The Möbius function of the consecutive pattern poset

Luca Ferrari (University of Firenze)

For the poset of classical permutation patterns, the first results about its Möbius function were obtained by Sagan and Vatter. Further results have been found by Steingrímsson and Tenner and by Burstein, Jelínek, Jelínková and Steingrímsson. The general problem in this case of classical patterns seems quite hard. In contrast, the poset of consecutive pattern containment has a much simpler structure. Here we compute the Möbius function of that poset. In most cases our results give an immediate answer. In the remaining cases, we give a polynomial time recursive algorithm to compute the Möbius function. In particular, we show that the Möbius function only takes the values -1, 0 and 1.

An interesting result to note in connection to this is Björner's one on the Möbius function of factor order. Although that poset is quite different from ours, there are interesting similarities. In particular, both deal with consecutive subwords and the possible values of the Möbius function are -1, 0 and 1 in both cases.

Denote with $\mathcal{P}$ the poset of permutations with respect to consecutive pattern containment, and take $\sigma, \tau \in \mathcal{P}$ such that $\sigma \leq \tau$. In order to present our results, we need a couple of definitions.

Suppose σ occurs in $\tau = a_1 a_2 \dots a_n$. If a_{i+1} is the leftmost letter of τ involved in any occurrence of σ in τ, we say that τ has a *left tail of length i with respect to σ*. Analogously, τ has a *right tail of length j with respect to σ* if a_{n-j} is the rightmost letter of τ involved in any occurrence of σ in τ. If it is clear from the context what σ is, we simply talk about left and right tails of τ.

For example, with respect to the pattern 123, the permutation 286134759 has a left tail of length 3, and a right tails of length 2, since all occurrences of 123 belong within the segment 1347.

The following definition is borrowed from the theory of codes.

Given a permutation τ, its *prefix* (resp. *suffix*) *pattern of length k* is the permutation of length k order isomorphic to the prefix (resp. suffix) of τ of length k. In other words, the prefix (resp. suffix) pattern of length k of τ is the unique permutation $\sigma \in S_k$ such that τ has a left (resp. right) tail of length 0 with respect to σ. In case the prefix and suffix patterns of length k of τ coincide, we say that τ has a *bifix pattern of length k*.

In the case where σ occurs precisely once in τ, we show that $\mu(\sigma, \tau)$ depends only on the lengths, a and b, of the two tails of τ with respect to σ. More precisely, $\mu(\sigma, \tau)$ is 1 if $a = b \leq 1$, it is -1 if $a = 0$ and $b = 1$ or vice versa, and 0 otherwise (in which case τ has a tail of length at least 2).

Our main result deals with intervals $[\sigma, \tau]$ where σ occurs at least twice in τ. This result implies that, as in the case of one occurrence, if τ has a tail of length at least 2, then $\mu(\sigma, \tau) = 0$. In the remaining cases, where the tails of τ have length at most 1, the main result gives a recursive algorithm for computing $\mu(\sigma, \tau)$, by producing, if possible, an element $\mathcal{C}$ in $[\sigma, \tau]$, where $|\mathcal{C}| < |\tau| - 2$, such that $\mu(\sigma, \tau) = \mu(\sigma, \mathcal{C})$. This element $\mathcal{C}$, if it exists, must be a bifix pattern of τ, and it must lie below the two elements covered by τ, but not below the element obtained by deleting one letter from each end of τ. If no such element $\mathcal{C}$ exists (which is most often the case), we have $\mu(\sigma, \tau) = 0$.

(This is joint work with Antonio Bernini and Einar Steingrimsson.)