration of the graph, delaying decisions about random edges as much as possible).

  (Based on joint papers with L. Becchetti, A. Clementi, R. Denni, F. Pasquale, I. Ziccardi)

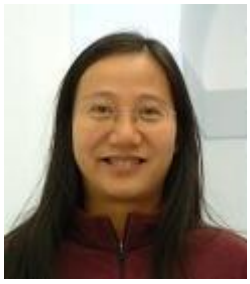

- Prudence Wong (University of Liverpool, UK)
- Title: Scheduling to optimize energy and electricity cost
- Abstract:

  Energy usage is a big concern these days in terms of computation and household usage. This motivates the revisit of classical scheduling problems to take energy into concern. In this talk, we will give an overview of several scheduling problems that attempt to optimize energy and electricity cost. In terms of processor scheduling, we investigate how to use speed scaling and sleep management to reduce energy usage effectively while providing certain level of quality of service. We also investigate how multi-processor scheduling can help reducing energy usage. In terms of household usage, we investigate the so called demand response management in electricity grid. We will also explore the relations of electricity grid scheduling and classical machine scheduling.

III. Tutorial Talk

- Evanthia Papadopoulou (University of Lugano, Switzerland)
- Title: Voronoi and Voronoi-like diagrams
- Abstract:

  Voronoi diagrams are versatile geometric partitioning structures that find diverse applications in Science and Engineering. Given a set of n simple geometric objects, called sites, their Voronoi diagram subdivides the surrounding space into regions of influence exerted by the given sites. These sites are often considered to be points, however, non-points such as line segments, circles, polygons, or polyhedra often model various realistic scenarios. Abstract Voronoi diagrams (AVDs) offer a unifying framework for many such constructs in the plane. In this talk, I will first survey fundamental differences between Voronoi diagrams of points and their counterparts of segments, circles, or AVDs. Because of these differences, some surprising open problems may still remain. For example, although linear-time algorithms for site-deletion in planar point Voronoi diagrams had been well-known to exist since the late 80's, until recently no corresponding algorithms existed for non-point diagrams. Towards bridging such gaps, I will introduce abstract Voronoi-like diagrams, a relaxed Voronoi structure, whose flexibility can help design simple, yet efficient algorithms. A Voronoi-like diagram is a graph on the arrangement of the underlying bisector system whose (non-leaf) vertices are locally Voronoi, i.e., they are vertices in a Voronoi diagram of three sites. Using Voronoi-like graphs we can devise simple randomized incremental constructions under the general AVD framework. I will show this technique and also its analysis, which introduces a simple alternative to the standard backwards analysis, applicable to order-dependent structures. We envision that Voronoi-like graphs will turn out useful in various generalized scenarios, including Voronoi diagrams with disconnected regions and Voronoi diagrams in 3D.

IV. Accepted Papers

Tight Competitive Analyses of Online Car-sharing Problems

Ya-Chun Liang[1], Kuan-Yun Lai[1], Ho-Lin Chen[2], Kazuo Iwama[1]
[1]National Tsing Hua University
[2]National Taiwan University

The car-sharing problem which has received much attention in recent years mainly focuses on an online model in which there are two locations: 0 and 1, and k total cars. Each request which specifies its pick-up time and pick-up location (among 0 and 1, and the other is the drop-off location) is released in each stage a fixed amount of time before its specified start (i.e. pick-up) time. The time between the booking (i.e. released) time and the start time is enough to move empty cars between 0 and 1 for relocation if they are not used in that stage. The model, called kS2L-F, assumes that requests in each stage arrive sequentially regardless of the same booking time and the decision (accept or reject) must be made immediately. The goal is to accept as many requests as possible. In spite of only two locations, the analysis does not seem easy and the (tight) competitive ratio is only known to be 2.0 for $k = 2$ and 1.5 for a restricted value of k, i.e., a multiple of three.

In this study, we remove all the holes of unknown competitive ratios; namely we prove that the competitive ratio is $\frac{2k}{k+\lfloor k/3 \rfloor}$ for all $k \geq 2$. Furthermore, if the algorithm can delay its decision until all requests have come in each stage, the competitive ratio is improved to roughly $4/3$. We can take this advantage even further, precisely we can achieve a competitive ratio of $\frac{2+R}{3}$ if the number of requests in each stage is at most Rk, $1 \leq R \leq 2$, where we do not have to know the value of R in advance. Finally, we demonstrate that randomization also helps to get (slightly) better competitive ratios.

Approximation Algorithms for the Star p-Hub Center Routing Problem

Hao-Ping Yeh[a], Wei Lu[a], Li-Hsuan Chen[a], Ling-Ju Hung[b], Ralf Klasing[c], Sun-Yuan Hsieh[a,d,*]

[a] *Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, Taiwan*
[b] *Department of Creative Technologies and Product Design, National Taipei University of Business, No. 321, Sec. 1, Jinan Rd., Zhongzheng District, Taipei City 100, Taiwan*
[c] *CNRS, LaBRI, Université de Bordeaux, 351 Cours de la Libération, 33405 Talence cedex, France*
[d] *Institute of Medical Information, Institute of Manufacturing Information and Systems, Center for Innovative FinTech Business Models, and International Center for the Scientific Development of Shrimp Aquaculture, National Cheng Kung University, No. 1, University Road, Tainan 701, Taiwan*

Abstract

Given a metric graph $G = (V, E, w)$, a specific vertex $c \in V$, and an integer p, let T be a depth-2 spanning tree of G rooted at c such that c is adjacent to p vertices called hubs and each of the remaining vertices is adjacent to a hub. The Star p-Hub Center Routing Problem is to find a spanning tree T of G and minimize the sum of distances between all pairs of vertices in T. In this paper, we prove that the Star p-Hub Center Routing Problem is NP-hard. It is shown that the Star p-Hub Center Routing Problem is at least as hard as the well-known NP-hard problem Exact Cover by 3-Sets Problem. Moreover, we present a 3-approximation algorithm and a 4-approximation algorithm for the same problem. Both algorithms run in time $O(n^2)$ where n is the number of vertices in the input graph. On the other hand, we give a counterexample for the 4-approximation algorithm (see Figure 1). In the counterexample, the output of the 4-approximation algorithm achieves an approximation ratio which is very close to 4.

Keywords: Hub Allocation, Analysis of Algorithms and Problem Complexity, Approximation Algorithms

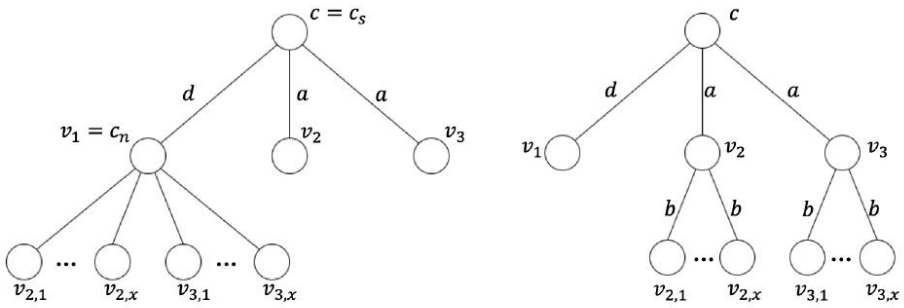

Figure 1: A counterexample for the 4-approximation algorithm, where a, b, d are positive numbers, $a \gg b$, $a > d$, and $a - d \approx 0$. The tree on the left side is the output of the 4-approximation algorithm. The tree on the right side is an optimal solution.

*Corresponding author.
Email addresses: P76081459@mail.ncku.edu.tw (Hao-Ping Yeh), p76041166@mail.ncku.edu.tw (Wei Lu), clh100p@cs.ccu.edu.tw (Li-Hsuan Chen), ljhung@ntub.edu.tw (Ling-Ju Hung), ralf.klasing@labri.fr (Ralf Klasing), hsiehsy@mail.ncku.edu.tw (Sun-Yuan Hsieh)

An Empirical Study of Finding a Most Frequent Fraction and Its Applications

Shi-Chun Tsai [*1], Meng-Tsung Tsai [†2], and Tsung-Ta Wu [‡1]

[1]National Yang Ming Chiao Tung University
[2]Academia Sinica

Given a multiset of n fractions $a_1/b_1, a_2/b_2, \ldots, a_n/b_n$ where a_i, b_i are integers for each $i \in [n]$, find a value v that appears most frequently in the multiset. Several approaches can solve this problem efficiently with incurring no precision error, such as:

- Sort the n fractions with a compare function implemented by cross multiplication, followed by a linear scan to find v.
- Reduce the n fractions to their lowest terms, followed by hashing the reduced fractions into a table to find v.

However, these standard approaches run much slower than an error-prone approach, i.e., dividing the numerators by the denominators using a sufficient number of digits of precision. We compare the above approaches empirically, and find that the performance gain of the error-prone approach can be as large as $3-5$ times for a wide range of parameter settings. This implies that some applications of finding a most frequent fraction also can be sped up drastically, such as testing whether a given set of integral points are in general position, etc.

[*]sctsai@nycu.edu.tw
[†]mttsai@iis.sinica.edu.tw
[‡]ttwu1998.cs09@nycu.edu.tw

The colorful knapsack center problem

Yi-Chang Liang, Hung-Lung Wang

Department of Computer Science and Information Engineering
National Taiwan Normal University, Taipei, Taiwan
(60847041s,hlwang)@gapps.ntnu.edu.tw

Abstract. In this paper, we are concerned with the colorful knapsack center problem, in which one is asked to deploy center facilities so that a requested number of given points are "covered", subject to a budget constraint. Formally, given is a 5-tuple $(X, d, w, g, c, \mathcal{P})$, where (X, d) is a metric space with d the distance measure on X, $w\colon X \to \mathbb{R}_{\geqslant 0}$ a weight function, g a nonnegative number called the budget, c a vector of integers called the coverage requirements, and $\mathcal{P}$ a vector of subsets of X specifying the colors the the points. Assume that the lengths of c and $\mathcal{P}$ are equal and finite. The objective is to find the smallest number ρ, called the radius, such that there is a subset $S \subseteq X$ satisfying $\sum_{s\in S} w(s) \leqslant g$ and for every i $|\{p \in P_i \mid \min_{s\in S} d(p, s) \leqslant \rho\}| \geqslant c_i$, where P_i and c_i stand for the ith entry of c and $\mathcal{P}$, respectively. We call the former constraint the budget constraint and the latter the coverage constraint.

