It was my great pleasure to review Marek Sokołowski's doctoral dissertation, supervised by Michał Pilipczuk. In the last four years, Marek has worked on several topics, most of them on graph algorithms and data structures. Marek chose to include five of his papers in his dissertation, opting for the theme of graph width parameters (treewidth, rank-width, twin-width) in relation with efficient data structures to store, or maintain under edge editions, graphs of low width.

The first paper, titled "Dynamic Treewidth," presents a data structure that takes linear initialization time in the number n of vertices of the graph, and maintains a tree-decomposition of width at most $6k + 5$ of the current graph, as well as the validity of a monadic second-order logic sentence on it, under edge insertions or deletions. Each update takes subpolynomial amortized time $n^{o(1)}$. If the treewidth goes above k, the data structure flags it with 'treewidth too large' and is exempted from its supposed behavior until the treewidth goes again below k. This result is technical and involves many insights and ingredients: (minimal) c-small k-closures, how the Dealternation Lemma of Bojańczyk and Pilipczuk yields them, relation of closures and linkedness, blockages/explorations/collected components, etc. On a very high level, after an edge uv is inserted in G, $\{u, v\}$ is added to a path from the topmost bags containing u and v, respectively, to the root of the current tree-decomposition $\mathcal{T} = (T, \mathrm{bag})$. This defines a prefix in the tree-decomposition (a subtree containing the root). An (in some sense minimal) c-small k-closure of this prefix—another prefix T' of T—is computed. An optimum tree-decomposition (of low height) if recomputed on the union X of bags of the subtree T'. The components of $G - X$ are rearranged in tree-decompositions that are attached to that of $G[X]$. As it is crucial to the running time that the tree-decomposition has subpolynomial height, a rebalancing procedure is also used. The amortized time of an update is controlled by a potential function defined on the tree-decomposition, which is roughly $2^{O(k \log k)}$ times the sum of the heights of every rooted subtree (hence lower, the more balanced the tree is).

The second contribution is the lift of this result to the dense counterparts of treewidth: rankwidth and cliquewidth. This is even more technical. It generally follows "Dynamic Treewidth," and introduces dense analogues of the objects developed there. Yet many steps require more work. Most prominently, the Dealternation Lemma, which was readily available for treewidth has to be established for rankwidth. This alone takes no less than twenty pages. This part, and other elements elsewhere, build on the work of Jeong, Kim and Oum efficiently finding rank-decompositions for subspace arrangements.

"Dynamic Treewidth" is the first non-trivial result (i.e., with sublinear updates) on this topic, despite the question/line of work being open for decades. For the work on dynamic rankwidth, there is another remarkable reward: it yields the fastest algorithm to compute the rank-width k of a graph. By adding the m edges one by one from the n-vertex edgeless graph, one attains running time $O_k(mn^{o(1)})$, whereas the previously known running time was $O_k(n^2)$. Actually the edge insertions can be batched, and one gets an even better running time of $O_k(n^{1+o(1)} + m)$.

The third contribution deals with a dynamization of Baker's technique. (Even if Baker's technique relies on creating induced subgraphs of small treewidth, the "Dynamic Treewidth" result and ideas are not utilized here.) This dynamization is exemplified on the weighted versions of Max Independent Set and Bounded-Degree Min Dominating Set on graphs excluding a fixed apex graph as a minor (such as planar graphs). An ad hoc recursive data structure is presented that supports edge editions, and edge reweighting.

If the dissertation (and Marek's results during his PhD) stopped here, it would already be a major success. The first two contributions are major results. The first one finally unblocks an elusive (and important) task, identified more than three decades ago. The second one almost closes a long line of work in computing rankwidth. It is astounding that Marek has done much more, in particular achieving powerful results on

- solving Max Independent Set in classes excluding some induced subgraphs,
- the study of sparsity and stability motivated by efficient first-order model checking,
- classical graph algorithms (4-edge-connectivity, Shortest Path, etc.),
- the recent topic of twin-width.

The last two contributions of the manuscript are actually on twin-width. This parameter was introduced in 2020, roughly at the same time as Marek was starting to do research. The fourth paper designs for graphs of twin-width at most d, a data structure on $O_d(n)$ bits that answers edge queries in time $O(\log\log n)$. If the graph is given as an adjacency matrix without large mixed minor, then the construction time is polynomial. If further a list of $O_d(n)$ rectangles partitioning the 1-entries is given, then the construction is almost linear in n (hence, sublinear in the number of matrix entries). This improved both the size of the data structure and the query time over the previous compact representations of graphs of bounded twin-width. The new data structure is a clever recursive construction: $m \times m$ matrices points to $m^{2/3} \times m^{2/3}$ matrices, thereby making a tree of height $O(\log\log n)$. This relies on analyzing how horizontal, vertical, and mixed cells can coexist without creating large mixed minors.

The fifth and last contribution shows that classes of bounded twin-width are quasipolynomially χ-bounded. The paper introduces almost mixed minors: like mixed minors, but the diagonal cells are left unconstrained. It is then shown how to break an adjacency matrix without d-almost mixed minor into blocks of considerably smaller clique number or no $(d-1)$-almost mixed minor. Again this is a beautiful piece of work, which brought a handful of key ideas to the then very recent topic of twin-width. This work later enabled Bourneuf and Thomassé to show that classes of bounded twin-width are even polynomially χ-bounded.

Verdict. *I deem the thesis as sufficient to grant a PhD.* This sentence (that I was asked to write, should I be positive about the dissertation) is an immense understatement. There are two excellent PhDs in Marek's manuscript, and about four in his whole doctoral output. The expertise in data structures (apparently sharpened through years of sport programming) is on display. I am very impressed with the depth of knowledge and understanding already possessed by Marek on various topics. It is exemplary how he worked with his colleagues in Warsaw, including his advisor and one competitive programming teammate, as well as a fellow PhD student from Bergen (with whom Marek coauthored the colossal contribution on Dynamic Rankwidth).

I'm looking forward to the PhD defense and to Marek's future great contributions.

Here are some sparse, minor comments that I took along the way or rare typos that I noticed. Those are so minor that I do not oppose to the thesis being published as is, and I do not need to see the dissertation after a possible correction.

- p19: "$\text{parent}_T(u)$" $\to$ "$\text{parent}_T(x)$"
- p35: "in the spirit" $\to$ "in spirit"
- p39: the expression after "so it is possible that" lacks what is being summed, $\text{height}_T(a)$.
- p39: neither big or shallow $\to$ neither big nor shallow
- p71: "that is both j-small and j-deep:" Don't you mean a full stop rather than a colon? I don't see how the latter part clarifies the former part.
- p71: In Lemma 3.5.3, don't you mean all numbers to be integers? At least it is the case if you assume d integral.
- p76: "If X is the set of descendants of x". You set a notation for that in Chapter 2.
- p103: you could recall that in $\vec{E}(T)$ edges are oriented from a node to its parent. I think you only suggest it via $\vec{L}(T)$ in Chapter 2.

- p103: a figure illustrating ENC 4 for each kind of 3-vertex path xyz would help.
- p176: Theorem 4.10.7, “an decomposition”
- p182: Corollary 5.2.3, you further assume that $\mathcal{C}$ is minor-closed.
- p183: “$\{uv\}$ otherwise” Isn’t it cleaner to further require that uv is an edge?
- p186: Definition 11, “a subset of V(H)” $\rightarrow$ “a subset of V(G)”

Édouard Bonnet
CNRS researcher
ENS Lyon, LIP