

COMPUTING FINITE TYPE INVARIANTS EFFICIENTLY

DROR BAR-NATAN, ITAI BAR-NATAN, IVA HALACHEVA, AND NANCY SCHERICH

(Communicated by David Futer)

ABSTRACT. We describe an efficient algorithm to compute finite type invariants of type k by first creating, for a given knot K with n crossings, a look-up table for all subdiagrams of K of size $\lceil \frac{k}{2} \rceil$ indexed by dyadic intervals in $[0, 2n - 1]$. Using this algorithm, any such finite type invariant can be computed on an n -crossing knot in time $\tilde{O}(n^{\lceil \frac{k}{2} \rceil})$, a lot faster than the previously best published bound of $\tilde{O}(n^k)$.

1. INTRODUCTION

Finite type invariants, also known as Vassiliev invariants [Vas90, Vas92], underlie many of the classical knot invariants, for instance they include the coefficients of the Jones, Alexander, and more generally HOMFLY-PT polynomials [BL93, BN95]. A knot invariant ζ is said to be of *finite type k* (equivalently of degree k) if it vanishes on all knots with at least $k + 1$ double points, where ζ is extended to knots with double points by the formula:

$$\zeta(\text{X}) = \zeta(\text{Y}) - \zeta(\text{Z}).$$

For example, the linking number of a two-component knot is a finite type invariant of type 1. In our Main Theorem, we provide an algorithm to compute finite type invariants from a planar projection of a knot in a surprisingly efficient time depending on the crossing number of the knot diagram.

For complexity measurements in this paper, we care only about polynomial degree (i.e. we ignore constants and $\log(n)$ terms). In this vein, recall that $f(n) = \tilde{O}(g(n))$ means that there exist positive numbers c, p, N so that for all $n > N$, we have that $f(n) < cg(n)(\log g(n))^p$.

Main Theorem. *Finite type invariants of type k can be computed on an n -crossing knot in time $\tilde{O}(n^{\lceil k/2 \rceil})$.*

The result is a surprising result, as before this theorem the fastest algorithm (known to the authors) to compute a type k invariant on a knot diagram with

Received by the editors December 30, 2024, and, in revised form, January 20, 2025, May 26, 2025, July 23, 2025, and July 24, 2025.

2020 *Mathematics Subject Classification.* Primary 57K16.

Key words and phrases. Finite type invariants, Gauss diagrams.

The first author was supported by NSERC-RGPIN-2018-04350 and by the Chu Family Foundation (NYC). The third author was supported by the Natural Science Foundation Grant No. DMS-2302664. This material is also based upon work supported by the National Science Foundation under Grant No. DMS-1929284 while the fourth author was in residence at the Institute for Computational and Experimental Research in Mathematics in Providence, RI, during the Braids program. This project is partially sponsored by the Provost Office of Elon University.

n crossings took time $\tilde{O}(n^k)$ [BNBNHS23], and it was commonly believed that this was the fastest possible. There are specific finite type invariants which can be computed much faster, such as the linking number and the coefficients of the Alexander polynomial, but these are special cases. The Main Theorem gives the current fastest known algorithm that works for *all* finite type invariants, and shows that the computational time can be reduced to roughly the square root of the previously fastest known algorithm. In a previous paper [BNBNHS23], we proved that finite type invariants can be computed efficiently using 3D methods. We argued that most knot invariants, including finite type invariants, should be more efficiently computed using 3D methods rather than 2D methods. However, the Main Theorem is significant as it is a 2D method that currently outperforms all known 3D methods to compute finite type invariants.

For impatient readers, the key formula in this paper is Equation (3). The preceding pages include the definitions leading up to the formula, and the proof that the formula can be evaluated in time $\tilde{O}(n^{\lceil k/2 \rceil})$. The proof uses a space-time tradeoff and the Main Theorem can be generalized to highlight the space requirements.

