Invent math (2010) 182: 419–447
DOI 10.1007/s00222-010-0275-6

Khovanov homology and the slice genus

Jacob Rasmussen

Received: 25 February 2004 / Accepted: 31 March 2005 / Published online: 8 September 2010

Abstract We use Lee's work on the Khovanov homology to define a knot invariant s. We show that $s(K)$ is a concordance invariant and that it provides a lower bound for the smooth slice genus of K. As a corollary, we give a purely combinatorial proof of the Milnor conjecture.

1 Introduction

In [8], Khovanov introduced an invariant of knots and links, now widely known as the Khovanov homology. This invariant takes the form of a graded homology theory $H_{Kh}(L)$, whose graded Euler characteristic is the unnormalized Jones polynomial of L. In [11], Lee showed that $H_{Kh}(L)$ is naturally viewed as the E^1 term of a spectral sequence which converges to $\mathbf{Q} \oplus \mathbf{Q}$. In this paper, we use this spectral sequence to define a knot invariant $s(K)$. The definition of $s(K)$ was motivated by a similar invariant $\tau(K)$ which is defined using knot Floer homology [17, 20]. In fact, the similarities between the two invariants extend far beyond their manner of definition.

Our main result is that the invariant s gives a lower bound for the slice genus:

The author was partially supported by an NSF Postdoctoral fellowship.

J. Rasmussen (✉)
Dept. of Mathematics, Princeton University, Princeton, NJ 08544, USA
e-mail: j.rasmussen@dpmms.cam.ac.uk

Present address:
J. Rasmussen
DPMMS, University of Cambridge, Cambridge CB5 8DW, UK

Theorem 1

$$|s(K)| \leq 2g_*(K)$$

where $g_*(K)$ denotes the smooth slice genus of K; in other words, the minimal genus of a smoothly embedded, orientable surface in B^4 which bounds K.

In fact,

Theorem 2 *The map s induces a homomorphism from $Conc(S^3)$ to $\mathbf{Z}$, where $Conc(S^3)$ denotes the smooth concordance group of knots in S^3.*

It is well known that analogous statements in the topological category hold for the Trotter-Murasugi knot signature $\sigma(K)$. For alternating knots, s does not provide any information beyond that given by σ:

Theorem 3 *If K is an alternating knot, then $s(K) = \sigma(K)$.*

There is, however, a class of knots for which $s(K)$ gives much better—indeed, sharp—information. We say that a knot is *positive* if it admits a planar diagram with all positive crossings.

Theorem 4 *If K is a positive knot,*

$$s(K) = 2g_*(K) = 2g(K)$$

where $g(K)$ is the ordinary genus of K.

As a corollary, we get a Khovanov homology proof of following result, which was first proved by Kronheimer and Mrowka using gauge theory [10]:

Corollary 1 (The Milnor Conjecture) *The smooth slice genus and unknotting number of the (p, q) torus knot are both equal to $(p-1)(q-1)/2$.*

The invariant s is sensitive to the smooth, rather than just the topological slice genus. The easiest way to see this is to exhibit a knot K with $\Delta_K(t) = 1$, but $s(K) \neq 0$. By a theorem of Freedman [5], such a K is topologically slice but cannot be smoothly slice by Theorem 1. Such knots are not difficult to find; perhaps the simplest example is the $(-3, 5, 7)$ pretzel knot. An elegant proof that this knot has $s = 2$ may be found in [19] or [23], in which it is shown that s satisfies a version of Bennequin's inequality. As Robert Gompf kindly pointed out to the author, such knots can be used to give a proof of the existence of an exotic $\mathbf{R}^4$ without using gauge theory. (See [7], pp. 377 and 522 for details.)

As the reader familiar with knot Floer homology will have already noted, the theorems above all hold with $2\tau(K)$ in place of $s(K)$. (See [17] for the first three, and [13] and [21] for the final one.) Indeed, the equality $s(K) = 2\tau(K)$ holds in all cases for which the author knows the value of $\tau(K)$. Based on these observations, we make the following (perhaps optimistic)

Conjecture *For any knot $K \subset S^3$, $s(K) = 2\tau(K)$.*[1]

Readers familiar with the Khovanov homology may also have observed that the notation $s(K)$ has already been used by Bar-Natan [1] to describe an apparent knot invariant which appears in one of his "phenomenological conjectures." This is no coincidence. Indeed, the author's interest in the subject was first aroused by the observation that Bar-Natan's s appeared to give a lower bound for the slice genus. Although we are unable to prove that the $s(K)$ defined here is the same as that determined by Bar-Natan's conjecture, we do give a fairly general condition (at least for small knots) under which the two agree.

The remainder of the paper is organized as follows. In Sect. 2, we review the Khovanov complex and Lee's construction of a spectral sequence from it. In Sect. 3, we define s and show that it behaves nicely with respect to the structure of the concordance group. Section 4 is devoted to the proof of Theorem 1. In Sect. 5, we prove Theorems 3 and 4, and discuss the relationship between $s(K)$ and $\tau(K)$ in more detail. Finally, Sect. 6 contains proofs of some technical results establishing the invariance of Lee's spectral sequence, which are needed in Sect. 2.