The colorful knapsack center problem is NP-hard, and several results related to approximation have been proposed. In [Danny Z Chen, Jian Li, Hongyu Liang, Haitao Wang. Matroid and knapsack center problems. Algorithmica 75 (2016)] a bi-criteria approximation was proposed for c having length 1: for any fixed $\epsilon > 0$, the algorithm computes a set of centers that uses three times the radius to fulfil the coverage constraint, but may violate the budget constraint by using at most $(1 + \epsilon)g$.

We simplify the analysis by employing the technique of linear programming, which is also widely applied in deriving approximation results for the problems of deloying center facilities. Specifically, our result is based on merging the formulation in [Sayan Bandyapadhyay, Tanmay Inamdar, Shreyas Pai, Kasturi Varadarajan. A Constant Approximation for Colorful k-Center. ESA 2019] with the partial enumeration technique applied by Chen, Li, Liang, and Wang. We also show that the algorithm can be extended for c having fixed length.

General Max-Min Fair Allocation

Sheng-Yen Ko[1], Ho-Lin Chen[2], Siu-Wing Cheng[3], Wing-Kai Hon[4], and Chung-Shou Liao[1]

[1] Department of Industrial Engineering and Engineering Management, National Tsing Hua University, Hsinchu 30013, Taiwan, s107034524@m107.nthu.edu.tw,csliao@ie.nthu.edu.tw

[2] Department of Electrical Engineering, National Taiwan University, Taipei 106, Taiwan, holinchen@ntu.edu.tw

[3] Department of Computer Science and Engineering, HKUST, Hong Kong, scheng@cse.ust.hk

[4] Department of Computer Science, National Tsing Hua University, Hsinchu 30013, Taiwan, wkhon@cs.nthu.edu.tw

In the general max-min fair allocation, also known as the Santa Claus problem, there are m players and n indivisible resources, each player has his/her own utilities for the resources, and the goal is to find an assignment that maximizes the minimum total utility of resources assigned to a player. We introduce an over-estimation strategy to help overcome the challenges of each resource having different utilities for different players. When all resource utilities are positive, we transform it to the machine covering problem and find a $\left(\frac{c}{1-\epsilon}\right)$-approximate allocation in polynomial running time for any fixed $\epsilon \in (0,1)$, where c is the maximum ratio of the largest utility to the smallest utility of any resource. When some resource utilities are zero, we apply the approximation algorithm of Cheng and Mao [ICALP 2019, pp. 38:1–38:13] for the restricted max-min fair allocation problem. It gives a $\left(1+3\hat{c}+O(\delta\hat{c}^2)\right)$-approximate allocation in polynomial time for any fixed $\delta \in (0,1)$, where $\hat{c}$ is the maximum ratio of the largest utility to the smallest positive utility of any resource. The approximation ratios are reasonable if c and $\hat{c}$ are small constants; for example, when the players rate the resources on a 5-point scale.

Keywords: Max-min allocation, hypergraph matching, approximation algorithms

High Multiplicity Strip Packing with Three Rectangle Types

Andrew Bloch-Hansen, Roberto Solis-Oba, Andy Yu
Department of Computer Science, Western University

In the *two-dimensional high multiplicity strip packing problem* (2DHMSPP), we are given K distinct rectangle types, where each rectangle type T_i has n_i rectangles each with width $0 < w_i \leq 1$ and height $0 < h_i \leq 1$. The goal is to pack these rectangles into a strip of width 1, without rotating or overlapping the rectangles, such that the total height of the packing is minimized. We consider the problem for the case when there are 3 different rectangle types, and we designed an algorithm that produces solutions of value at most $OPT + \frac{3}{2} + \epsilon$ for any constant $\epsilon > 0$.

Forty years of research supports the more general two-dimensional strip packing problem, with Harren et al.'s algorithm's approximation ratio of $(\frac{5}{3} + \epsilon)OPT$ and Jansen and Solis-Oba's APTAS being the best known algorithms for the problem, but very little research has been done on the high multiplicity variant. High multiplicity problems are important because the number of distinct object types is typically small in practice and algorithms for high multiplicity problems generally produce solutions closer to optimal than algorithms for the original problems.

2DHMSPP can be relaxed to the *two-dimensional fractional strip packing problem* (2DFSPP), which permits horizontal cuts on the rectangles, and can be represented using the Configuration Linear Program. Since 2DFSPP is identical to the fractional bin packing problem, we can use an algorithm of Karmarkar and Karp to compute a basic feasible solution to the linear program in polynomial time of height at most $LIN + \epsilon$, where LIN is the height of an optimal fractional packing and $\epsilon > 0$. A basic feasible solution to this linear program consists of a set of at most K configurations stacked on top of one another.

A simple algorithm for 2DHMSPP is to compute a basic feasible solution to the above linear program and then simply replace each fractional rectangle with a whole one of the corresponding type, shifting surrounding rectangles upwards as necessary. This algorithm computes a solution to 2DHMSPP of height at most $LIN + K + \epsilon$. Improving on this algorithm is challenging. Our algorithm computes solutions of value at most $LIN + \frac{3}{2} + \epsilon$; this bound is almost tight as there are instances in which a fractional packing and an integer one differ by 1.

Our algorithm takes as input the fractional packing computed by the linear program. We first partition the packing into a common section, in which a particular rectangle type appears in all K configurations, and an uncommon section. Transforming fractional rectangles into whole ones is trivial for the common section, so our contributions focus on the uncommon section of the packing. We sort the fractional rectangles within each configuration by non-decreasing fractional values. We then partition the uncommon portion of the packing into vertical columns whose boundaries are defined by the points where two rectangles of different types are packed side-by-side. We group the vertical columns into different cases depending on the heights of the fractional rectangles within each column.

When the heights of fractional rectangles within a vertical column are "small", we transform them into fractional rectangles of full height (but the same area) and pack them in a new region of height 1. Fractional rectangles of the same type are merged into whole rectangles whenever possible. When the heights of fractional rectangles within a vertical column are "large", we replace them with whole rectangles of the corresponding types.

A major technical challenge for the design of the algorithm is how to define the notions of "small" and "large" and how to combine the two above techniques when the heights of fractional rectangles in a vertical column sum to a "moderate" value (neither "small" nor "large"), as we then use the first technique on some of the fractional rectangles and the second technique on others within a vertical column. We have to carefully choose when and how to combine our techniques to ensure that we do not increase the height of the packing by too much.

We note that our algorithm might vertically split a rectangle and transform its pieces using different techniques. We structure the packing so that the empty space left by moving re-shaped fractional rectangles can be used to ensure that any such vertically split rectangles are re-united with their counterparts so that the output of our algorithm is an integer packing.

Note that algorithms for the general strip packing problem need to determine the position of every rectangle in a solution; however, doing so for 2DHMSPP requires identifying the positions of an exponential number of rectangles with respect to the input size, as the input to 2DHMSPP is described using only $3K$ numbers. Designing a polynomial time algorithm for 2DHMSPP is very complicated as the number of operations must be polynomial in this compact representation. We note that it is not even known whether 2DHMSPP is in NP. We use a compact representation for the packings and show that our algorithm has time complexity $O(K^3)$ plus the time needed to compute a solution to the linear program.

Self-Improving Voronoi Construction for a Hidden Mixture of Product Distributions

Siu Wing Cheng
HKUST, Hong Kong, China
scheng@cse.ust.hk

Man Ting Wong
HKUST, Hong Kong, China
mtwongaf@connect.ust.hk

July 2021

Abstract

We proposed a self-improving algorithm for computing Voronoi diagrams under a given convex distance function with constant description complexity. The n input points are drawn from a hidden mixture of product distributions; we are only given an upper bound $m = o(\sqrt{n})$ on the number of distributions in the mixture, and the property that for each distribution, an input instance is drawn from it with a probability of $\Omega(1/n)$. For any $\epsilon \in (0, 1)$, after spending $O(mn \log^{O(1)}(mn) + m^{\epsilon} n^{1+\epsilon} \log(mn))$ time in a training phase, our algorithm achieves an $O(\frac{1}{\epsilon} n \log m + \frac{1}{\epsilon} n 2^{O(\log^* n)} + \frac{1}{\epsilon} H)$ expected running time with probability at least $1 - O(1/n)$, where H is the entropy of the distribution of the Voronoi diagram output. The expectation is taken over the input distribution and the randomized decision of the algorithm. For the Euclidean metric, the expected running time improves to $O(\frac{1}{\epsilon} n \log m + \frac{1}{\epsilon} H)$.

Rectangular Partitions of a Rectilinear Polygon

Hwi Kim* Jaegun Lee* Hee-Kap Ahn*

We investigate the problem of partitioning an axis-aligned rectilinear polygon P into axis-aligned rectangles using disjoint open line segments drawn inside P under two optimality criteria. In the *minimum ink* partition, the total length of the line segments drawn inside P is minimized. In the *thick* partition, the minimum side length over all resulting rectangles is maximized. If there are more than one such partition, we break ties among them by favoring one with the fewest number of rectangles.

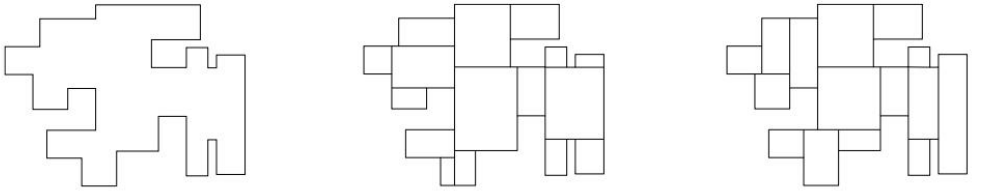

Figure 1: Left: rectilinear polygon P. Middle: a minimum ink partition of P. It minimizes the total length of the line segments used to partition P into rectangles. Right: a thick partition of P. It maximizes the minimum side length of the resulting rectangles.

Previous works. Lingas et al. proposed the minimum ink partition problem. They gave a sketch of an $O(n^4)$-time algorithm using dynamic programming for rectilinear polygons with n vertices and no holes. For a rectilinear polygon containing holes, they showed that the corresponding decision problem for the minimum ink partition problem is strongly NP-complete.

The thick partition problem was studied by O'Rourke and Tewari. They conjectured that the problem is NP-hard if holes are allowed, and claimed an $O(n^{42})$-time algorithm for rectilinear polygons with n vertices without holes. When the line segments of a partition are restricted to be incident to polygon vertices, they gave an $O(n^5)$-time and $O(n^4)$-space algorithm, by using arguments similar to the one by Lingas et al.