Space-aware version of the Main Theorem. *For any integer $f \leq \frac{k}{2}$, any finite type invariant of type k can be computed on an n -crossing knot in time $\tilde{O}(n^{k-f})$ using storage space $\tilde{O}(n^f)$.*

2. BACKGROUND

2.1. Gauss diagrams. A *Gauss diagram* of an n -crossing long knot diagram parametrized by an open interval $I_{\mathbb{R}} \subseteq \mathbb{R}$ is given by the interval $I_{\mathbb{R}}$ along with n decorated arrows (equivalently, oriented perfect matchings of $2n$ points, where the arrows are the edges of the matching). Usually $I_{\mathbb{R}} = (-1, 2n)$ but $I_{\mathbb{R}}$ can be any open interval. To distinguish between intervals of real numbers and integers, we use the notation of a subscript of \mathbb{R} or \mathbb{Z} on the interval. So $I_{\mathbb{Z}} = I_{\mathbb{R}} \cap \mathbb{Z}$. Each arrow corresponds to one of the n crossings of the knot and has endpoints at the integer values in $I_{\mathbb{R}}$ (the endpoints are in $I_{\mathbb{Z}}$). The head of an arrow is at the point in $I_{\mathbb{Z}}$ which parametrizes the lower strand of the crossing and the tail of the arrow is at the point which parametrizes the upper strand of the crossing (we assume all crossing points occur at the integer location of the parametrization). Each arrow is decorated with a sign corresponding to the sign of the crossing. We say that such a diagram is *numbered* by $I_{\mathbb{Z}}$, where usually $I_{\mathbb{Z}} = [0, 2n]_{\mathbb{Z}}$. Figure 1 (A) shows an example of a Gauss diagram. For a Gauss diagram D , the quantity $|D|$ is the number of arrows in D . Let $\mathcal{GD} = \langle \text{Gauss diagrams} \rangle$ denote the \mathbb{Z} -module of \mathbb{Z} -linear combinations of Gauss diagrams, and let \mathcal{GD}_k denote the subspace spanned by Gauss diagrams with k or fewer arrows.

Let D be a Gauss diagram with n arrows numbered by $[0, 2n]_{\mathbb{Z}}$. A *k -arrow subdiagram* of D is a diagram consisting of $I_{\mathbb{R}}$ and a subset of k decorated arrows from D . A subdiagram corresponds to a choice of k crossings in the knot diagram represented by the Gauss diagram D . An example is shown in Figure 1(B). Notice that a subdiagram D' of D keeps the original numbering along the interval $I_{\mathbb{R}}$ and the $2k$ endpoints of D' will be spread out amongst the $2n$ points in $I_{\mathbb{R}}$. A k -arrow subdiagram of D is *not* a proper Gauss diagram because of this numbering issue. To make a k -arrow subdiagram D' into a Gauss diagram, we can apply the forgetful map ψ to D' which rennumbers $[0, 2n]_{\mathbb{Z}}$ monotonically to $[0, 2k]_{\mathbb{Z}}$, in




FIGURE 1. (A) An example of the Gauss diagram of a long knot diagram. (B) A 2-arrow subdiagram of a Gauss diagram. (C) The forgetful map ψ applied to a subdiagram yielding a renumbered subdiagram.

essence forgetting how D' was realized as a subdiagram of D . We will call $\psi(D')$ a *renumbered subdiagram* of D . An example is shown in Figure 1(C).

Given two disjoint subdiagrams E and F of a Gauss diagram D , with $e = |E|$ and $f = |F|$, we say that the *pattern* $P = \pi_D(E, F)$ of E and F within D is the $2e$ -element subset of $[0, 2(e + f)]_{\mathbb{Z}}$ numbering the ends of the images of the arrows of E within $\psi(E \cup F)$ (and so the complementary subset $[0, 2(e + f)]_{\mathbb{Z}} \setminus P$ numbers the arrow ends of F). In the special case where $D = E \cup F$, we simply write $P = \pi(E, F)$.

