

Pattern Avoidance in Permutations

Nik Ruškuc, St Andrews

`nik@mcs.st-and.ac.uk`

`http://www-groups.mcs.st-and.ac.uk/~nik`



Cambridge, 15 May, 2003

Sorting With a Stack

Consider the set of all finite permutations that can be sorted by a stack.

Proposition. A permutation $\pi = \pi_1\pi_2\dots\pi_n$ can be sorted by a stack if there do not exist $i, j, k \in \{1, \dots, n\}$ such that $i < j < k$ and $\pi_k < \pi_i < \pi_j$.

Proposition. The number of permutations of length n which can be sorted by a stack is equal to the n th Catalan number C_n .



Pattern Involvement

Let us have two sequences: $\sigma = \sigma_1\sigma_2\dots\sigma_m$ and $\tau = \tau_1\tau_2\dots\tau_n$ of (say) positive integers.

- σ and τ are order isomorphic ($\sigma \cong \tau$) if $m = n$ and for all $i, j = 1, \dots, n$ we have

$$\sigma_i \leq \sigma_j \iff \tau_i \leq \tau_j.$$

- σ is involved in τ ($\sigma \preceq \tau$) if τ contains a subsequence $\tau_{i_1}, \tau_{i_2}, \dots, \tau_{i_m}$ ($i_1 < \dots < i_m$) order isomorphic to σ .
- If σ is not involved in τ we say that τ avoids σ .

Example. $231 \preceq 142635$. $231 \not\preceq 142365$.

Closed Classes: Definition

Notation. S_n – the set of all permutations on $\{1, \dots, n\} = [n]$ (written as sequences of images).

$$S = \bigcup_{i=0}^{\infty} S_n.$$

Proposition. The following two conditions are equivalent for a class (set) $\mathcal{C} \subseteq S$ of permutations:

- (i) $\pi \in \mathcal{C} \ \& \ \rho \preceq \sigma \Rightarrow \rho \in \mathcal{C}.$
- (ii) $\sigma \in S \setminus \mathcal{C} \ \& \ \sigma \preceq \tau \Rightarrow \tau \in S \setminus \mathcal{C}.$
- (iii) There exists a (unique) antichain B of permutations such that $\sigma \in \mathcal{C}$ iff σ avoids all permutations in B .

Definition. If \mathcal{C} satisfies one (and hence all) of the above conditions we say that \mathcal{C} is a closed class. The set B described in (iii) is called the basis for \mathcal{C} .

Closed Classes: Examples

Example. The class of permutations sorted by a stack is closed. Its basis is $\{231\}$.

Example. Permutations obtained by splitting $12\dots n$ into two subsequences and interleaving them arbitrarily. For instance: $21356487 \in \mathcal{C}$, $321 \notin \mathcal{C}$. The basis of this class is $\{321\}$.

Example. The riffle-shuffle class \mathcal{R} : all the permutations obtained by splitting $12\dots n$ into two subsequences $12\dots i$ and $i+1, \dots, n$ and interleaving them.

For instance: $12673485 \in \mathcal{R}$, $2143, 2413 \notin \mathcal{R}$. The basis of this class is $\{321, 2143, 2413\}$.

Closed Classes: General Questions

Given a closed class \mathcal{C} (by its basis, sorting mechanism, generating mechanism, etc.), we ask the following questions:

Basis Problem. Determine the basis of \mathcal{C} . Failing this, is the basis finite? What is the number of basis permutations of length n ($n = 1, 2, 3, \dots$).

Membership Problem. Does there exist an algorithm which for every $\pi \in \mathbf{S}$ determines whether or not $\pi \in \mathcal{C}$? Is there an efficient (linear, polynomial, etc.) algorithm for doing this?

Enumeration Problem. Determine the numbers $e_n(\mathcal{C}) = |\mathcal{C} \cap \mathbf{S}_n|$ of all permutations in \mathcal{C} of length n .