Our contribution. We present an $O(n^3)$-time algorithm using $O(n^2)$ space that returns a minimum ink partition of P. Our algorithm is based on the work by Lingas et al. By analyzing the subproblems carefully and exploiting certain geometric and combinatorial coherence among them, we could improve upon their algorithm. We present an $O(n^3 \log^2 n)$-time algorithm using $O(n^3)$ space that returns a thick partition using line segments incident to vertices of P. We also show that if the input rectilinear polygon has holes, the corresponding decision problem for the thick partition problem using line segments incident to vertices of the polygon is NP-complete.

*Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, South Korea. Email address: {hwikim, jagunlee, heekap}@postech.ac.kr.

Covering Convex Polygons by Two Congruent Disks*

Jongmin Choi† Dahye Jeong† Hee-Kap Ahn‡

The problem of covering a region R by a predefined shape Q (such as a disk, a square, a rectangle, a convex polygon, etc.) in the plane is to find k homothets[1] of Q with the same homothety ratio such that their union contains R and the homothety ratio is minimized. The homothets in the covering are allowed to overlap, as long as their union contains the region. This is a fundamental optimization problem arising in analyzing and recognizing shapes, and it has real-world applications, including computer vision and data mining.

In this paper we consider the covering problem for a convex polygon in which we find two congruent disks of minimum radius whose union contains the convex polygon. Our problem can be considered as the (geodesic) two-center problem for a convex polygon. We present an $O(n \log n)$-time deterministic algorithm for the two-center problem for a convex polygon P with n vertices. That is, given a convex polygon with n vertices, we can find in $O(n \log n)$ time two congruent disks of minimum radius whose union covers the polygon. This improves upon the $O(n \log^3 n \log \log n)$ time bound of Kim and Shin.

*This research was supported by the Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2017-0-00905, Software Star Lab (Optimal Data Structure and Algorithmic Applications in Dynamic Geometric Environment)) and (No. 2019-0-01906, Artificial Intelligence Graduate School Program(POSTECH)).

†Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, South Korea. Email address: {icothos,dahyejeong}@postech.ac.kr

‡Department of Computer Science and Engineering, Graduate School of Artificial Intelligence, Pohang University of Science and Technology, Pohang, South Korea. Email address: (heekap@postech.ac.kr

[1]For a shape Q in the plane, a (positive) homothet of Q is a set of the form $\lambda Q + v := \{\lambda q + v \mid q \in Q\}$, where $\lambda > 0$ is the homothety ratio, and $v \in \mathbb{R}^2$ is a translation vector.

Minimum-Link Shortest Paths for Polygons amidst Rectilinear Obstacles

Mincheol Kim* Hee-Kap Ahn†

We consider the problem for shortest paths connecting two axis-aligned rectilinear simple polygons in the domain consisting of axis-aligned rectilinear obstacles in the plane. The bounding boxes, one defined for each polygon and one defined for each obstacle, are disjoint. The problem is that, given two axis-aligned rectilinear simple polygons S and T in a rectilinear domain in the plane such that S, T, and the obstacles in the domain are pairwise box-disjoint, find a minimum-link rectilinear shortest path from S to T.

We present an algorithm that computes a minimum-link rectilinear shortest path connecting the two polygons and avoiding the obstacles in $O((N+n)\log(N+n))$ time using $O(N+n)$ space, where n is the number of vertices in the domain and N is the total number of vertices of the two polygons.

We first consider a simpler problem for an axis-aligned line segment S and a point t contained in the domain consisting of axis-aligned rectangular obstacles. We partition the domain into at most eight regions using eight xy-monotone paths from S. We observe that every shortest path from S to a point in a region is x-, y-, or xy-monotone. Moreover, we define a set of $O(n)$ baselines for each region, and show that there is a minimum-link shortest path from S to t consisting of segments contained in the baselines. Based on these observations, our algorithm applies a plane sweep technique with a sweep line moving from S to t and computes the minimum numbers of links at the intersections of the baselines and the sweep line efficiently. After the sweep line reaches t, our algorithm reports a minimum-link shortest path that can be obtained from a reverse traversal from t using the number of links stored in baselines. During the sweep, our algorithm maintains a data structure storing baselines (and their minimum numbers of links) and updates the structure for the segments (events) on the boundary of the region.

Then we extend our algorithm to handle a line segment T (not a point t) and box-disjoint rectilinear obstacles (not necessarily rectangles). We observe that every shortest path contained in a region from S to any point of T is x-, y-, or xy-monotone, so our algorithm partitions the domain into at most eight regions again. Then T intersects at most five regions. Our algorithm computes a minimum-link shortest path from S to T' for the portion T' of T contained in each region, and then returns the minimum-link shortest path among the paths. Then, we consider that the input objects are rectilinear simple polygons S and T with N vertices. We add $O(N)$ additional baselines and $O(N)$ events induced by S and T during the plane sweep algorithm. Then the number of events becomes $O(N+n)$ and the time to handle each event takes $O(\log(N+n))$, so we obtain $O((N+n)\log(N+n))$ time and $O(N+n)$ space.

*Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Korea. rucatia@postech.ac.kr

†Graduate School of Artificial Intelligence, Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Korea. heekap@postech.ac.kr

Duality Theorem in Kontsevich's Pebble Game and Its Generalization

Ching Hsiang Lin, Wing-Kai Hon

Department of Computer Science,
National Tsing Hua University, Hsinchu, Taiwan.

Abstract

The paper studies a solitaire game generalized from the pebble game of Kontsevich (1981), where instead of playing the game on a two-dimensional chessboard, we extend the game to be played on any directed graph. We show an interesting duality theorem, such that if two initial configurations c and c' of the game are a *dual*, then the game with configuration c is solvable if and only if the game with configuration c' is solvable.

Keywords: pebble game, solitaire, duality theorem, conversion algorithm

Profitable Prediction Market Making (Abstract)

Po-An Chen*, Yiling Chen**, Chi-Jen Lu***, and Chuang-Chieh Lin†

An *automated prediction market maker* is an institution that adaptively sets prices for each security and is willing to accept trades at these prices all the time. A common goal of a prediction market making is to upper bound the worst-case loss since it makes sense for a market maker not to lose an arbitrarily large amount of money. Nevertheless, a general market maker provides liquidity to traders and seeks to "profit" from the the buy and sell prices of an asset. Profit making for a market maker has also been considered in some market making works, none of which exploits a connection to *no-regret online learning* (like those that bound the worst-case loss [1]) that is our main technical approach to achieve not only an upper bounded worst-case loss but also a lower bounded profit for some classes of instances.

In [1], the correspondence between an online linear optimization/learning problem using Follow the Regularized Leader and a market making problem using a duality-based cost function market maker has been extensively explored. There, the goal of regret minimization for learners corresponds to the goal of minimizing the worst-case loss for market makers, which is a common objective for prediction market design. With the insights from online learning about designing no-regret algorithms under a "predictable" or "more regular" loss environment (e.g., specifically, low variation or low deviation) [3, 4], which corresponds to some "patterns" of trade sequences in market making, *we aim to achieve market making that furthermore guarantees profits, i.e., negative regrets, under appropriate patterns of trade sequences, which may require conditions other than those suggested by just low deviation or variation.*

The main methods of no-regret online learning (optimistic agile-update online mirror descents [4] and in particular, double agile-update online mirror descents [3] as a special case) are first reviewed, and then come our results that include designing optimistic lazy-update online mirror descents and analyzing it in the "Be-The-Regularized-Leader" framework because with linear losses, optimistic lazy-update online mirror descents are equivalent to Be the Regularized Leader with the supposedly unknown current loss vector being estimated due to the "predictability".

Following the framework of regret analysis with the help of Be the Regularized Leader, we focus on analyzing Be the Regularized Leader with the known current loss vector (i) when *in each time step, a leader is "strong" compared with the other non-minimizers in terms of its much little current cumulative loss.* The regret will be negative in this case, and the more frequent changes of leaders the more negative of the regret. Finally, if the immediately previous loss vector is a good estimator of the current loss vector, the regret can stay negative if the estimation error is small. On the other hand, (ii) we are using the modified double-update multiplicative update algorithm of [2] for our purpose of *catching the changes of "dominant experts" quickly enough* to beat a fixed best expert in hindsight in cumulative losses thereby to obtain negative regrets.

References

1. J. Abernethy, Y. Chen, and J. W. Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(1), 2012.
2. C.-K. Chiang, C.-J. Lee, and C.-J. Lu. Online learning in a gradually evolving world, 2011. Unpublished manuscript.
3. C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, and S. Zhu. Online optimization with gradual variations. In *Proc. Conference on Learning Theory*, 2012.
4. A. Rakhlin and K. Sridharan. Online learning with predictable sequences. In *Proc. Conference on Learning Theory*, 2013.

* Institute of Information Management, National Yang Ming Chiao Tung University. poanchen@nctu.edu.tw
** Department of Computer Science, Harvard University. yiling@seas.harvard.edu
*** Institute of Information Science, Academia Sinica. cjlu@iis.sinica.edu.tw
† Department of Computer Science and Information Engineering, Tamkang University. 158778@mail.tku.edu.tw

Correcting matrix products over the ring of integers

Yu-Lun Wu, Hung-Lung Wang

Department of Computer Science and Information Engineering
National Taiwan Normal University, Taipei, Taiwan
(60947037s,hlwang)@gapps.ntnu.edu.tw

Abstract. Matrix operations are important to many scientific fields. However, erroneous results of matrix operations may occur due to several reasons, like software bugs and bit-flips in memory or communication error. In this work, we focus on matrix multiplications and are concerned with how errors in the product can be corrected. Formally, given three n-by-n matrices A, B, and C, assuming that C has at most k elements differ from $A \times B$, one is asked to find the correct result of $A \times B$.

In [Leszek Gąsieniec, Christos Levcopoulos, Andrzej Lingas, Rasmus Pagh, Takeshi Tokuyama. Efficiently correcting matrix products. Algorithmica 79 (2017)], an $\tilde{O}(kn^2)$-time algorithm was proposed for correcting matrix products over any ring, where the notation $\tilde{O}$ suppresses polylogarithmic terms in n and k. Motivated by this result, we are interested in employing the properties of integers to derive a more efficient deterministic algorithm. Similar work has been conducted in [Ivan Korec, Jiří Widermann. Deterministic verification of integer matrix multiplication in quadratic time. SOFSEM 2014], under the Blum-Shub-Smale (BSS) model of computation. However, the ability of accomplishing the operations on any real numbers in constant time is treated unrealistic.

The main issue of the algorithm designed under the BSS model is that exponentially large values, to the maximum entry in the given matrices, have to be manipulated in the execution. We improve this bound to be polynominal by taking advantage of the generator of a cyclic group. This results in an $O(k^{0.75}n^2)$-time deterministic algorithm.