Conversely, if E and F are Gauss diagrams with $e = |E|$ and $f = |F|$ and P is an appropriately-sized *pattern* (namely, $P \subset [0, 2(e + f)]_{\mathbb{Z}}$ and $|P| = 2e$), then there is a unique Gauss diagram D with $|D| = e + f$ containing E and F as disjoint subdiagrams such that $\pi_D(E, F) = P$. We denote $D = E \#_P F$ and say that D is the *superimposition* of E and F with pattern P . We extend $\#_P$ to a bilinear map $\#_P : \mathcal{GD}_e \times \mathcal{GD}_f \rightarrow \mathcal{GD}_{e+f}$. Two examples are shown in Figure 2.

We denote by $\varphi_k : \{\text{knot diagrams}\} \rightarrow \mathcal{GD}_k$ the map which sends a knot diagram to the sum of all of the *renumbered subdiagrams* of its Gauss diagram which have exactly k arrows. Note that $\varphi_k = \psi \circ \bar{\varphi}_k$ where $\bar{\varphi}_k$ sends a knot diagram to the sum of all of the subdiagrams of its Gauss diagram which have exactly k arrows, and ψ renumbers each summand to make it a Gauss diagram. Let $\varphi_{\leq k} := \sum_{i=1}^k \varphi_i$ be the map that sends a knot diagram to the sum of all of the renumbered subdiagrams




FIGURE 2. Two examples of superimposing the same diagrams along different patterns.

of its Gauss diagram which have *at most* k arrows. The maps φ_k and $\varphi_{\leq k}$ are *not* invariants of knots but every finite type invariant factors through $\varphi_{\leq k}$, as follows from the next theorem.

Theorem 2.1 (Goussarov-Polyak-Viro [GPV00], see also [Rou07]). *A \mathbb{Q} -valued knot invariant ζ is of type k if and only if there is a linear functional ω on $\mathbb{Q} \otimes \mathcal{GD}_k$ such that $\zeta = \omega \circ \varphi_{\leq k}$.*

We show that $\varphi_{\leq k}$ can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$. This result, combined with the above theorem, proves that all finite type invariants of type k can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$, which is the main result of this paper.

It is surprising that $\varphi_{\leq k}$ can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$ because, at first glance, it would seem that one must require time $\tilde{O}(n^k)$.¹ A Gauss diagram with n arrows has $\sum_{i=1}^k \binom{n}{i}$ subdiagrams with k or fewer arrows. Because $\binom{n}{i} = \tilde{O}(n^i)$, a Gauss diagram with n arrows has $\sum_{i=1}^k \binom{n}{i} = \tilde{O}\left(\sum_{i=1}^k n^i\right) = \tilde{O}(n^k)$ subdiagrams with k or fewer arrows. So $\bar{\varphi}_{\leq k}$ evaluated on a knot diagram with n crossings will be a sum of $\tilde{O}(n^k)$ subdiagrams. Note that the outputs of $\bar{\varphi}_k$ and φ_k have the same number of summands, but φ_k will have repeated terms and $\bar{\varphi}_k$ will not. So where do the computational savings come from? The idea is to break the computation of $\varphi_{\leq k}$ into two parts, one of which can be quickly pre-computed in a look-up table. The creation of this look-up table uses counting techniques taking advantage of dyadic intervals. These techniques are completely self-contained, and unrelated to finite type invariants and knot theory. In the next section, we describe these techniques that we will apply to prove that $\varphi_{\leq k}$ can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$.

3. COMPUTATIONAL PRELIMINARIES: COUNTING TECHNIQUES USING DYADIC INTERVALS

3.1. Counting with a look-up table. For the purpose of this paper, a *look-up table* is a lexicographically-ordered list of (key \mapsto value) entries, or more formally a lexicographically-ordered associative array. Below, Theorem 3.1 shows how to use

¹We slightly abuse notation and use \tilde{O} also to mean “equal up to constants and logarithms”.

a look-up table to count elements of a set inside $[n]^\ell$, where $[n] := [1, n]_{\mathbb{Z}}$. While we will need a generalized version of this theorem, the proof of Theorem 3.1 showcases nicely how dyadic intervals are used to get computational savings.