Remarks

- Not all closed classes are finitely based, because there exist infinite antichains of permutations. For example the permutations:

$$(2, 3), 5, 1, 7, 4, 9, 6, \dots, 2i - 1, 2i + 2, \dots, \\ \dots 2n + 1, 2n - 2, (2n + 2, 2n + 3), 2n$$

($n = 2, 3, 4, \dots$) form such an antichain (Spielman, Bona).

- Finite basis implies a polynomial membership testing algorithm.

Enumeration

- If \mathcal{C} is a class defined by a single basis permutation of length 3 then $e_n(\mathcal{C}) = C_n$.
- Gessel: a formula for the enumeration of the class with basis $\{1234\}$.
- Bona: basis $\{1342\}$.
- Regev: an asymptotic formula for basis $\{123\dots n\}$.
- Simion, Schmidt: all basis permutations of length 3.
- West: one basis permutation of length 3, and one of length 4.

Stanley–Wilf Conjecture

Conjecture. (Stanley–Wilf) For every closed class \mathcal{C} there exists q such that $e_n(\mathcal{C}) \leq q^n$.

Remark. It is sufficient to consider classes with bases of size one.

Theorem. (Noga, Friedgut) For $e_n(\mathcal{C}) \leq q^{n\gamma(n)}$ where γ is a function that increases very (very) slowly.

Conjecture. (Gessel) The enumeration sequence of every finitely based closed class satisfies a recursive formula with polynomial coefficients.

Remark. Not true for infinitely based closed classes (Atkinson, Murphy).

Two Stacks in Series

\mathcal{S}^2 : permutations that can be sorted with two stacks connected in series.


Example. $231 \in \mathcal{S}^2$, $2435761 \notin \mathcal{S}^2$.

Fact. \mathcal{S}^2 is not finitely based.

Questions. Find the enumeration sequences for \mathcal{S}^2 and its basis.

Conjecture. The membership problem for \mathcal{S}^2 is NP-complete (Atkinson). The membership problem for \mathcal{S}^2 is polynomial (Murphy).*

* Delete as appropriate.





Dunedin–St Andrews Collaboration



- M.D. Atkinson, M.M. Murphy, N. Ruskuc: two ordered stacks in series: basis and enumeration.
- M.D. Atkinson, M.M. Murphy, N. Ruskuc: partially well ordered closed classes; the class of one-stack sortable permutations is p.w.o.
- M. Albert, M.D. Atkinson, N. Ruskuc: closed classes and languages.
- M.M. Murphy: extensive structure theory, new methods for defining closed classes, beginnings of a theory of antichains.



Bounded Classes


Suppose that we are generating permutations using a mechanism with finite memory of size M .

If the system is full it must output before taking more input. So we cannot generate any permutation of the form

$$(M + 1, t_1, \dots, t_M), \quad (t_1, \dots, t_M < M + 1). \quad (*)$$


Let Ω_M be the class with basis consisting of all permutations of the form (*).

Every closed subclass of Ω_M is called bounded.



Graphs as Generators

$S_{2,2}$ – two stacks of size 2 in series:



A closed class generated by a graph is called a graph class.

Regular Languages

Let A be a finite alphabet.


A language is any set of (finite) words over A .

A language is regular if it is accepted by an automaton.

Regular expressions: well-formed expressions over $A \cup \{ (,), *, +, \cdot, \cap, - \}$.

Theorem. (Kleene) A language is regular if and only if it is defined by a regular expression.

Example. The language accepted by the automaton M is defined by the regular expression $(a + b)^* - (a + b)^* aa(a + b)^*$.



Regular Languages: Membership and Enumeration

Fact. For every regular language L there exists an algorithm which tests the membership in L in linear time.

Proof. Feeding the word through an automaton accepting L .

For a language L let $e_n(L) = |L \cap A^n|$ be the number of words in L of length n .