Query Learning of Symbolic Weighted Finite Automata

Kaito Suzuki*, Diptarama Hendrian, Ryo Yoshinaka, and Ayumi Shinohara

Graduate School of Information Sciences, Tohoku University, Japan

We propose a query learning algorithm for symbolic weighted finite automata (SWFAs), along with an algorithm for minimizing and checking equivalence of them. SWFAs can be seen as the unification of two notable extensions of classical finite automata, symbolic finite automata (SFAs) and weighted finite automata (WFAs). SWFAs represent *formal power series* on a semiring, which are functions from a set of strings over a possibly infinite alphabet to a semiring. Transition edges of an SWFA are labeled by finitely describable functions, called *guard functions*, in a fixed class $\mathcal{G}$ from the alphabet to a semiring. This design makes SWFAs suitable to deal with large or infinite alphabet efficiently. Although SWFAs and their extensions are considered in some recent studies due to their generality, the learnability of SWFAs has not been studied yet.

Query learning is a learning framework which assumes the existence of an oracle, called a *minimally adequate teacher (MAT)*. The MAT answers *membership queries (MQs)* and *equivalence queries (EQs)* from the learner. For an MQ, the MAT receives an input and returns the corresponding output value. For an EQ, the MAT receives a hypothesis and returns "yes" if it correctly represents the target function otherwise returns a counterexample. Our query learning algorithm for SWFAs is a combination of the one for SFAs [1] and the one for WFAs [2]. We prove that SWFAs over a field are exactly learnable under the MAT model if the class of guard functions $\mathcal{G}$ admits a MAT learner Λ and is closed under linear combination. Our algorithm creates instances of Λ for transition edges of an SWFA and let each instance of Λ learn a guard function by pretending to be a MAT for them. That is, our algorithm answers queries from instances of Λ, which we call $\mathcal{G}$-EQs and $\mathcal{G}$-MQs. The query complexity of our algorithm is $O(n^3\mathcal{E})$ for EQs and $O(n^4\mathcal{M} + n^4\mathcal{E}(n + \log m))$ for MQs, where n is the minimal number of states to represent the target power series, m is the length of the longest counterexample against EQs, and $\mathcal{E}$ and $\mathcal{M}$ are the number of $\mathcal{G}$-EQs and $\mathcal{G}$-MQs required to learn a guard function on each transition edge by Λ, respectively. The query complexity of our algorithm does not depend on the size of the alphabet unlike that of WFAs. Thus, our algorithm is well suited for learning power series over a large or infinite alphabet.

We also propose an algorithm to minimize SWFAs when $\mathcal{G}$ is closed under linear combination and admits equivalence checking. With this algorithm, minimization and equivalence checking of SWFAs over a field can be achieved in cubic time in the number of states of given SWFAs as with WFAs [3]. Therefore, answers to EQs on SWFAs are efficiently computable.

The full version of the above results will be available in [4]. In addition, we conducted some experiments and found that the practical performance of our algorithm is much better than the theoretical worst-case query complexity. We considered SWFAs over rational numbers and used randomly-generated SWFAs as oracles, the class of polynomial functions as $\mathcal{G}$, and simple polynomial curve fitting algorithm with queries as Λ. In this setting, we observed that the number of EQs of our algorithm is almost linear in n.

One of the interesting applications of our algorithm is extracting an SWFA from a recurrent neural network (RNN) to obtain a faster surrogate. A prior work [5] use WFAs for this purpose, but the range of applicable RNNs is limited by the size of the alphabet. SWFAs and our algorithm are suitable to extract automata from RNNs over a large or infinite alphabet, like real-valued inputs. Such an application requires additional devices for our algorithm to work well with numerical error, which is left for future work.

References

[1] G. Argyros and L. D'Antoni. The learnability of symbolic automata. In *CAV 2018*, pages 427–445, 2018.

[2] L. Bisht, N. H. Bshouty, and H. Mazzawi. On optimal learning algorithms for multiplicity automata. In *COLT 2006*, pages 184–198, 2006.

[3] M. P. Schützenberger. On the definition of a family of automata. *Inf. Control*, 4(2-3):245–270, 1961.

[4] K. Suzuki, D. Hendrian, R. Yoshinaka, and A. Shinohara. Query learning algorithm for symbolic weighted finite automata. In *ICGI 2020/21*, pages 202–216, 2021.

[5] T. Okudono, M. Waga, T. Sekiyama, and I. Hasuo. Weighted automata extraction from recurrent neural networks via regression on state spaces. In *AAAI 2020*, pages 5306–5314, 2020.

*Correspondence to: kaito_suzuki@shino.ecei.tohoku.ac.jp

Accountable Ring Signature From Isogeny Group Action*

Yao-Ching Hsieh
ychsieh@ntu.edu.tw

Mi-Ying Huang
miying.huang@usc.edu

Yu-Hsuan Huang
asd00012334.cs04@nctu.edu.tw

Background. Research efforts have been put into signatures that preserve privacy. These include *ring signatures* [RST01] and *group signatures* [CvH91]. The former allows a signer to produce a signature on behalf of a *ring*, a set of players chosen by the signer, while the latter are signed on behalf of a *group*, a set of players chosen by a prescribed *master* that is able to *open*, i.e. revealing signer identities with his *master secrete key*. Joining the good of both, if a signer gets to choose his own ring and a master non-interactively chosen by the signer could open signer identities, the so-called *accountable ring signature* [XY04] would be even more desirable.

Later on, a (pre-quantum) accountable ring signature is proposed by [BCC+15] and, as what they have pointed out, accountable ring signatures would imply group signatures. Early works on constructing group signature schemes are mostly based on bilinear groups [LLLS13]. Though such schemes are considered elegant and efficient, they use pairing-based assumptions with no chance to achieve post-quantum security. Group signatures using (post-quantum) lattice-based assumption are instantiated following [LLNW14]. On the contrast, the above-mentioned approaches do not extend to isogeny-based cryptography, which has been absent for constructing group signatures for a while.

Isogeny Group Action. The notion of *hard homogeneous space* was brought up by 1001[*Cou*06], revealing a pathway using group action without discrete logarithm for cryptographic construction. As early as [Sto12], an identification scheme based on hard homogeneous spaces has been instantiated. In the considered protocol, a set $\mathcal{E} = \{E/\mathbb{F}_p \text{ elliptic curve with prescribed order}\}/\cong_{\mathbb{F}_p}$ of curves is acted by their (identical) ideal class group $\mathsf{Cl}(\mathcal{O})$ of the $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O} = \mathsf{End}_p(E)$ for any $E \in \mathcal{E}$. The prover shows to verifier that he knows the secrete key $\mathsf{sk} \in \mathsf{Cl}(\mathcal{O})$ corresponding to the public key $\mathsf{pk} = \mathsf{sk} \cdot E_0 \in \mathcal{E}$, for some prescribed initial curve $E_0 \in \mathcal{E}$. First, the prover sends $\phi = b \cdot \mathsf{pk}$ for some randomly sampled $b \in \mathsf{Cl}(\mathcal{O})$, waiting for verifier to challenge a random $\mathsf{ch} \in \{0, 1\}$ and then respond with $r_{\mathsf{ch}} = \begin{cases} b \text{ if } \mathsf{ch} = 0, \\ b \cdot \mathsf{sk} \text{ otherwise,} \end{cases}$ for verifier to check consistency, i.e. $r_0 \cdot \mathsf{pk} \stackrel{?}{=} \phi$ or $r_1 \cdot E_0 \stackrel{?}{=} \phi$. Later its efficiency was improved in [CLM+18] for the underlying group action, but a non-unique representation for elements in the class group $\mathsf{Cl}(\mathcal{O})$ was introduced, which potentially reveals secrete keys in the identification. To address this, works based on rejection sampling 1001[*FG*19], and computing $\mathsf{Cl}(\mathcal{O})$ structure [BKV19] were proposed. The line of works based on isogeny group actions has proven to be useful to construct signatures.

A recent advancement by Beullens et al. [BKP20] gives a group-action based construction for the so-called *OR-proofs*, in which the prover tries to convince verifer that he knows $\mathsf{sk} \in \mathsf{Cl}(\mathcal{O})$ corresponds to one of the public keys $E_1, \ldots, E_n \in \mathcal{E}$, i.e. $\exists i \in [n]$ s.t. $\mathsf{sk} \cdot E_0 = E_i$. This, when non-interactivized, yields the first construction of ring signature based on isogeny group action, shedding light on privacy-preserving signatures for isogenies.

Contribution. We present our construction of isogeny-group-action-based accountable ring signatures. We first construct its interactive variant with 4 challenges. Then apply the Fiat-Shamir style transformation to non-interactivize it, except we tell the master to vote on each parallel repetition to decide who is the signer. In the end, we have obtained rather strong security: (1) the signer remains anonymous under full key exposure, and (2) a malicious signer cannot incriminate an honest non-signer even if ring members could be adaptively corrupted. Proving this security is not just a matter of applying existing results on Fiat-Shamir transform. As a part of our technical challenges, we have to extract a secrete key from the adversary incriminating honest parties under a weaker form of special-soundness, in which the desired keys could be extracted only if responses from all the 4 challenges are collected.

The interactive variant goes as follows: First, the prover samples $\Delta_i, \Delta'_i, b \in \mathsf{Cl}(\mathcal{O})$ for each i, a shuffled list of indices $\sigma_1, \ldots, \sigma_n$, and sends $\phi = (\{E_i^\alpha\}_{i\in[n]}, \{E_i^\beta\}_{i\in[n]}, \{E_{\sigma_i}^\gamma\}_{i\in[n]}, E^{\mathsf{Open}}, E^{\mathsf{Check}})$ where $E_i^\alpha = \Delta_i E_i$, $E_i^\beta = \Delta'_i E_i^\alpha$, $E_i^\gamma = bE_i^\beta$, $E^{Open} = \Delta_k \Delta'_k s E_m$, $E^{Check} = \Delta_k \Delta'_k bs E_m = bE^{Open}$, $(E_m, s_m) \in \mathcal{E} \times \mathsf{Cl}(\mathcal{O})$ is the master key pair, $E_i \in \mathcal{E}$ is the public key for the i-th player, and k is the index of signer. The prover waits for verifier's challenge $\mathsf{ch} \in \{0, 1, 2, 3\}$ and then respond with r_{ch} for the verifier to check consistency, where r_0, r_1, r_2, r_3 is $\{\Delta_i\}_{i\in[n]}$, $\{\Delta'_i\}_{i\in[n]}$, b, $b\Delta'_k\Delta_k\mathsf{sk}$ respectively. Then, the master could reveal signer's identity k by finding an index i such that $s_m E_i^\beta = E^{\mathsf{Open}}$.