Theorem 3.1. *Let Q be an enumerated subset of $[n]^\ell$ with $|Q| = q = n^f$, where $0 \leq f \leq \ell$ (typically $0 < f < \ell$ and so Q is “big” yet “much smaller” than $[n]^\ell$). In time $\tilde{O}(q)$, a look-up table of size $\tilde{O}(q)$ can be created so that computing $|Q \cap R|$ will take time $\tilde{O}(1)$ for any rectangular box $R \subset [n]^\ell$.*

The straightforward approach to this theorem would be to create a look-up table with key-value pairs of the form $(R \mapsto |Q \cap R|)$ for all possible rectangular boxes R . A rectangular box in $[n]^\ell$ is determined by choosing two interval endpoints in each coordinate of $[n]^\ell$, so there are $\tilde{O}(n^{2\ell})$ possible rectangular boxes in $[n]^\ell$. This look-up table would take much longer than $\tilde{O}(q)$ to create. The trick is to create a restricted look-up table using only rectangular boxes which intersect Q non-trivially and whose sides are dyadic intervals. The structure is reminiscent of the data structures quadtrees/octrees which are used in computer graphics.

To prove Theorem 3.1, let us first start with some preliminaries on dyadic intervals.

3.2. Dyadic intervals. For the purpose of this paper, a *dyadic interval* is a half-closed interval of the form $[2^p q, 2^p(q + 1))_{\mathbb{Z}}$ for $p, q \in \mathbb{Z}_{\geq 0}$. A dyadic interval can be expressed via a binary expansion up to a certain accuracy, i.e. a fixed sequence of 0’s and 1’s followed by a fixed number of *’s or ‘free’ entries, and the integers contained in the interval are all possible completions of the expansions, i.e. all possible ways of replacing the *’s with 0’s and 1’s. For example $[2^2 2, 2^2 3)_{\mathbb{Z}} = \{8, 9, 10, 11\} = \{1000_2, 1001_2, 1010_2, 1011_2\}$, expressed in decimal and binary expansion respectively. Using the binary expansion, every number in the interval is of the form $10**$ where the least significant two bits are free and the most significant two bits are fixed to be 10. We write u for a dyadic interval where u is a binary sequence followed by some number of *’s.

A dyadic interval u is *maximal* in $(b, c)_{\mathbb{Z}}$ if $u \subset (b, c)_{\mathbb{Z}}$ and any larger dyadic interval containing u is not contained in $(b, c)_{\mathbb{Z}}$. Notice if two dyadic intervals overlap, then one is contained in the other. Therefore, the maximal dyadic intervals of $(b, c)_{\mathbb{Z}}$ are disjoint. Every interval decomposes uniquely as a disjoint union of maximal dyadic intervals, as described in the following elementary Lemma.

Lemma 3.2. *For any interval $(b, c)_{\mathbb{Z}} \subset \mathbb{Z}_{\geq 0}$, there are at most $2 \log_2(c - b)$ maximal dyadic intervals contained in $(b, c)_{\mathbb{Z}}$, and $(b, c)_{\mathbb{Z}}$ is a disjoint union of its maximal dyadic intervals.*

For a binary number, a truncation process can be used to compute dyadic intervals containing that binary number.

Truncation process. Given a binary number x with m bits of accuracy, starting from right to left, replacing one bit at a time with a * generates a list of m dyadic intervals which contain x . For example, let $x = 7 = 00111_2$ where $m = 5$, then the truncation process generates the list of dyadic intervals $\{00111, 0011*, 001**, 00***, 0****\}$. If a number x (in decimal notation) is at most ℓ , then the number of bits needed to describe x in binary expansion is at most $\log_2(\ell)$, and so x is contained in at most $\log_2(\ell)$ dyadic intervals coming from the truncation process with $\log_2(\ell)$ bits of accuracy.

3.3. Proof and generalization of Theorem 3.1. We proceed with the proof of Theorem 3.1, as well as its generalization, Proposition 3.3, which we will use in the proof of the Main Theorem.

Proof of Theorem 3.1. A dyadic rectangular box in $[n]^\ell$ is a product of ℓ dyadic intervals of $[n]$. We describe a process to create a restricted look-up table of key-value pairs of the form $(R_d \mapsto |Q \cap R_d|)$ for only R_d that are dyadic rectangles in $[n]^\ell$ and such that $|Q \cap R_d| > 0$.