Fact. If L is regular then $e_n(L)$ satisfy a linear recurrence with constant coefficients.

Regular Languages: Membership and Enumeration

Example. Consider the language accepted by the automaton M . Let $L^{(i)}$ be the language accepted by M with s_i as the only accept state. Then:

$$\begin{aligned}L_0^{(0)} &= 1, \quad L_0^{(1)} = 0, \quad L_0^{(2)} = 0, \\L_{n+1}^{(0)} &= L_n^{(0)} + L_n^{(1)}, \\L_{n+1}^{(1)} &= L_n^{(0)}, \\L_{n+1}^{(2)} &= L_n^{(1)} + 2L_n^{(2)}.\end{aligned}$$

Hence

$$L_{n+1} = L_{n+1}^{(0)} + L_{n+1}^{(1)} = 2L_n^{(0)} + L_n^{(1)} = L_n + L_{n-1}.$$

Encoding Ω_M

For $\pi = \pi_1\pi_2 \dots \pi_n \in \mathbf{S}$ let

$$E(\pi) = p_1p_2 \dots p_n$$

where

$$p_i = |\{j : j \geq i, \pi_j \leq \pi_i\}|$$

(the rank of π_i among $\{\pi_i, \dots, \pi_n\}$). For example

$$E(24163875) = 23131321.$$

Fact. If $A \subseteq \Omega_M$ then $E(A) \subseteq [M]^*$.

Regular Classes

Definition. A set $A \subseteq \Omega_M$ is regular if $E(A)$ is a regular language over $[M]^*$.

Fact. Ω_M itself is regular.

Proof. $E(\Omega_M) = [M]^* - [M]^*F$, where F is a finite set of words.

Fact. Not every subclass of Ω_M is regular.