Finally, we take this opportunity to fix two conventions which we will use throughout. First, we will always work with $\mathbf{Q}$ coefficients. Although Khovanov's complex can be defined with coefficients in $\mathbf{Z}$, Lee's theorem (Theorem 2.2) does not hold in this context. Second, we will often abuse our notation, letting L refer both to a planar diagram of a link and to the underlying link itself. The reader should have little trouble determining from context which meaning is intended.

2 Review of Khovanov homology

In this section, we briefly recall the construction of the Khovanov complex [8] and Lee's extension of it [11].

2.1 The cube of resolutions

Given an oriented link diagram L with crossings labeled 1 through k, we can form the cube of all possible resolutions of L. This is a k-dimensional cube

[1] The conjecture has been disproved by Hedden and Ording, who showed that $s \neq 2\tau$ for certain 2-strand cables of the trefoil knot.

Fig. 1 0- and 1-resolutions of a crossing

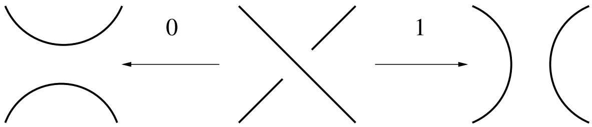

with its vertices and edges decorated by 1-manifolds and cobordisms between them. More specifically, each crossing of L can be resolved in two different ways, as illustrated in Fig. 1. To each vertex v of the cube $[0, 1]^k$, we associate the planar diagram D_v obtained by resolving the i-th crossing of L according to the i-th coordinate of v. Then D_v is a planar diagram without crossings, so it is a disjoint union of circles.

Let e be an edge of the cube. The coordinates of its two ends differ in one component—say the l-th. We call the end which has a 0 in this component the *initial end*, and denote it by $v_e(0)$. The other end is called the *terminal end*, written $v_e(1)$. We assign to e the cobordism $S_e : D_{v_e(0)} \to D_{v_e(1)}$, which is a product cobordism except in a neighborhood of the l-th crossing, where it is the obvious saddle cobordism between the 0 and 1-resolutions.

The Khovanov complex is constructed by applying a $1 + 1$ dimensional TQFT $\mathcal{A}$ to the cube of resolutions. In other words, one replaces each vertex v by the group $\mathcal{A}(D_v)$, and each edge e by the map $\mathcal{A}(S_e)$. The underlying group of $C_{Kh}(L)$ is the direct sum of the groups $\mathcal{A}(D_v)$ for all vertices v, and the differential on the summand $\mathcal{A}(D_v)$ is a sum of the maps $\mathcal{A}(S_e)$ for all edges e which have v as their initial end. More precisely, for $x \in \mathcal{A}(D_v)$

$$d(x) = \sum_{i=1}^{c_0(v)} (-1)^{s(e_i)} \mathcal{A}(S_{e_i}) \tag{1}$$

Here $c_0(v)$ is the number of crossings in v which have a 0-resolution, and e_i is the edge which corresponds to changing the i-th such crossing to a 1-resolution. The signs $(-1)^{s(e_i)}$ are chosen in such a way that $d^2 = 0$. (There are many different ways to do this, but it is easy to see that they all give rise to isomorphic chain complexes.) The *homological grading* of an element $x \in \mathcal{A}(D_v)$ is defined to be $\mathrm{gr}(v) = |v| - n_-$, where $|v|$ is the number of 1's in the coordinates of v and n_- is the number of negative crossings in the diagram for L. Note that d increases the homological grading by 1—strictly speaking, the Khovanov homology is a cohomology theory!

2.2 Khovanov's TQFT

We now give a more explicit description of the TQFT $\mathcal{A}$. Let V be a vector space spanned by two elements, $\mathbf{v}_+$ and $\mathbf{v}_-$. The vector space associated by $\mathcal{A}$ to a single circle is defined to be V, so that if D is a diagram composed of n disjoint circles, $\mathcal{A}(D) = V^{\otimes n}$. Thus we can think of $C_{Kh}(L)$ as being the vector space spanned by the space of "states" for L, where a state consists of a complete resolution of L, together with a labeling of each component of the resolution by either $\mathbf{v}_+$ or $\mathbf{v}_-$.

The cobordisms S_e come in two forms: either two circles can merge into one, or one can split into two. In the first case, $\mathcal{A}(S_e)$ is given by a map $m: V^{\otimes 2} \to V$, where the two factors in the tensor product correspond to the labels on the two circles that merge, and the copy of V in the image corresponds to the label on the single resulting circle. Likewise, in the second case, $\mathcal{A}(S_e)$ is given by a map $\Delta: V \to V^{\otimes 2}$. The formulas for these maps are

$$\begin{aligned} &m(\mathbf{v}_+ \otimes \mathbf{v}_+) = \mathbf{v}_+ && \Delta(\mathbf{v}_+) = \mathbf{v}_+ \otimes \mathbf{v}_- + \mathbf{v}_- \otimes \mathbf{v}_+ \\ &m(\mathbf{v}_+ \otimes \mathbf{v}_-) = m(\mathbf{v}_- \otimes \mathbf{v}_+) = \mathbf{v}_- && \Delta(\mathbf{v}_-) = \mathbf{v}_- \otimes \mathbf{v}_- \\ &m(\mathbf{v}_- \otimes \mathbf{v}_-) = 0 \end{aligned} \tag{2}$$