To create the table, run through the enumerated elements $x \in Q$. For each x , from the truncation process described in Section 3.2, x is contained in at most $(\log_2(n))^\ell$ dyadic rectangular boxes R_d in $[n]^\ell$. In the table, increment the value for R_d by 1 for each such R_d containing x , or create such an entry if it didn't already exist. Since there are q elements in Q , creating this table takes time $\tilde{O}(q(\log_2(n))^\ell) = \tilde{O}(q)$, and there are $\tilde{O}(q)$ elements in the table. Using standard binary sorting techniques, accessing and modifying values in the table takes time $\tilde{O}(\log_2 q) = \tilde{O}(1)$.

A general non-dyadic rectangular box R in $[n]^\ell$ is the disjoint union of at most $(2\log_2(n))^\ell = \tilde{O}(1)$ maximal dyadic rectangular boxes, by Lemma 3.2. To count $|Q \cap R|$, one retrieves the values in the look-up table with key each maximal dyadic rectangular box in R , and then sums together all those stored values to get $|Q \cap R|$. Since retrieval takes time $\tilde{O}(1)$ and the sum is over $\tilde{O}(1)$ elements, after the look-up table is completed, it takes time $\tilde{O}(1)$ to compute $|Q \cap R|$. \square

Viewing $|Q \cap R|$ as the sum $\sum_{x \in Q \cap R} 1$, Theorem 3.1 can be generalized, using almost exactly the same proof, to compute weighted sums $\sum_{x \in Q \cap R} \theta_x$ where the weights θ_x are valued in some \mathbb{Z} -module. This generalization is stated below. It is the version of Theorem 3.1 that will be used to prove the main result in Section 4.

Proposition 3.3. *Let Q be an enumerated subset of $[n]^\ell$ with $|Q| = q = n^f$, where $0 \leq f \leq \ell$. Let M be a free \mathbb{Z} -module of rank $\tilde{O}(1)$, and let $\theta : [n]^\ell \rightarrow M$ be a map that is zero outside of Q . Then, in time $\tilde{O}(q)$, a look-up table of size $\tilde{O}(q)$ can be created so that computing $\sum_R \theta$ will take time $\tilde{O}(1)$ for any rectangular box $R \subset [n]^\ell$.²*

4. MAIN RESULT

In this section, we state and prove the Main Theorem of this article. The strategy to quickly compute φ_k on a diagram D is to view a k -arrow subdiagram of D as the superimposition of two smaller subdiagrams E and F . The subdiagram E is placed inside D first, and instead of placing F , a look-up table is created to count in how many ways F could have been placed. Using Proposition 3.3, this look-up table can be computed and accessed very quickly, which ultimately gives the computational savings for the final result.

²There is a small caveat to Proposition 3.3 which requires that in some basis for M , the coefficients of θ must be computable in time $\tilde{O}(1)$ and their size must be $\tilde{O}(1)$, which may not always be the case. However, for the purposes of this paper and the application for which this proposition is used, this issue does not arise and so we chose to leave this technical detail out of the statement.

Main Theorem. *Finite type invariants of type k can be computed on an n -crossing knot in time $\tilde{O}(n^{\lceil k/2 \rceil})$*

Proof. By Theorem 2.1, it suffices to show that $\varphi_{\leq k}$ can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$. Since $\varphi_{\leq k} = \sum_{i=1}^k \varphi_i$, it suffices to prove that φ_k can be computed in time $\tilde{O}(n^{\lceil k/2 \rceil})$. This shows that $\varphi_{\leq k}$ can be computed in time $\tilde{O}(k \cdot n^{\lceil k/2 \rceil}) = \tilde{O}(n^{\lceil k/2 \rceil})$

For a knot K with n crossings, let K be represented as a Gauss diagram with n arrows. By definition, $\varphi_k(K)$ is the sum of all renumbered subdiagrams with exactly k arrows,

$$\varphi_k(K) = \sum_{D \subset K, |D|=k} \psi(D),$$

where $D \subset K$ means D is a subdiagram of K .