*For the full version, see: https://eprint.iacr.org/2021/1368

References

[BCC+15] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit, *Short accountable ring signatures based on ddh*, European Symposium on Research in Computer Security, Springer, 2015, pp. 243–265.

[BKP20] Ward Beullens, Shuichi Katsumata, and Federico Pintore, *Calamari and Falafl: Logarithmic (linkable) ring signatures from isogenies and lattices*, International Conference on the Theory and Application of Cryptology and Information Security, Springer, 2020, pp. 464–492.

[BKV19] Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren, *CSI-fish: Efficient isogeny based signatures through class group computations*, International Conference on the Theory and Application of Cryptology and Information Security, Springer, 2019, pp. 227–247.

[CLM+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes, *CSIDH: an efficient post-quantum commutative group action*, International Conference on the Theory and Application of Cryptology and Information Security, Springer, 2018, pp. 395–427.

[Cou06] Jean Marc Couveignes, *Hard homogeneous spaces.*, IACR Cryptol. ePrint Arch. **2006** (2006), 291.

[CVH91] David Chaum and Eugène Van Heyst, *Group signatures*, Workshop on the Theory and Application of of Cryptographic Techniques, Springer, 1991, pp. 257–265.

[FG19] Luca De Feo and Steven D. Galbraith, *SeaSign: Compact isogeny signatures from class group actions*, 759–789.

[LLLS13] Fabien Laguillaumie, Adeline Langlois, Benoît Libert, and Damien Stehlé, *Lattice-based group signatures with logarithmic signature size*, International conference on the theory and application of cryptology and information security, Springer, 2013, pp. 41–61.

[LLNW14] Adeline Langlois, San Ling, Khoa Nguyen, and Huaxiong Wang, *Lattice-based group signature scheme with verifier-local revocation*, International workshop on public key cryptography, Springer, 2014, pp. 345–361.

[RST01] Ronald L Rivest, Adi Shamir, and Yael Tauman, *How to leak a secret*, International Conference on the Theory and Application of Cryptology and Information Security, Springer, 2001, pp. 552–565.

[Sto12] Anton Stolbunov, *Cryptographic schemes based on isogenies*, Ph.D. thesis, 01 2012.

[XY04] Shouhuai Xu and Moti Yung, *Accountable ring signatures: A smart card approach*, Smart Card Research and Advanced Applications VI, Springer, 2004, pp. 271–286.

Making Three Out of Two: Three-Way Online Correlated Selection*

(Abstract)

Yongho Shin[1] and Hyung-Chan An[†1]

[1]Department of Computer Science, Yonsei University, South Korea
[1]{yshin, hyung-chan.an}@yonsei.ac.kr

July 22, 2021

Two-way online correlated selection (two-way OCS) is an online algorithm that, at each timestep, takes a pair of elements from the ground set and irrevocably chooses one of the two elements, while ensuring negative correlation in the algorithm's choices. Whilst OCS was initially invented by Fahrbach, Huang, Tao, and Zadimoghaddam to solve the edge-weighted online bipartite matching problem, it is an interesting technique on its own due to its capability of introducing a powerful algorithmic tool, namely negative correlation, to online algorithms. As such, Fahrbach et al. posed two tantalizing open questions in their paper, one of which was the following: Can we obtain *n-way OCS* for $n > 2$, in which the algorithm can be given $n > 2$ elements to choose from at each timestep?

In this paper, we affirmatively answer this open question by presenting a *three-way OCS*. Our algorithm uses two-way OCS as its building block and is simple to describe; however, as it internally runs two instances of two-way OCS, one of which is fed with the output of the other, the final output probability distribution becomes highly elusive. We tackle this difficulty by approximating the output distribution of two-way OCS by a *flat*, less correlated function. Previously known probability bounds on two-way OCS are *almost* convex, and therefore this flat function can serve as a safe "surrogate" of the real distribution. This yields a new probability bound on our three-way OCS, for a special case in which we consider the algorithm's choices only for a consecutive interval of timesteps. This bound is further generalized to non-consecutive timesteps by carefully defining a series of surgical operations that enables us to consider non-consecutive timesteps simply as a union of consecutive subintervals. Our three-way OCS can also be used to improve the previous competitive ratio of 0.5086 due to Fahrbach et al. to give a new 0.5093-competitive algorithm for edge-weighted online bipartite matching.

References

[1] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 412–423. IEEE, 2020.

[2] Yongho Shin and Hyung-Chan An. Making three out of two: Three-way online correlated selection. *arXiv preprint arXiv:2107.02605*, 2021.

*This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2019R1C1C1008934).

†Corresponding author. Department of Computer Science, Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul 03722, South Korea.

Rearranging a Sequence of Points onto a Line

Taehoon Ahn* Jongmin Choi* Chaeyoon Chung* Hee-Kap Ahn†
Sang Won Bae‡ Sang Duk Yoon§

Consider a trajectory data of a moving object which is represented by a sequence of pairs, each consisting of a time stamp and the coordinates of the object at the time. We measure the quality of the trajectory with respect to a linear path by their similarity. We define a rearrangement of the trajectory data onto the path that represents the similarity, and give efficient algorithms to compute the optimal rearrangement in various settings.

Given a sequence of n weighted points $\langle p_1, p_2, \ldots, p_n \rangle$ in the plane, we consider the problem of finding a rearrangement of the points, q_i for each p_i, onto a line such that any two consecutive points q_i and q_{i+1} are at distance no more than their weight difference, and the maximum distance between p_i and q_i over all i is minimized. We present efficient algorithms that compute optimal rearrangements for three variants of the problem either under the L_1 metric or under the Euclidean metric. When the line is fully specified or partially specified by only its orientation, our algorithms take near-linear time. When we need to find a target line, onto which the input sequence can be rearranged with the optimal rearrangement cost, we present an $O(n^3 \operatorname{polylog} n)$-time algorithm.

*Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Korea. {sloth, icothos, chaeyoon17}@postech.ac.kr

†Department of Computer Science and Engineering, Graduate School of Artificial Intelligence, Pohang University of Science and Technology, Pohang, Korea. heekap@postech.ac.kr

‡Department of Computer Science, Kyonggi University, Suwon, Korea. swbae@kgu.ac.kr

§Department of Service and Design Engineering, Sungshin Women's University, Seoul, Korea. sangduk.yoon@sungshin.ac.kr

Intersecting Disks using Two Congruent Disks

Byeonguk Kang* Jongmin Choi* Hee-Kap Ahn†

The Euclidean k-center problem is a fundamental problem in the field of computational geometry. The Euclidean k-center problem is to find k smallest congruent balls such that every input point is contained in one of the k balls. The special case of the problem for $k = 2$ in the plane, also known as the planar 2-center problem. Motivated by facility location problems for mobile demand points and geometric optimization for imprecise points, we consider the 2-center problem on disks. The 2-center problem on disks is a generalization of the 2-center problem in which given a set $\mathcal{D}$ of n disks of nonnegative radii in the plane, find two smallest congruent disks C_1 and C_2 satisfying $D \cap (C_1 \cup C_2) \neq \emptyset$ for every $D \in \mathcal{D}$. We call this problem the *2-center problem on disks*. The 2-center problem on disks and its related problems are shown in Figure 1.

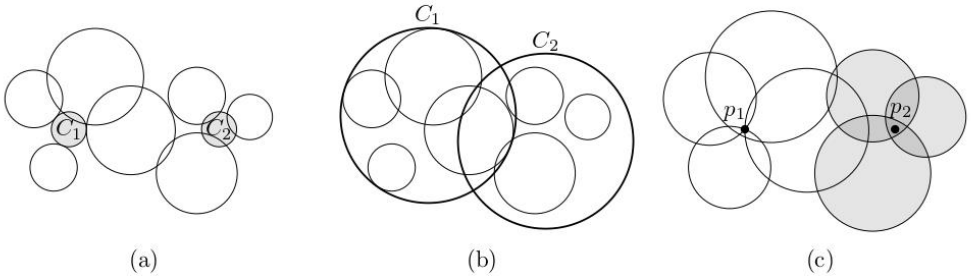

Figure 1: (a) The 2-center problem on disks in the plane: every input disk intersects C_1 or C_2. (b) The restricted 2-cover problem on disks in the plane: every disk is fully contained in C_1 or C_2. (c) The 2-piercing problem on disks in the plane: every disk is pierced by p_1 or p_2.

Previous work. Ahn et al. gave a deterministic algorithm that returns an optimal pair of congruent disks in $O(n^2 \log^4 n \log\log n)$ time, and a randomized algorithm with $O(n^2 \log^3 n)$ expected time. They showed that their algorithms also work for the *restricted 2-cover problem* and the *2-piercing problem* on disks in the plane. The restricted 2-covering problem on disks in the plane is to find two smallest congruent disks C_1 and C_2 such that every input disk is contained in C_1 or C_2. The 2-piercing problem on disks in the plane is to find two points p_1 and p_2 such that every input disk is pierced by p_1 or p_2. See Figure 1 (b, c).

Our contribution. We give a deterministic algorithm for the 2-center problem that returns an optimal pair of congruent disks in $O(n^2 \log^3 n/\log\log n)$ time. We also present a randomized algorithm with $O(n^2 \log^2 n/\log\log n)$ expected time. These results improve the previously best deterministic and randomized algorithms, making a step closer to the optimal algorithms for the problem. We show that the same algorithms also work for the restricted 2-covering problem and the 2-piercing problem on disks.

*Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Korea. Email address: {kbu417, icothos}@postech.ac.kr.

†Department of Computer Science and Engineering, Graduate School of Artificial Intelligence, Pohang University of Science and Technology, Pohang, Korea. Email address: heekap@postech.ac.kr

Greedy is Optimal for Online Restricted Assignment and Smart Grid Scheduling for Unit Size Jobs

Fu-Hong Liu*† Hsiang-Hsuan Liu‡ Prudence W.H. Wong§

In this paper, we study online scheduling of unit-sized jobs in two related problems, namely, restricted assignment problem and smart grid problem. The input to the two problems are in close analogy but the objective functions are different. We show that the greedy algorithm is an optimal online algorithm for both problems by showing that both objective functions have led to the same property of the greedy algorithm. The property is crucial for the optimality of the greedy algorithm.

Smart grid scheduling. We consider online scheduling of unit-sized requests with the following input. A consumer sends in a power request j with unit power requirement, unit duration of service, and feasible timeslots $F(j)$ that j can be served. The operator of the smart grid selects a timeslot from $F(j)$ for each request j. The *load* of the grid at each timeslot t is the number of requests allocated to t. The *energy cost* is modeled by a strictly increasing convex function $f(t)$ on load(t). The objective is to minimize the total energy cost over time, i.e., minimize $\sum_t f(\text{load}(t))$.