For reference, we also record two other maps ι and ϵ used to define $\mathcal{A}$. These maps are not needed at the moment, but they make an appearance in Sect. 4 when we study cobordisms. Corresponding to the addition of a 0-handle (the birth of a circle in a diagram), there is a map $\iota: \mathbf{Q} \to V$, and corresponding to the addition of a two handle (the death of a circle) there is a map $\epsilon: V \to \mathbf{Q}$. These maps are given by

$$\begin{aligned} &\epsilon(\mathbf{v}_-) = 1 \qquad && \iota(1) = \mathbf{v}_+ \\ &\epsilon(\mathbf{v}_+) = 0 \end{aligned}$$

$\mathcal{A}$ is especially nice because it is a graded TQFT. We define a grading p on V by setting $p(\mathbf{v}_\pm) = \pm 1$ and extend it to $V^{\otimes n}$ by $p(v_1 \otimes v_2 \otimes \cdots \otimes v_n) = p(v_1) + p(v_2) + \cdots + p(v_n)$. Then it is easy to see that if $\mathbf{v}$ is a homogeneous element of $V^{\otimes n}$, $p(S_e(\mathbf{v})) = p(\mathbf{v}) - 1$. Next, we define a grading q on $C_{Kh}(L)$ by $q(\mathbf{v}) = p(\mathbf{v}) + \mathrm{gr}(\mathbf{v}) + n_+ - n_-$, where $n_\pm$ are the number of positive and negative crossings in the diagram L. (The term $n_+ - n_-$ is included so that the q-grading remains invariant for different diagrams of the same knot.) Then $q(d(\mathbf{v})) = q(\mathbf{v})$, so $C_{Kh}(L)$ splits into a direct sum of complexes, one for each q grading. In fact, its graded Euler characteristic is the unnormalized Jones polynomial of L, but we will not make use of this here.

In [8], Khovanov proves that the homology of $C_{Kh}(L)$ (thought of as a bigraded group) is an invariant of the underlying link L. We denote this homology group by $H_{Kh}(L)$. We remark that although the definition of $C_{Kh}(L)$ required an oriented link L, the choice of orientation plays a relatively minor role. Indeed, the only place it is used is in determining n_+ and n_-. Thus the operation of global orientation reversal has no effect on $H_{Kh}(L)$, and if we change the orientation on some components but not others, the effect is to change the homological and q gradings by an overall shift.

2.3 Lee's TQFT

In [11], Lee considers a similar construction, but with another TQFT $\mathcal{A}'$ in place of $\mathcal{A}$. The underlying vector spaces for these two TQFT's are the same, but the maps $m' : V \otimes V \to V$ and $\Delta' : V \to V \otimes V$ induced by cobordisms are slightly different. They are given by

$$\begin{aligned} m'(\mathbf{v}_+ \otimes \mathbf{v}_+) = m'(\mathbf{v}_- \otimes \mathbf{v}_-) = \mathbf{v}_+ \qquad & \Delta'(\mathbf{v}_+) = \mathbf{v}_+ \otimes \mathbf{v}_- + \mathbf{v}_- \otimes \mathbf{v}_+ \\ m'(\mathbf{v}_+ \otimes \mathbf{v}_-) = m'(\mathbf{v}_- \otimes \mathbf{v}_+) = \mathbf{v}_- \qquad & \Delta'(\mathbf{v}_-) = \mathbf{v}_- \otimes \mathbf{v}_- + \mathbf{v}_+ \otimes \mathbf{v}_+ \end{aligned} \tag{3}$$

(The maps ι and ϵ corresponding to the addition of 0 and 2-handles are the same as for $\mathcal{A}$.) We denote the resulting complex by $C_{Lee}(L)$ and its homology by $H_{Lee}(L)$.

Using the obvious identification between the underlying groups of $C_{Kh}(L)$ and $C_{Lee}(L)$, we can define a q-grading on the latter group as well. It is clear from equation 3 that this grading does not behave quite so well with respect to the differential d'. Indeed, $\Delta'(\mathbf{v}_-)$ is not even homogeneous. It is easy to see, however, that if $\mathbf{v} \in C_{Lee}(L)$ is a homogeneous element, then the q-grading of every monomial in $d'(\mathbf{v})$ is greater than or equal to the q-grading of $\mathbf{v}$. In other words, the q-grading defines a filtration on the complex $C_{Lee}(L)$. This fact leads to the following theorem, which is implicit in [11]:

Theorem 2.1 *There is a spectral sequence with E^1 term $H_{Kh}(L)$ which converges to $H_{Lee}(L)$. The E^1 and higher terms of this spectral sequence are invariants of the link L.*

The first part of the theorem is more or less immediate from the observations above. The filtration on C_{Lee} gives rise to a spectral sequence converging to H_{Lee}. The differential on its E^0 term is the part of d' which preserves (rather than raises) the q-grading. Comparing Eqs. 2 and 3, we see that the E^0 term is the complex C_{Kh}.