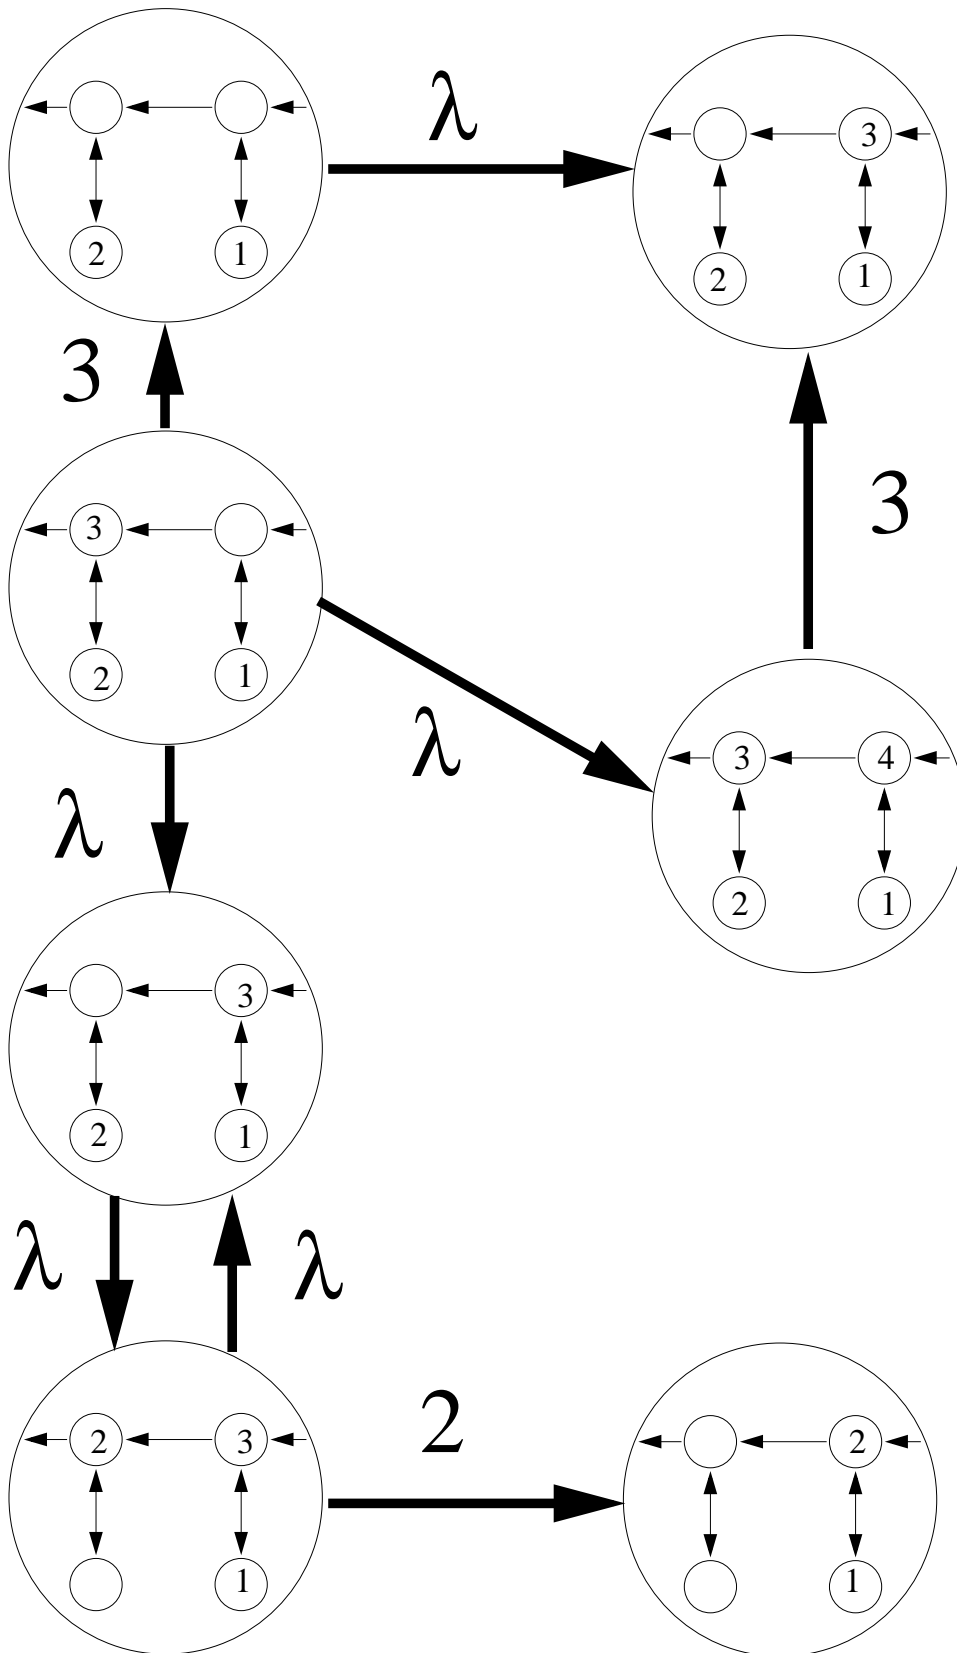
Proof. Ω_M is not partially well ordered, and hence it has uncountably many closed subclasses (Atkinson, Murphy, Ruskuc).

Examples

Example. Let \mathcal{X} be the closed class with basis $\{312, 321, 231\}$. 312 and 321 ensure that $\mathcal{X} \subseteq \Omega_2$. 231 then implies that the subword 22 never appears in the encoding. Thus $E(\mathcal{X}) = E(\Omega_2) - (1 + 2)^*22(1 + 2)^*$.

Theorem. (**Atkinson, Livsey, Tuley**) Every graph class is regular.

Proof. Construct an automaton. Here are some states and transitions for $\mathcal{S}_{2,2}$.