Restricted assignment problem. Considering a set of jobs and a set of machines, each job is associated with a unit processing time and a subset of machines that the job can be scheduled on. We consider online scheduling of the jobs. The objective is to minimize the maximum number of jobs assigned to any machine while satisfying the assignment restriction constraints.

Theorem 1 *When the input to the grid scheduling problem and the restricted assignment problem is a set of unit-sized jobs, the greedy algorithm is an optimal online algorithm having the best possible competitive ratio.*

Typically, an online algorithm is proved to be an optimal online algorithm through bounding its competitive ratio and showing a lower bound with matching competitive ratio. However, our analysis does not take this approach. Instead, we prove the optimality without giving the exact bounds on the competitive ratio.

In this paper, we develop an *adversary design* technique. We design some adversarial input which is bad for the greedy strategy. Also by this technique, we design an adversary for any online algorithm that performs bad enough to be worse than the performance of a greedy strategy. By these adversarial inputs, we can show that greedy is the best online strategy for this problem.

More specifically, given any online algorithm and a load configuration, we show the existence of two job instances J_1 and J_2 such that (i) J_1 and J_2 admit the same optimal offline schedule represented by the given load configuration; (ii) the cost of the schedule produced by the given online algorithm on J_1 is at least the cost of the schedule produced by the greedy algorithm on J_2. This means that when we consider any job instance for the greedy algorithm, there is always another job instance such that the ratio versus the (same) optimal offline schedule of the greedy algorithm is not larger than any online algorithm. Hence, we can show that the competitive ratio of the greedy algorithm is the smallest possible. The existence of the two job sets relies on a property about the relative costs of two comparable schedules. We show that this property holds for both objective functions for the grid scheduling problem and the restricted assignment problem, hence, the optimality holds for both problems.

*Department of Industrial Engineering and Engineering Management, National Tsing Hua University, Taiwan
†E-mail: homer.fhliu@gmail.com
‡Department of Information and Computing Sciences, Utrecht University, The Netherlands
§Department of Computer Science, University of Liverpool, UK

Online Independent Set with Amortized Late Accept/Reject

Jonathan Toole-Charignon*† Hsiang-Hsuan Liu‡

In an *online* setting, a problem instance is presented piece-by-piece as a sequence of inputs rather than as a single complete input. In the case of a graph problem, this is typically done by revealing the input graph vertex-by-vertex, with edges being revealed once both of their incident vertices have been presented. An online algorithm must make irrevocable decisions on each input without knowledge of future inputs. The performance of an online algorithm is most commonly measured by the *competitive ratio*, where an online algorithm has competitive ratio c if, for all instances of a problem, the performance of the online algorithm is no more than c times worse than the offline optimal solution.

The Online Independent Set problem consists of constructing an independent set that is as large as possible, where the algorithm constructing the solution must choose whether to include a vertex in its solution upon that vertex's reveal.

Within this standard online model, any deterministic algorithm for Online Independent Set has a trivial competitive ratio of $n-1$, where n is the number of vertices in the input graph. Because of this, research on Online Independent Set (and other graph problems with similar results) has focused on alternative models that relax one or more of the restrictions involved in the standard online model. In particular, Boyar et al. introduced the notions of Late Accept (LA) and Late Reject (LR), where the online algorithm is respectively allowed to include a previously unselected vertex into its solution, or to reject a previously selected vertex. These two notions are also combined in the Late Accept/Reject (LAR) model, with the caveat that late-rejections are irrevocable.

We extend the models introduced by Boyar et al. through Amortized Late Accept/Reject (ALAR), where the caveat above is removed, and we establish a relationship between the number of amortized late rejections per vertex and the competitive ratio. This directly addresses Boyar et al.'s open question on the "trade-offs between the number of late operations employed and the quality of the solution".

Our main result is an upper bound for a special graph class on the required number of amortized late rejections per vertex given a target competitive ratio.

Theorem. *In ALAR, given a target competitive ratio of $c > 1$, there exists a deterministic online algorithm for* Online Independent Set *on bipartite graphs that is c-competitive while allowing r amortized late rejections per vertex, where r is upper-bounded by: $r \leq \frac{c}{(c-1)(c+1)}$*

This result is obtained by reducing any possible adversary to an adversary on a complete bipartite graph that forces the online algorithm to switch between the two partitions of the graph in order to maintain the target competitive ratio. This result also provides a corresponding bound on the competitive ratio given a restriction on the allowed number of amortized late rejections per vertex.

Corollary. *In ALAR, given a value $r > 0$, there exists a deterministic online algorithm on bipartite graphs that is c-competitive while allowing at most r amortized late rejections per vertex, where $c > 1$ is upper-bounded by: $c \leq \frac{1}{2r}(\sqrt{4r^2+1}+1)$*

We can use this corollary for a direct numerical comparison, by noting that the upper bound on the competitive ratio for ALAR while allowing at most one amortized late rejection per vertex is $\varphi \approx 1.618$, whereas the equivalent upper bound for LAR is $\frac{3\sqrt{3}}{2} \approx 2.598$.

*Department of Information and Computing Sciences, Utrecht University, The Netherlands
†Email: j.c.f.toole-charignon@uu.nl
‡Department of Information and Computing Sciences, Utrecht University, The Netherlands

Learning-Augmented Algorithms for Online TSP

Hsiao-Yu Hu, Ya-Chun Liang, Jian-Xi Shao and Chung-Shou Liao

National Tsing Hua University

In recent years, a rapidly developing line of research which incorporated machine-learned predictions into online algorithms has been widely discussed to improve the performance [4, 5, 6]. The goal is to design online algorithms that take advantage of learning models while guaranteeing theoretical loss bounds possibly caused by poor predictions.

In this study, we consider one of the classical online combinatorial optimization problems, the online traveling salesman problem (OLTSP) [1, 2, 3], in which a salesman needs to serve a sequence of requests that arrive in an online fashion and return to the origin as quickly as possible. In other words, the input of an online instance of the problem is a set of pairs in which their locations are to be visited in a metric space and their arrival time is released online. The objective is to minimize the total completion time.

We investigate two types of learning-augmented online algorithms for the OLTSP: offline and online predictions. The former ones make a whole prediction in a priori knowledge before an online instance is coming, similar to what the almighty adversary does. The latter ones allow an input instance of the online requests to be forecast one by one only. We explore both the two types of online algorithms with predictions, and present such online algorithms with theoretical guarantees for the OLTSP in a metric space and on a real line.

Keywords: Learning-augmented, online algorithm, competitive analysis

References:

[1] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Algorithms for the On-Line Travelling Salesman. Algorithmica 29, pages 560–581, 2001.

[2] Antje Bjelde, Yann Disser, Jan Hackfeld, Christoph Hansknecht, Maarten Lipmann, Julie Meißner, Kevin Schewior, Miriam Schlöter, and Leen Stougie. Tight Bounds for Online TSP on the Line. In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'17), pages 994–1005, 2017.

[3] Michiel Blom, Sven O. Krumke, Willem E. de Paepe, and Leen Stougie. The Online TSP Against Fair Adversaries. INFORMS Journal on Computing 13(2), pages 138-148, 2001.

[4] Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online Scheduling via Learned Weights. In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'20), pages 1859–1877, 2020.

[5] Thodoris Lykouris and Sergei Vassilvitskii. Competitive Caching with Machine Learned Advice. In International Conference on Machine Learning (ICML), pages 3302–3311, 2018.

[6] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving Online Algorithms via ML Predictions. In Advances in Neural Information Processing Systems (NeurIPS), pages 9684–9693, 2018.

Single-Pass Streaming Algorithms to Partition Graphs into Few Forests

Cheng-Hung Chiang
timmychiang@iis.sinica.edu.tw

Meng-Tsung Tsai
mttsai@iis.sinica.edu.tw

We devise a single-pass $O(n)$-space deterministic streaming algorithm to partition any n-node undirected simple graph G into $O(\alpha \log n)$ forests where α is the minimum number of forests which G can be partitioned into. We then apply this result to obtain single-pass streaming algorithms for other graph problems, including low outdegree orientation, partitioning graphs into few planar subgraphs, and finding small dominating sets.

Noname manuscript No.
(will be inserted by the editor)

On the Computational Power of Phosphate Transfer Reaction Networks

Chun-Hsiang Chan · Cheng-Yu Shih · Ho-Lin Chen

Received: date / Accepted: date

Abstract Phosphate transfer reactions [11] involve the transfer of a phosphate group from a donor molecule to an accepter, which is ubiquitous in biochemistry. Besides natural systems, some synthetic molecular systems such as seesaw gates are also equivalent to (subsets of) phosphate transfer reaction networks. In this paper, we study the computational power of phosphate transfer reaction networks (PTRNs). PTRNs are chemical reaction networks (CRNs) with only phosphate transfer reactions. Previously, it is known [4] that a function can be deterministically computed by a CRN if and only if it is semilinear. However, the computational power of programmable phosphate transfer networks is unknown.

In this paper, we present a formal model to describe PTRNs and study the computational power of these networks. We prove that when each molecule can only carry one phosphate group, the output must be the total initial count in a subset S_1 minus the total initial count of another subset S_2. On the other hand, when every molecule can carry up to three phosphate groups, or two phosphate groups with different functions, PTRNs can "simulate" arbitrary CRNs. Finally, when each molecule can carry up to two functionally iden-

Research supported by MOST (Taiwan) grant number 107-2221-E-002-031-MY3 and 104-2221-E-002-045-MY3. A preliminary version of this paper [3] (2-page abstract) has appeared in the proceedings of FNANO 2016.

Chun-Hsiang Chan
University of Michigan. The work was done when he was in the Department of Electrical Engineering, National Taiwan University
E-mail: kennyhchan@gmail.com

Cheng-Yu Shih
Department of Electrical Engineering, National Taiwan University
The first two authors contribute equally to this paper and are listed in alphabetical order.
E-mail: r07921036@ntu.edu.tw

Ho-Lin Chen
Corresponding author. Department of Electrical Engineering, National Taiwan University
E-mail: holinchen@ntu.edu.tw

Node Failure Survivability: An Efficient Logical Topology Mapping Algorithm for IP-over-WDM Optical Networks

Dun-Wei Cheng[1], Jo-Yi Chang[2], Chen-Yen Lin[3], Limei Lin[4], Yanze Huang[5], Krishnaiyan Thulasiraman[6], and Sun-Yuan Hsieh[7](Senior Member, IEEE)

Keywords: *Network Survivability, Survivable Mapping Design, Disjoint Paths, Node Failure.*

Abstract

A survivable mapping problem (SMP) in IP-over-WDM network with logical graph and physical graph is the problem finding a mapping in physical layer so that any failure in physical topology does not break the logical topology's connection. To determine whether a survivable mapping against failure exists is an NP-complete problem, and therefore many heuristic algorithms have been proposed in the literature. In this paper, a heuristic mapping design strategy is proposed to enable the lightpaths to more efficiently endure node failure. Experimental results demonstrate that the proposed algorithm can provide notably survivable mapping in IP-over-WDM networks.