Fixing a choice of $e, f \in \mathbb{Z}_{\geq 0}$ with $e + f = k$, every k -arrow subdiagram D can be viewed in $\binom{k}{e}$ ways as the superimposition of two smaller subdiagrams E and F of sizes e and f respectively. Note that $e, f, k = \tilde{O}(1)$. As discussed above, rather than breaking a specific k -arrow subdiagram down as a superimposition, one can instead build up all k -arrow subdiagrams by first choosing an e -arrow subdiagram E and then choosing an f -arrow subdiagram F that lies in the complement of E in K . Summing over the possible choices of E, F , and superimpositions gives the next formula³:

$$(1) \quad \varphi_k(K) = \sum_{\substack{D \subset K \\ |D|=k}} \psi(D) = \binom{k}{e}^{-1} \sum_{\substack{E \subset K \\ |E|=e}} \sum_{\substack{\text{patterns} \\ P}} \sum_{\substack{F \subset K, |F|=f \\ \pi_K(E, F)=P}} \psi(E) \#_P \psi(F).$$

The sum above overcounts every k -arrow subdiagram D of K by a factor of $\binom{k}{e}$ as every choice of splitting D into two pieces (i.e. a choice of E) occurs exactly once in the sum. Hence we include the pre-factor $\binom{k}{e}^{-1}$.

Define $\theta_K : ([0, 2n - 1]_{\mathbb{Z}})^f \rightarrow \mathcal{GD}_f$ by

$$(F_0, F_1, \dots, F_{2f-1}) \mapsto \begin{cases} \psi(F) & \text{if } (F_0, F_1, \dots, F_{2f-1}) \text{ are the ends of a subdiagram } F \subset K \\ 0 & \text{otherwise.} \end{cases}$$

Now, the innermost sum from Equation (1) can be rewritten as a sum of values of θ_K as follows:

$$(2) \quad \sum_{\substack{F \subset K, |F|=f \\ \pi_K(E, F)=P}} \psi(E) \#_P \psi(F) = \psi(E) \#_P \left(\sum_{\prod_{i=0}^{2f-1} [m_{i,E,P}^-, m_{i,E,P}^+]} \theta_K \right),$$

where $m_{i,E,P}^-$ is the minimal possible place for the i 'th arrow-end of a Gauss diagram F within $[0, 2n - 1]_{\mathbb{Z}}$ given that E along with F should make the pattern P , and $m_{i,E,P}^+$ is the maximal such value. Both of these quantities are easily determined by E and P , though we spare the reader the explicit formulas.

³We note that all the scalars in this formula and in the formulas that follow are integers of magnitude $\tilde{O}(n^k)$. Additions of such integers take time $\tilde{O}(1)$, and hence we need not worry about the time cost of working with large integers.

The upshot is that the sum on the right can be computed very quickly with a look-up table. Proposition 3.3 applies by taking Q to be the set of all

$$(F_0, F_1, \dots, F_{2f-1}) \in ([0, 2n - 1]_{\mathbb{Z}})^f$$

that are the endpoints of a subdiagram of K . Here, $|Q| = \binom{n}{f} = \tilde{O}(n^f)$ and the conclusion from Proposition 3.3 is that a look-up table can be created in time $\tilde{O}(n^f)$ so that computing the sum on the right hand side of Equation (2) takes time $\tilde{O}(1)$.

Thus we arrive at our final equation,

$$(3) \quad \varphi_k(K) = \binom{k}{e}^{-1} \sum_{\substack{E \subset K \\ |E|=e}} \sum_{\substack{\text{patterns} \\ P}} \psi(E) \#_P \left(\sum_{\prod_{i=0}^{2f-1} [m_{i,E,P}^-, m_{i,E,P}^+]} \theta_K \right).$$

To understand the computational complexity of Equation (3), we need to understand the complexity of each sum. We already showed that the innermost sum can be computed in time $\tilde{O}(n^f)$ by building a look-up table from Proposition 3.3. For the middle sum, P can be any subset of $[0, 2(e + f)]_{\mathbb{Z}}$ of size $2e$. There are $\binom{2(e+f)}{2e} = \binom{2k}{2e} = \tilde{O}(1)$ such subsets. That does not add to the complexity of the total sum.