To prove the second statement, we check that the spectral sequence is invariant under the Reidemeister moves. Suppose L and $\tilde{L}$ are two diagrams related by the i-th Reidemeister move. In [11], Lee defines maps $\rho_i' : C_{Lee}(L) \to C_{Lee}(\tilde{L})$ which induce isomorphisms on homology. In Sect. 6, we show that these maps induce isomorphisms on E^1 terms of spectral sequences, thus completing the proof of the theorem.

2.4 Calculation of H_{Lee}

Lee's second major result is that the homology group $H_{Lee}(L)$ is surprisingly simple. To show this, she introduces a new basis $\{\mathbf{a}, \mathbf{b}\}$ for V, where $\mathbf{a} =$

$\mathbf{v}_- + \mathbf{v}_+$ and $\mathbf{b} = \mathbf{v}_- - \mathbf{v}_+$. With respect to this new basis, the maps m' and Δ' are given by

$$\begin{aligned} m'(\mathbf{a} \otimes \mathbf{a}) &= 2\mathbf{a} & \Delta'(\mathbf{a}) &= \mathbf{a} \otimes \mathbf{a} \\ m'(\mathbf{a} \otimes \mathbf{b}) &= m'(\mathbf{b} \otimes \mathbf{a}) = 0 & \Delta'(\mathbf{b}) &= \mathbf{b} \otimes \mathbf{b} \\ m'(\mathbf{b} \otimes \mathbf{b}) &= -2\mathbf{b} \end{aligned}$$

and the maps ϵ' and ι' are given by

$$\epsilon'(\mathbf{a}) = \epsilon'(\mathbf{b}) = 1 \qquad \iota'(1) = (\mathbf{a} - \mathbf{b})/2$$

Using this basis, she proves

Theorem 2.2 (Theorem 5.1 of [11]) *$H_{Lee}(L)$ has rank 2^n, where n is the number of components of L.*

Indeed, Lee exhibits an explicit bijection between the set of possible orientations for L and a set of generators of $H_{Lee}(L)$, under which the generator corresponding to an orientation o is a single state in the new basis. This state lies in the vertex of the cube of resolutions determined by taking the oriented resolution of L with respect to the orientation o. More precisely, given an orientation o of L, let D_o be the corresponding oriented resolution, and let S be the state obtained by labeling the circles in D_o with $\mathbf{a}$'s and $\mathbf{b}$'s according to the following rule. To each circle C we assign a mod 2 invariant, which is the mod 2 number of circles in D_o which separate it from infinity. (In other words, draw a ray in the plane from C to infinity, and take the number of other times it intersects the other circles, mod 2.) To this number, we add 1 if C has the standard (counterclockwise) orientation, and 0 if it does not. Label C by $\mathbf{a}$ if the resulting invariant is 0, and by $\mathbf{b}$ if it is 1.

We denote the state described above by $\mathfrak{s}_o$, and refer to it as the *canonical generator* associated to o. The name is justified by the following result, whose proof is given in Sect. 6.

Proposition 2.3 *Suppose L and $\tilde{L}$ are related by the i-th Reidemeister move. Then an orientation o on L induces an orientation $\tilde{o}$ on $\tilde{L}$, and $\rho'_{i*}([\mathfrak{s}_o])$ is a nonzero multiple of $[\mathfrak{s}_{\tilde{o}}]$.*

Remark Note that orientations on L play two distinct roles in the discussion above. We must fix an orientation o on L to determine the exact values of the q and homological gradings. To get the generators of $H_{Lee}(L)$, however, we consider all possible orientations on L, regardless of which orientation we fixed. The two are related in a small way: if o is the orientation on L used to determine the gradings, then an easy calculation shows that the homological grading of $\mathfrak{s}_o$ is equal to 0. (In general, the homological gradings of all

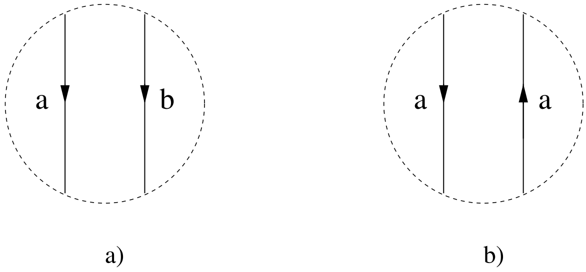

Fig. 2 Local behavior of the state $\mathfrak{s}_o$

the other generators are determined by the linking numbers of the various components of L.)

We end this section with an elementary but important observation.

Lemma 2.4 (Coherent orientations) *Suppose there is a region in the state diagram for $\mathfrak{s}_o$ containing exactly two segments, as shown in Fig. 2. Then either the orientations of the two are the same and the labels are different (like part a of the figure) or the orientations are different and the labels are the same (like part b).*

Proof We consider three possible cases: either the two segments belong to the same circle in D_o, or they belong to two circles, one of which is contained inside the other, or they belong to two circles, neither of which is contained inside the other. In each case, it is easy to verify that the claim holds. □

Corollary 2.5 *If two circles in the state diagram for $\mathfrak{s}_o$ share a crossing, they have different labels.*

3 Definition and basic properties of the invariant

Let K be a knot in S^3. By Theorems 2.1 and 2.2, we know that there is a spectral sequence associated to K which converges to $\mathbf{Q} \oplus \mathbf{Q}$. This spectral sequence is a relatively complicated object, but we can extract some simpler invariants of K from it. Let $s_{\max}$ and $s_{\min}$ (with $s_{\max} \geq s_{\min}$) be the q-gradings of the two surviving copies of $\mathbf{Q}$ which remain in the E^∞ term of the spectral sequence. Like all q-gradings for a knot, $s_{\max}$ and $s_{\min}$ are odd integers. Since the isomorphism type of the spectral sequence is an invariant of K, $s_{\max}$ and $s_{\min}$ are invariants as well.