1. Dun-Wei Cheng was with Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN E-mail: dunwei.ncku@gmail.com.
2. Jo-Yi Chang was with Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN E-mail: insjoy20@gmail.com.
3. Chen-Yen Lin was with Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN E-mail: P76021077@mail.ncku.edu.tw.
4. Limei Lin is with the College of Mathematics and Informatics and the Key Laboratory of Network Security and Cryptology, Fujian Normal University, Fuzhou, Fujian, 350117 P.R. China. E-mail: putianlinlimei@163.com.
5. Yanze Huang is with the School of Mathematics and Physics, Fujian Provincial Key Laboratory of Big Data Mining and Applications, Fujian University of Technology, Fuzhou, Fujian, 350118 P.R. China. E-mail: yzhuang@fjut.edu.cn.
6. Krishnaiyan Thulasiraman was with Department of Computer Science, Computer Science University of Oklahoma, USA E-mail: thulasi@ou.edu.
7. Sun-Yuan Hsieh was with Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN E-mail: hsiehsy@mail.ncku.edu.tw.

The Construction of Multiple Independent Spanning Trees on Generalized Recursive Circulant Graphs

DUN-WEI CHENG1[1] , KAI-HSUN YAO[2],
AND SUN-YUAN HSIEH[3](Senior Member, IEEE)

Keywords: *Independent Spanning Trees, Generalized Recursive Circulant Graphs, Interconnection Networks.*

Abstract

The generalized recursive circulant networking can be widely used in the design and implementation of interconnection networks. It consists of multi-processors, each is connected through bidirectional, point-to-point communication channels to different neighbors.

A set of spanning trees are said to be independent if they are rooted at the same vertex and for each of the remaining vertex there exists internally disjoint paths connected to the root. Independent spanning trees(ISTs) are largely used to improve fault-tolerant ability and secure information distribution.

In this work, we proposed a novel method to construct independent spanning trees on generalized recursive circulant graphs(GRC graphs). On each vertex, we first apply the shortest path routing concept to collect the needed movements which indicate a connection from a node to another one. Then finding the vertex's parents through the strategy depends on the IST we select. Eventually, the strategy leads to forming multiple internally disjoint paths from the vertex to the root.

The proposed method can construct maximal ISTs on GRC graphs in O(Nh) time complexity, where N, h denotes cardinality and dimension of the graphs. Also it loosened the restricted conditions in previous research and extended the result to a more general vertex setting by designing the specific algorithm to deal with the constraint issue.

Searching independent spanning trees is a challenging problem. In this paper, we proposed a method to find the maximal independent spanning trees on generalized recursive circulant graphs. By finding the independent spanning tees, we can use it in an interconnection network to enhance its fault-tolerant property or achieve reliable broadcasting and design secure distributed protocols.

[1] Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN. (e-mail:dunwei.ncku@gmail.com)
[2] Department of Computer Science and Information Engineering, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN. (e-mail:p76071226@gs.ncku.edu.tw)
[3] Department of Computer Science and Information Engineering, Institute of Medical Informatics, Institute of Manufacturing Information and System, Center for Innovative FinTech Business Models and International Center for the Scientific Development of Shrimp Aquaculture, National Cheng Kung University, No. 1, University Road, Tainan 701, TAIWAN. (e-mail: hsiehsy@mail.ncku.edu.tw)

Hardness Results on Generalized Puyopuyo*

Hiroshi Eto[1], Hironori Kiya[2], and Hirotaka Ono[3]

[1] Department of Economic Engineering, Kyushu University, 744 Motooka Nishi-ku, Fukuoka 819-0395, Japan h-eto@econ.kyushu-u.ac.jp
[2] Institutes of Innovation for Future Society, Nagoya University, Furocho, Chikusa-ku, Nagoya, 464-8601, Japan kiya.hironori@gmail.com
[3] Department of Mathematical Informatics, Nagoya University, Furocho, Chikusa-ku, Nagoya, 464-8601, Japan ono@nagoya-u.jp

Puyopuyo is a Tetris-like falling-block video game, which has been played in Japan more than 30 years. The rule of Puyopuyo is as follows: The board forms two-dimensional 12×6 cells with four directions, *top*, *bottom*, *left*, and *right*. A *puyo* is a colored or colorless unit whose shape is a unit circle. Here, "unit" implies that one puyo can occupy one cell. Puyos are located at cells monotonically from the bottom to the top of the board, that is, there is no empty cell vertically between two cells occupied by puyos. The number of colors is typically 3, 4 or 5. Two puyos with an identical color are connected if they are in either vertically or horizontally adjacent cells. If more than three puyos are connected, they form a cluster, and the puyos of a cluster are cleared (i.e., disappear). If there is a puyo on a cluster and the cluster is cleared, the puyo drops down according to gravity, which may change the configuration of puyos. If the new configuration contains a cluster, it again disappears. That is, clearing a cluster may cause clearing another cluster, and this phenomenon is called a *clearing chain* or *chain* simply, and the length of chains is defined as the number of clearing. Other than colored puyos, there is a colorless puyo called a *nuisance* puyo, or *N-puyo*, for short. Since N-puyos are colorless, they do not form a cluster, but an N-puyo is cleared if a neighboring colored puyo is cleared. These are basic properties of puyos. In a normal play of Puyopuyo, pairs of puyos appear at the top of the board in a one-by-one manner, and they fall according to gravity. A player can move such a falling pair of puyos horizontally, which controls its dropping point. Also a falling pair of puyos can be rotated while falling and it is placed either vertically or horizontally. The score of a play depends on the number of cleared puyos or the number of chains.

In this paper, we investigate the computational complexity of a generalized variant of Puyopuyo. The variant generalizes the size of the board and the number of colors but falling pairs of puyos are given in the offline manner. We focus on two problems of the generalized Puyopuyo; board clearing and maximizing chains. Both problems are already known to be NP-complete. More precisely, the former is NP-complete even for puyos of 2 colors with N-puyo setting [3], and the latter is NP-complete even for puyos of 4 colors with N-puyo setting [2]. The latter result is mentioned to be improved to the setting of puyos of 3 colors with N-puyo and the setting of puyos of 5 colors without N-puyo [1], though the detail is not published. In this paper, we strengthen these results from several aspects. Our results are as follows: (1) The chain maximization is NP-complete even for the setting of puyos of 4 colors without N-puyo. (2) The chain maximization cannot be approximated within any polynomial factor in polynomial time, unless P=NP. (3) The board clearing is NP-complete even for the setting of puyos of 4 colors without N-puyo.

References

1. Kiba, Y., Muneshige, N., Uejima, A.: Irosu-to-ojama-puyo-wo-seigenshita-ippannka-puyopuyo-no-rensasuhanteimondaino-np-kanzensei (in japanese). In: Abstracts of The Operations Research Society of Japan (Fall 2011). pp. 370–371. The Operations Research Society of Japan (sep 2011)
2. Matsukane, T., Takenaga, Y.: NP-completeness of maximum chain problem on puyopuyo. The IEICE transactions on information and systems D (Japanese edition) **J89-D**(3), 405–413 (2006)
3. Muta, H.: PUYOPUYO is NP-complete. IEICE Technical Report pp. COMP2005–14 (2005)

* Supported by KAKENHI.

Largest similar copies of convex polygons in polygonal domains

Taekang Eom * Seungjun Lee* Hee-Kap Ahn†

In this paper, we aim to find a largest similar copy of a given convex polygon P with k vertices that can be inscribed in a polygonal domain Q consisting of n points and line segments.

The earliest result was perhaps the SoCG'89 paper by Chew and Kedem. They considered the problem and gave an incremental technique for handling all the combinatorial changes to the edge Delaunay triangulation of the polygonal domain Q (shortly eDT) under the distance function induced by the input polygon P while P is rotating. They gave an upper bound $O(k^4 n\lambda_4(kn))$ on the combinatorial changes, that is, the number of *critical orientations*, together with a deterministic $O(k^4 n\lambda_4(kn)\log n)$-time algorithm, where $\lambda_s(n)$ the length of the longest Davenport–Schinzel sequence of order s including n distinct symbols. A few years later, the bound was improved to $O(k^4 n\lambda_3(n))$ by them, and thus the running time of the algorithm became $O(k^4 n\lambda_3(n)\log n)$ [CGTA93].

Sharir and Toledo [CGTA94] studied this problem and applied the motion-planning algorithm to solve this problem. They gave an algorithm with running time $O(k^2 n\lambda_4(kn)\log^3(kn)\log\log(kn))$ that uses the parametric search technique of Megiddo. $O(k^4 n\lambda_3(n)\log n)$-time algorithm and $O(k^2 n\lambda_4(kn)\log^3(kn)\log\log(kn))$-time algorithm are comparable to each other.

We improve the upper bound on the combinatorial changes considered during the rotation and the algorithm to compute a largest similar copy. We present an upper bound $O(k^2 n^2\lambda_4(k))$ on the combinatorial changes, and this directly improves the time bound for the algorithm to $O(k^2 n^2\lambda_4(k)\log n)$. This improves the previously best known results in more than 25 years.

Our strategy to improve the upper bound on the combinatorial changes follows the approach of Chew and Kedem, which consists of two parts. In the first part, we analyze the combinatorial changes for a fixed k. As Chew and Kedem, we consider functions such that combinatorial changes are induced by breakpoints of the lower envelope of the functions. By partitioning the functions into finer groups than Chew and Kedem did, we improve the upper bound of the combinatorial changes to $O(n^2)$ from the previous bound $O(n\lambda_3(n))$.

In the second part, we analyze the combinatorial changes to eDT with respect to k. A combinatorial change to eDT corresponds to a quadruplet of pairs, each pair consisting of an element of Q and an element of P touching each other in some placement of a scaled copy of P simultaneously. To count the quadruplets inducing combinatorial changes to eDT, we consider the triplets of such pairs and define a function for each triplet. For the lower envelope L of the functions, a combinatorial change corresponds to an intersection of two such functions appearing on L. There are $O(k^3 n^2)$ such functions and two functions intersect at most four times. To reduce the bound, we classify the functions into types such that any two functions belonging to the same type intersect each other less than four times. By applying the partition method in the first part and the classification of the functions, we show that the upper bound of combinatorial changes becomes $O(k^2 n^2\lambda_4(k))$.

*Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Korea. {tkeom0114, juny2400}@postech.ac.kr

†Department of Computer Science and Engineering, Graduate School of Artificial Intelligence, Pohang University of Science and Technology, Pohang, Korea. heekap@postech.ac.kr

Pattern Matching in Doubling Spaces

Corentin Allair
École Polytechnique, Paris, France
corentin.allair@polytechnique.edu

Antoine Vigneron
Department of Computer Science and Engineering
UNIST, Republic of Korea
antoine@unist.ac.kr

July 23, 2021

A metric space has *doubling dimension* δ if any ball can be covered by at most 2^δ balls of half its radius. When $\delta = O(1)$, we say that this space is *doubling*. For instance, the Euclidean space $\mathbb{R}^d$ has doubling dimension $O(d)$, hence doubling spaces are generalizations of fixed-dimensional Euclidean spaces.

We study pattern matching problems in doubling spaces. Given two doubling spaces (X, d_X) and (Y, d_Y) of doubling dimension δ, and sizes $|X| = k$ and $|Y| = n$, where $k \leqslant n$, our goal is to find a subspace of Y that resembles the *pattern* X. More precisely, we consider the *ρ-distortion problem* and the *minimum distortion problem*.

Given $\rho \geqslant 1$, the ρ-distortion problem is to find, if it exists, a mapping $\sigma : X \to Y$ such that

$$(1/\rho)d_X(x, x') \leqslant d_Y(\sigma(x), \sigma(x')) \leqslant \rho d_X(x, x')$$

for all $x, x' \in X$. The ρ-distortion problem is analogous to the problem of matching two point-sets in Euclidean space under rigid transformations. If, in addition, we allow scaling, then an analogous problem in general metric spaces is the minimum distortion problem.

We first give a hardness result: We show that for any $\rho \geqslant 1$, the k-clique problem reduces to ρ-distortion in doubling dimension $\log_2 3$. It implies that this problem cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f, unless the exponential time hypothesis (ETH) is false. On the positive side, we present a near-linear time approximation algorithm for small values of k. More precisely, if $0 < \varepsilon \leqslant 1$, our algorithm returns in $2^{O(k^2 \log k)}(\rho^2/\varepsilon)^{2k\delta} n \log n$ time a solution to the $(1+\varepsilon)\rho$-distortion problem whenever a solution to the ρ-distortion problem exists.

We also show how to extend these results to the minimum distortion problem. In particular, we show that the minimum distortion problem cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f, unless ETH is false, and we give a $(1+\varepsilon)$-approximation algorithm running in $2^{O(k^2 \log k)}(\mathrm{dist}(X, Y)^{2k\delta}/\varepsilon^{2k\delta+O(1)})n^2 \log n$ time, where $\mathrm{dist}(X, Y)$ denotes the minimum distortion between X and Y.

Lazy Data Types

Jihoon Hyun, Sewon Park, and Martin Ziegler

KAIST, School of Computing

The standard evaluation/interpretation of a term is generically defined via structural induction, computationally corresponding to recursion over the expression tree. Here, if a partial operation at some leaf renders the corresponding subterm undefined, then the induction/recursion breaks down and makes the entire term undefined. In order to still yield a (possibly multi-)defined value in (at least some) such cases, both Mathematics and Computer Science often consider alternative, tailored notions of evaluation/interpretation over some ad-hoc extended structures:

Example 1. a) $1/x$ or $\ln(x)$ are mathematically undefined when $x = 0$.

b) $t \equiv \sqrt{x^2}$ is undefined over the *reals* when $x < 0$. One benefit (or arguably even purpose) of introducing *complex* numbers is to make $t(-1)$ well-defined.

c) $t \equiv \frac{x^2-4}{x-2}$ is mathematically undefined when $x = 2$: according to the standard interpretation. But the underlying singularity is removable, and doing so yields value $t(x = 2) = 4$.

d) $t \equiv (x \leq 1) \vee \ln(x) > 0$ is undefined when $x \leq 0$: according to the standard interpretation of Propositional Calculus over *Boolean* structure. However over *Kleene* Logic, $t \equiv \texttt{true}$ holds for *all* (real) x.

e) The dichotomy $(x < 0) \vee (x \geq 0) \equiv \texttt{true}$ is unprovable in Constructive Mathematics, but trivial in Classical Logic.

f) Consider the following code fragment making a logical assignment involving access to the array $x[]$: BOOL b = $\big(n < \mathrm{size}(x)\big)$ && $(x[n] == 0)$. It is valid in `C++` with left-to-right short-cut evaluation, but not in right-to-left nor in eager evaluation as possible for instance in `Pascal` [3].

g) Classical bisection for numerical Root Finding tests the sign of $w := f(z)$ at $z := (x + y)/2$ to proceed either with $y := z$ or with $x := z$. However, for $w = 0$, sign calculation is considered 'unstable', and in fact uncomputable [5, Exercise 4.2.9]. Put differently, a test "$w > 0$" is mathematically assigned the 'value' `unknown`—which computationally however never actually gets returned.

h) *Tri*section still makes Root Finding total by testing the signs of both $f(u)$ and $f(v)$ simultaneously, where $u := (2x+y)/3$ and $v := (x+2y)/3$ [2, p. 336]. Note that, if f has a unique root, then at least one test is guaranteed to succeed, i.e., return `true` or `false`. This amounts to a multivalued operation on tuples of Kleenean Logic $\mathbb{K} = \{\texttt{true}, \texttt{false}, \texttt{unknown}\}$, see [4,1].

We devise a systematic and unified perspective on the above ad-hoc approaches, combining the theory of *representations* [5, §3] with the *Exact Real Computation* paradigm `arXiv:1608.05787`.

References

1. Martín Hötzel Escardó, Martin Hofmann, and Thomas Streicher. On the non-sequential nature of the interval-domain model of real-number computation. *Mathem. Structures in Computer Science*, 14(6):803–814, 2004.
2. Peter Hertling. Topological complexity with continuous operations. *Journal of Complexity*, 12:315–338, 1996.
3. B. Kernigham. Why Pascal is not my favorite programming language. Technical Report 100, Bell laboratories, 1981.
4. Horst Luckhardt. A fundamental effect in computations on real numbers. *Theoretical Computer Science*, 5(3):321–324, 1977.
5. Klaus Weihrauch. *Computable Analysis*. Springer, Berlin, 2000.

Efficient Construction of Cryptarithm Catalogues over Deterministic Finite Automata

Koya Watanabe[1], Diptarama Hendrian[1], Ryo Yoshinaka[1], Takashi Horiyama[2], and Ayumi Shinohara[1]

[1] Graduate School of Information Sciences, Tohoku University
{koya_watanabe@shino.ecei., diptarama@, ryoshinaka@, ayumis@}tohoku.ac.jp
[2] Faculty of Information Science and Technology, Hokkaido University
horiyama@ist.hokudai.ac.jp

Cryptarithm is a famous mathematical puzzle where given an arithmetic equation written with letters rather than numerals, a player must discover an assignment of numerals on letters that makes the equation hold true. A well-known example of a cryptarithm instance is `send`+`more` = `money`. Solving a cryptarithm is rather easy: one can find a solution by trying at most 10! assignments: i.e., cryptarithms of addition are solvable by brute force in linear time under the *decimal* system. However, when the base is not fixed, the problem becomes NP-hard [1].

Nozaki et al. [3] proposed a *cryptarithm deterministic finite automata (DFAs)* that accept (uniquely) solvable addition cryptarithm instances under a reasonable encoding for different bases. Their automata can be seen as complete catalogues of cryptarithm with which one can solve, count, and enumerate cryptarithm instances. Their idea of representing cryptarithm instances by an automaton is to assign a tentative solution set to each state. Since there are $k!$ possible assignments of numerals to an alphabet of size k, the size of the automaton will quickly explode by increasing the base k. To tackle this problem, they proposed an extension of DFAs with which solvable instances can be represented much more compactly than the naive ones, taking advantage of the symmetry among assignments. The compressed DFA is constructed by merging states equivalent modulo permutation. As an expense of this memory saving technique, their method for constructing compact automata required much more time than the naive construction. As a result, they succeeded in constructing the DFAs for different bases up to 7.

We propose an even more compact data structure representing solvable cryptarithm instances and an efficient construction algorithm, by resolving the two issues posed by them as future work. One is to leverage the symmetry between the first and second summand terms in cryptarithm instances based on the commutative property of addition. We managed to formalize the idea for merging further states and moreover symmetric edges, without impairing the competence of the automata as cryptarithm catalogues. Since the edges occupy much of the memory space of the automata, reducing edges is quite effective in memory saving. The other issue is to accelerate the construction of cryptarithm automata, as their state merging technique slowed the construction as mentioned above. We define a representative as the "canonical form" among states to be merged, and design an algorithm to convert states into their canonical forms. The use of canonical states significantly reduces construction time.

We implemented the method and succeeded in constructing cryptarithm automata for bases up to 9. Our experiments demonstrated significant efficiency improvement by our technique. For example, while the method by Nozaki et al. took more than four hours and 3.7 GB for constructing the base-7 cryptarithm automaton according to their paper, ours took less than a minute and 0.4 GB, i.e., more than 100 times faster and nearly 10 times smaller, though the comparison is not precise for the difference of their and our experimental environments. Compared to the naive cryptarithm DFA for base 7, our elaborated DFA has nearly 100 times less states and edges. Constructing the decimal cryptarithm DFA was difficult due to the memory usage and time consumption, but we were able to construct the decimal system cryptarithm DFA where instances include at most 7 different letter.

One may (actually we did) think that πDDs [2], which are efficient data structures for manipulating permutation sets, would be useful for compactly representing tentative solution sets assigned to the states of the automata. However, πDDs saved very little memory usage, while computational over head was considerable, because actual tentative solution sets have about three elements only in average in the base-7 cryptarithm automaton and most memory was used for edges rather than states in our implementation.

References

1. Eppstein, D.: On the NP-completeness of cryptarithms. ACM SIGACT News **18**(3), 38–40 (1987)
2. Minato, S.: πDD: A new decision diagram for efficient problem solving in permutation space. In: Sakallah, K.A., Simon, L. (eds.) Theory and Applications of Satisfiability Testing (SAT 2011). LNCS, vol. 6695, pp. 90–104 (2011)
3. Nozaki, Y., Hendrian, D., Yoshinaka, R., Shinohara, A.: Enumeration of cryptarithms using deterministic finite automata. In: International Conference on Implementation and Application of Automata. pp. 286–298. Springer (2018)