Deleting/Inserting Elements

Lemma. (Albert, Atkinson, Ruskuc) For any $\mathcal{A} \subseteq \Omega_M$ let:

- \mathcal{A}_d be the set of all permutations obtained by deleting a single entry in a permutation from \mathcal{A} ;
- \mathcal{A}_D be the set of all permutations obtained by deleting any number of entries in a permutation from \mathcal{A} ;
- \mathcal{A}_i be the set of all permutations obtained by inserting a single entry in a permutation from \mathcal{A} ;
- \mathcal{A}_I be the set of all permutations obtained by inserting any number of entries in a permutation from \mathcal{A} .

If \mathcal{A} is regular then so are \mathcal{A}_d , \mathcal{A}_D , \mathcal{A}_i and \mathcal{A}_I .

Regular Classes: General Results

(Albert, Atkinson, Ruskuc)

Theorem 1. There is an algorithm which decides whether or not a given regular set $L \subseteq [M]^*$ is the encoding of a (regular) closed class in Ω_M .

Theorem 2. A closed bounded class \mathcal{C} is regular if and only if its basis B is regular. Moreover, there are algorithms which compute $E(B)$ from $E(\mathcal{C})$ and vice versa.

Regular Classes: General Results

Corollary 2A. There is an algorithm which decides whether or not a given closed, bounded class is finitely based or not.

Facts. (I) The membership problem in any regular bounded class is decidable in linear time.
(II) The enumeration sequence of a regular bounded class satisfies a linear recurrence with constant coefficients.

Natural Classes

Let $\pi : \mathbb{N} \rightarrow \mathbb{N}$ be a permutation.

Let \mathcal{X} be the set of all finite permutations involved in π . Clearly \mathcal{X} is closed.

Closed classes obtained in this way are called natural.

Example. If $\pi = (2, 1, 3, 4, 5, 6, 7, \dots)$ then \mathcal{X} consists of all permutations of the forms $123 \dots n$ and $2134 \dots n$.

Sum Complete Classes

For permutations $\alpha = \alpha_1 \dots \alpha_m$ and $\beta = \beta_1 \dots \beta_n$, define

$$\alpha \oplus \beta = \alpha_1 \dots \alpha_m \beta'_1 \dots \beta'_n,$$

where $\beta'_i = \beta_i + m$.

A closed class \mathcal{X} is said to be sum complete if

$$\alpha\beta \in \mathcal{X} \Rightarrow \alpha \oplus \beta \in \mathcal{X}.$$

Examples. All the following classes are sum complete: one stack sortable, two stacks in series sortable, graph classes, etc. The natural class \mathcal{X} defined by $\pi = (2, 1, 3, 4, \dots)$ is not sum complete, since $21 \in \mathcal{X}$ but $21 \oplus 21 = 2143 \notin \mathcal{X}$.


Proposition. (Atkinson) Every sum complete class is natural.

Proposition. (Atkinson) A closed class is sum complete if and only if all its basis permutations are sum indecomposable.

Murphy's Alternative

Theorem. (Murphy) If \mathcal{X} is a natural class defined by a permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$, then at least one of the following two statements is true:

- (1) $\mathcal{X} = \mathcal{F} \oplus \mathcal{C}$ where \mathcal{F} is finite and \mathcal{C} is sum complete.
- (2) π is periodic and (hence) \mathcal{X} is regular.



Deciding the Join Property

Theorem. There is an algorithm which determines for any finitely based class whether or not it is natural.

By taking $\pi : \mathbb{Z} \rightarrow \mathbb{Z}$ we can define integral classes.

Question. Is there an algorithm which decides whether or not a finitely based class is integral?

A closed class \mathcal{C} is said to have the join property if for every $\alpha, \beta \in \mathcal{C}$ there exists $\gamma \in \mathcal{C}$ such that $\alpha \preceq \gamma$ and $\beta \preceq \gamma$.

Fact. Every natural (resp. integral) class has the join property.

Question. Is there an algorithm which decides for any finitely based closed class whether or not it has the join property?