Before making this definition formal, we digress to establish some terminology related to filtrations. Suppose C is a chain complex. A *finite length filtration* of C is a sequence of subcomplexes

$$C = \mathcal{F}^n C \supset \mathcal{F}^{n+1} C \supset \mathcal{F}^{n+2} C \supset \cdots \supset \mathcal{F}^m C = \{0\}.$$

To such a filtration, we associate a *grading* defined as follows: $x \in C$ has grading i if and only if $x \in \mathcal{F}^i C$ but $x \notin \mathcal{F}^{i+1} C$. If $f : C \to C'$ is a map between two filtered chain complexes, we say that f *respects the filtration* if $f(C_i) \subset C'_i$. More generally, we say that f is a *filtered map of degree* k if $f(C_i) \subset C'_{i+k}$.

A filtration $\{\mathcal{F}^i C\}$ of C induces a filtration

$$H_*(C) = \mathcal{F}^n H_*(C) \supset \mathcal{F}^{n+1} H_*(C) \supset \mathcal{F}^{n+2} H_*(C) \supset \cdots \supset \mathcal{F}^m H_*(C) = \{0\}$$

of $H_*(C)$ defined as follows: a class $[x] \in H_*(C)$ is in $\mathcal{F}^i H_*(C)$ if and only if has a representative which is an element of $\mathcal{F}^i C$. If $f : C \to C'$ is a filtered chain map of degree k, then it is easy to see that the induced map $f_* : H_*(C) \to H_*(C')$ is also filtered of degree k.

A finite length filtration$\{\mathcal{F}^i C\}$ on C induces a spectral sequence which converges to the associated graded group of the induced filtration on $H_*(C)$. In other words, the group which survives at filtration grading i in the spectral sequence is naturally identified with $\mathcal{F}^i H_*(C)/\mathcal{F}^{i+1} H_*(C)$.

Let us denote by s the grading on $H_{Lee}(K)$ induced by the q-grading on $C_{Lee}(K)$. Then the informal definition above is equivalent to

Definition 3.1

$$s_{\min}(K) = \min\{s(x) \mid x \in H_{Lee}(K), x \neq 0\}$$
$$s_{\max}(K) = \max\{s(x) \mid x \in H_{Lee}(K), x \neq 0\}$$

Since H_{Kh} of the unknot U has rank two and is supported in q-gradings ± 1, we have $s_{\max}(U) = 1$, $s_{\min}(U) = -1$.

If we wanted to, we could have defined $s_{\max}$ and $s_{\min}$ purely in terms of the filtration on $H_{Lee}(K)$, and avoided any mention of the spectral sequence. Indeed, the fact that $s_{\min}$ and $s_{\max}$ are knot invariants follows directly from the next proposition, whose proof may be found in Sect. 6.

Proposition 3.2 *The maps ρ'_{i*} and $(\rho'_{i_*})^{-1}$ both respect the induced filtration on H_{Lee}.*

In practice, however, s is most easily computed using Lee's spectral sequence, so we have chosen to retain this fact in the definition.

3.1 The invariant s

Our first task in this section is to prove

Proposition 3.3

$$s_{\max}(K) = s_{\min}(K) + 2$$

which justifies

Definition 3.4

$$s(K) = s_{\max}(K) - 1 = s_{\min}(K) + 1$$

Since $s_{\max}$ and $s_{\min}$ are odd, $s(K)$ is always an even integer.

Before proving the proposition, we need some preliminary results.

Lemma 3.5 *Let n be the number of components of L. There is a direct sum decomposition $C_{Lee}(L) \cong C^o_{Lee}(L) \oplus C^e_{Lee}(L)$, where $C^o_{Lee}(L)$ is generated by all states with q-grading congruent to $2+n$ mod 4, and $C^e_{Lee}(L)$ is generated by all states with q-grading congruent to n mod 4. If o is an orientation on L and $\overline{o}$ is the reverse orientation, then $\mathfrak{s}_o + \mathfrak{s}_{\overline{o}}$ is contained in one of the two summands, and $\mathfrak{s}_o - \mathfrak{s}_{\overline{o}}$ is contained in the other.*

Proof Following Lee [11], we write

$$m' = m + \Phi_m$$

$$\Delta' = \Delta + \Phi_\Delta$$

where m and Δ preserve the q-grading and Φ_m and Φ_Δ raise it by 4. This proves the first statement.