For the fixed choice of $e + f = k$, the outer sum in Equation (3) has at most $\tilde{O}\left(\binom{n}{e}\right) = \tilde{O}(n^e)$ terms. Therefore, the total time of computing $\varphi_k(K)$, which includes first creating the look-up table and then computing the sum over all subdiagrams $D \subseteq K$ of size k , is $\tilde{O}(n^e + n^f)$. This sum can be minimized for $e = \lceil \frac{k}{2} \rceil$ and $f = \lfloor \frac{k}{2} \rfloor$, in which case we get computation time $\tilde{O}(n^{\lceil \frac{k}{2} \rceil})$. \square

Note that the lookup table within the proof of the Main Theorem is of size $\tilde{O}(n^f)$, where at the end, f is set to be $\lfloor \frac{k}{2} \rfloor$, and so our algorithm uses storage space $\tilde{O}(n^{\lfloor \frac{k}{2} \rfloor})$. If storage space is limited we can set f differently and get the following:

Main Theorem (space-aware version). *For any integer $f \leq \frac{k}{2}$, any finite type invariant of type k can be computed on an n -crossing knot in time $\tilde{O}(n^{k-f})$ using storage space $\tilde{O}(n^f)$.* \square

ACKNOWLEDGMENTS

We would like to thank ICERM for hosting the first and third authors for a week long visit. Finally we thank A. Referee for their extremely valuable comments.

REFERENCES

[BL93] Joan S. Birman and Xiao-Song Lin, *Knot polynomials and Vassiliev's invariants*, Invent. Math. **111** (1993), no. 2, 225–270, DOI 10.1007/BF01231287. MR1198809

[BN95] Dror Bar-Natan, *On the Vassiliev knot invariants*, Topology **34** (1995), no. 2, 423–472, DOI 10.1016/0040-9383(95)93237-2. MR1318886

[BNBNHS23] Dror Bar-Natan, Itai Bar-Natan, Iva Halacheva, and Nancy Scherich, *Yarn ball knots and faster computations*, J. Appl. Comput. Topol. **8** (2024), no. 1, 175–192, DOI 10.1007/s41468-023-00144-7. MR4707460

[GPV00] Mikhail Goussarov, Michael Polyak, and Oleg Viro, *Finite-type invariants of classical and virtual knots*, Topology **39** (2000), no. 5, 1045–1068, DOI 10.1016/S0040-9383(99)00054-3. MR1763963

- [Rou07] F. Roukema, *Goussarov-Polyak-Viro combinatorial formulas for finite type invariants*, Preprint, [arXiv:0711.4001](https://arxiv.org/abs/0711.4001), 2007.
- [Vas90] V. A. Vassiliev, *Cohomology of knot spaces*, Theory of singularities and its applications, Adv. Soviet Math., vol. 1, Amer. Math. Soc., Providence, RI, 1990, pp. 23–69. MR1089670
- [Vas92] V. A. Vassiliev, *Complements of discriminants of smooth maps: topology and applications*, Translations of Mathematical Monographs, vol. 98, American Mathematical Society, Providence, RI, 1992. Translated from the Russian by B. Goldfarb, DOI 10.1090/mmono/098. MR1168473

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TORONTO, TORONTO, ONTARIO M5S 1A1,
CANADA

Email address: drorbn@math.toronto.edu

URL: <http://www.math.toronto.edu/drorbn>

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CALIFORNIA 90095

Email address: itaibn@math.ucla.edu

DEPARTMENT OF MATHEMATICS, NORTHEASTERN UNIVERSITY, BOSTON, MASSACHUSETTS 02115

Email address: i.halacheva@northeastern.edu

URL: <https://sites.google.com/site/ivahalacheva3/>

DEPARTMENT OF MATHEMATICS AND STATISTICS, ELON UNIVERSITY, 100 CAMPUS DR, ELON,
NORTH CAROLINA 27244

Email address: nscherich@elon.edu

URL: <http://www.nancyscherich.com>