For the second, let $\iota : C_{Lee}(L) \to C_{Lee}(L)$ be the map which acts by the identity on C^e_{Lee} and by multiplication by -1 on C^o_{Lee}. We claim that $\iota(\mathfrak{s}_o) = \pm\mathfrak{s}_{\overline{o}}$. To see this, we define a new grading on V with respect to which $\mathbf{v}_-$ has grading 0 and $\mathbf{v}_+$ has grading 2. Let $i : V \to V$ be given by $i(\mathbf{v}_-) = \mathbf{v}_-$, $i(\mathbf{v}_+) = -\mathbf{v}_+$, so that $i(\mathbf{a}) = \mathbf{b}$ and $i(\mathbf{b}) = \mathbf{a}$. Then the induced map $i^{\otimes n} : V^{\otimes n} \to V^{\otimes n}$ acts as the identity on elements whose new grading is congruent to 0 mod 4 and as multiplication by -1 on elements whose new grading is congruent to 2 mod 4. The new grading differs from the q-grading on D_o by an overall shift, so

$$\iota(\mathfrak{s}_o) = \pm i^{\otimes n}(\mathfrak{s}_o) = \pm\mathfrak{s}_{\overline{o}}$$

It follows that $\mathfrak{s}_o + \iota(\mathfrak{s}_o) = \mathfrak{s}_o \pm \mathfrak{s}_{\overline{o}}$ is contained in one summand, while $\mathfrak{s}_o - \iota(\mathfrak{s}_o) = \mathfrak{s}_o \mp \mathfrak{s}_{\overline{o}}$ is contained in the other. □

Corollary 3.6

$$s(\mathfrak{s}_o) = s(\mathfrak{s}_{\overline{o}}) = s_{\min}(K)$$

Corollary 3.7 $s_{\max}(K) > s_{\min}(K)$.

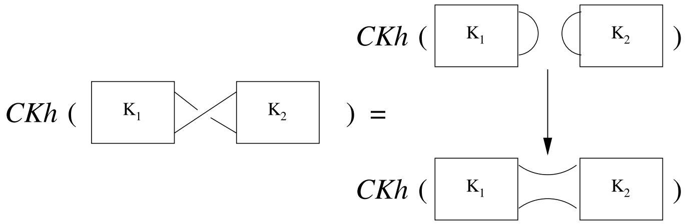

Fig. 3 A short exact sequence for $C_{Lee}(K_1\#K_2)$

Proof Since $C_{Lee}(K)$ decomposes as a direct sum, its affiliated spectral sequence decomposes too. The homology of each summand is $\mathbf{Q}$, so each must account for one of the surviving terms in the spectral sequence. The two summands are supported in different q-gradings, so the surviving terms must have different q-gradings as well. $\square$

Lemma 3.8 *Let K_1 and K_2 be oriented knots. Then there is a short exact sequence*

$$0 \to H_{Lee}(K_1\#K_2) \xrightarrow{p_*} H_{Lee}(K_1) \otimes H_{Lee}(K_2) \xrightarrow{\partial} H_{Lee}(K_1\#K_2^r) \to 0$$

where K_2^r denotes the reverse of K_2. The maps p_* and ∂ are filtered of q-degree -1.

Proof Consider the diagram for $K_1\#K_2$ shown in Fig. 3. From it, we get a short exact sequence

$$0 \to C_{Lee}(K_1\#K_2^r)\{2\} \xrightarrow{i} C_{Lee}(K_1\#K_2) \xrightarrow{p} C_{Lee}(K_1 \amalg K_2)\{1\} \to 0$$

where we use the notation $C\{a\}$ to indicate the complex C with the q-grading shifted up by a. With this convention, i and p are filtration preserving maps. Since $H_{Lee}(K_1\#K_2)$ and $H_{Lee}(K_1\#K_2^r)$ have rank two and $H_{Lee}(K_1 \amalg K_2) \cong H_{Lee}(K_1) \otimes H_{Lee}(K_2)$ has rank four, i_* must be the zero map. Thus the long exact sequence in homology splits to give the short exact sequence of the statement. Finally, we can remove the shifts in q-degree at the cost of making p_* and ∂ filtered maps of degree -1. $\square$

Proof of Proposition 3.3 Consider the exact sequence of the previous lemma with $K_1 = K$ and K_2 the unknot. Denote the canonical generators of K by $\mathfrak{s}_a$ and $\mathfrak{s}_b$, according to their label near the connected sum point, and the canonical generators of U by $\mathbf{a}$ and $\mathbf{b}$. Without loss of generality, we may assume that $s(\mathfrak{s}_a - \mathfrak{s}_b) = s_{\max}(K)$. From Fig. 3, we see that $\partial((\mathfrak{s}_a - \mathfrak{s}_b) \otimes \mathbf{a}) = \mathfrak{s}_a$.

Since ∂ is a filtered map of degree -1, we conclude that

$$s((\mathfrak{s}_a - \mathfrak{s}_b) \otimes \mathbf{a}) \leq s(\mathfrak{s}_a) + 1$$
$$s_{\max}(K) - 1 \leq s_{\min}(K) + 1$$

Since we already know that $s_{\max}(K) \neq s_{\min}(K)$, this gives the desired result. □

3.2 Properties of s

We check that s behaves nicely with respect to orientation reversal, mirror image and connected sum.

Lemma 3.9 *Let K be an oriented knot, and let K^r denote the same knot with the reversed orientation. Then $s(K) = s(K^r)$.*

Proof As observed in Sect. 2, C_{Kh} and C_{Lee} are invariant under the operation of global orientation reversal. □

Proposition 3.10 *Let $\overline{K}$ be the mirror image of K. Then we have*

$$s_{\max}(\overline{K}) = -s_{\min}(K)$$
$$s_{\min}(\overline{K}) = -s_{\max}(K)$$
$$s(\overline{K}) = -s(K)$$

Proof Suppose that C is a filtered complex with filtration

$$C = \mathcal{F}^n C \supset \mathcal{F}^{n+1} C \supset \mathcal{F}^{n+2} C \supset \cdots \supset \mathcal{F}^m C = \{0\}.$$

Then the dual complex C^* has a dual filtration

$$C^* = \mathcal{F}^{-m} C^* \supset \cdots \supset \mathcal{F}^{-n+2} C^* \supset \mathcal{F}^{-n+1} C^* \supset \mathcal{F}^{-n} C^* = \{0\}$$

where $\mathcal{F}^{-i} C^* = \{x \in C^* \mid \langle x, y \rangle = 0, \forall y \in \mathcal{F}^i C\}$.

To prove the proposition, we observe that the filtered complex $C_{Lee}(\overline{K})$ is isomorphic to $(C_{Lee}(K))^*$. Indeed, it is easy to see from Eq. 3 that there is an isomorphism

$$r : (V, m', \Delta') \to (V^*, \Delta'^*, m'^*)$$

which sends $\mathbf{v}_\pm$ to $\mathbf{v}^*_\mp$. Then if $\mathbf{v}$ is a state of the diagram $\overline{K}$, we define $R(\mathbf{v})$ to be state of K obtained by applying r all the labels of $\mathbf{v}$. It is straightforward to check that the map $R : C_{Lee}(\overline{K}) \to (C_{Lee}(K))^*$ is the desired isomorphism. (Compare with Sect. 7.3 of [8], where it is shown that $C_{Kh}(\overline{K}) \cong (C_{Kh}(K))^*$.)

We now appeal to the following general result, whose proof is left to the reader:

Lemma 3.11 *If C and C' are dual filtered complexes over a field, then their associated spectral sequences E^i and E'^i are dual, in the sense that $E^i \cong (E'^i)^*$.*

Thus if the two surviving generators in the spectral sequence for $H_{Lee}(K)$ have filtration gradings $s_{\min}$ and $s_{\max}$, the surviving generators in the spectral sequence for $H_{Lee}(\overline{K})$ will have gradings $-s_{\max}$ and $-s_{\min}$. □

Proposition 3.12

$$s(K_1\#K_2) = s(K_1) + s(K_2)$$

Proof We use the short exact sequence of Lemma 3.8. Denote the canonical generators of K_i by $\mathfrak{s}_a^i$ and $\mathfrak{s}_b^i$, according to their label near the connected sum point. It is not difficult to see that $H_{Lee}(K_1\#K_2)$ has a canonical generator $\mathfrak{s}_o$ which maps to $\mathfrak{s}_a \otimes \mathfrak{s}_b$ under p_*. Thus

$$s(\mathfrak{s}_o) - 1 \leq s(\mathfrak{s}_a^1 \otimes \mathfrak{s}_b^2)$$
$$s_{\min}(K_1\#K_2) - 1 \leq s_{\min}(K_1) + s_{\min}(K_2)$$

Applying the same argument to $\overline{K}_1$ and $\overline{K}_2$, and using the fact that $s_{\min}(K) = -s_{\max}(K)$, we see that

$$s_{\max}(K_1\#K_2) + 1 \geq s_{\max}(K_1) + s_{\max}(K_2)$$

Substituting $s_{\max}(K) = s_{\min}(K) + 2$, we get

$$s_{\min}(K_1\#K_2) + 3 \geq s_{\min}(K_1) + s_{\min}(K_2) + 4$$

which, when combined with the previous equation, implies that

$$s_{\min}(K_1\#K_2) = s_{\min}(K_1) + s_{\min}(K_1) + 1$$
$$s_{\max}(K_1\#K_2) = s_{\max}(K_1) + s_{\max}(K_1) - 1$$

This proves the claim. □

4 Behavior under cobordisms

Let L_0 and L_1 be two links in $\mathbf{R}^3$. A cobordism from L_0 to L_1 is a smooth, oriented, compact, properly embedded surface $S \subset \mathbf{R}^3 \times [0, 1]$ with

$S \cap (\mathbf{R}^3 \times \{i\}) = L_i$. In this section, we define and study a map ϕ'_S : $H_{Lee}(L_0) \to H_{Lee}(L_1)$ induced by such a cobordism. Our construction follows Sect. 6.3 of [8], where Khovanov describes a similar map for the homology theory H_{Kh}.

4.1 Orientations

We begin by fixing some conventions about orientations. Let S be cobordism from L_0 to L_1, and let o_S be an orientation on S. We say that orientations o_0 on L_0 and o_1 on L_1 are *compatible* with o_S if o_1 is the usual orientation induced by o_S and o_0 is the *reverse* of the induced orientation. More generally, we say that o_1 and o_2 are compatible if there is some orientation on S with which they are both compatible. (Thus if S is the product cobordism from L_0 to L_0, o_0 is compatible with itself.) Given such an S, we seek to define $\phi'_S : H_{Lee}(L_0) \to H_{Lee}(L_1)$, where L_0 and L_1 are equipped with orientations compatible with the given orientation on S.

4.2 Elementary cobordisms

Following Khovanov, we decompose S into a series of elementary cobordisms, each represented by a single move from one planar diagram to another. (See [4] for a more detailed treatment of this material.) For $i \in [0, 1]$, let $L_i = S \cap (\mathbf{R}^3 \times \{i\})$. After a small isotopy of S, we can assume that L_i is a link in $\mathbf{R}^3$ for all but finitely many values of i. Next, we fix a projection $p : \mathbf{R}^3 \to \mathbf{R}^2$. After a further small isotopy of S, we can assume that p defines a regular projection of L_i for all but finitely many values of i, and that this set of special values is disjoint from the first set where L failed to be a link. The isotopy type of the oriented planar diagram $p(L_i)$ remains constant except when L passes through one of the two types of special values, where it changes by some well-defined local move. Each of these moves corresponds to an elementary cobordism, so we can write the whole cobordism S as a composition of elementary cobordisms.

The necessary moves may be subdivided into two types: Reidemeister moves and Morse moves. There is one Reidemeister-type move for each of the ordinary Reidemeister moves, as well as one for each of their inverses. These moves do not change the topology of the surface S_i. The Morse moves correspond to the addition of a 0, 1 or 2-handle to S_i. They are illustrated in Fig. 4.

4.3 Induced maps

Given a cobordism S from L_0 to L_1, we want to assign to it an induced map $\phi'_S : H_{Lee}(L_0) \to H_{Lee}(L_1)$. ϕ'_S should be a filtered map of degree $\chi(S)$. In

Fig. 4 Local pictures for Morse moves

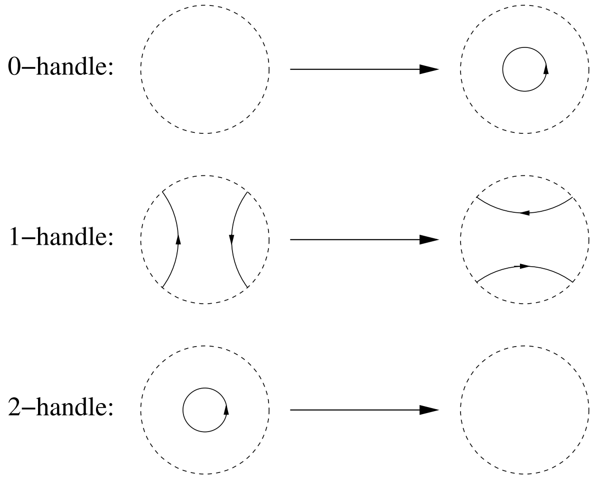

addition, we would like this assignment to be functorial, in the sense that if S is the composition of two cobordisms S_1 and S_2, ϕ_S is the composition of ϕ_{S_1} and ϕ_{S_2}. Thus it suffices to consider the case when S is an elementary cobordism.

First, suppose that S is an elementary cobordism corresponding to the i-th Reidemeister move or its inverse. Recall from Sect. 2.3 that there is a map $\rho'_i : C_{Lee}(L_0) \to C_{Lee}(L_1)$ which induces an isomorphism on homology. We define ϕ'_S to be ρ'_{i*} or its inverse. By Proposition 3.2, this is a filtered map of degree 0.

Now suppose S is an elementary cobordism associated to a Morse move. To define ϕ'_S in this case, recall that $C_{Lee}(L_0)$ is generated by elements of the form $\mathbf{v} \in \mathcal{A}'(D_{0,v})$, where $D_{0,v}$ is a complete resolution of the diagram L_0. $D_{0,v}$ naturally determines a complete resolution $D_{1,v}$ of L_1, and the cobordism S induces a cobordism S_v from $D_{0,v}$ to $D_{1,v}$. Let $\phi : C_{Lee}(L_0) \to C_{Lee}(L_1)$ be the map given at the chain level by $\phi(\mathbf{v}) = \mathcal{A}'(S_v)(\mathbf{v})$ for $\mathbf{v} \in D_{0,v}$. We define ϕ'_S to be the induced map on homology. It is easy to see that ϕ'_S is a filtered map of degree 1 for a 0- or 2-handle addition and degree -1 for a 1-handle.

In general, given a cobordism S, we decompose it as a union of elementary cobordisms: $S = S_1 \cup S_2 \cup \cdots \cup S_k$ and define the induced morphism $\phi'_S : H_{Lee}(L_0) \to H_{Lee}(L_1)$ to be the composition $\phi'_{S_k} \circ \cdots \circ \phi'_{S_1}$. This is a filtered map of degree $\chi(S)$. Although we will not need this fact here, it follows from [3] that up to sign, the map ϕ'_S does not depend on the particular decomposition of S we chose.

4.4 Canonical generators

The map ϕ'_S behaves nicely with respect to canonical generators. More precisely, we have

Proposition 4.1 *Let S be a cobordism from L_0 to L_1 with no closed components. If o is an orientation on L_0, then*

$$\phi'_S([\mathfrak{s}_o]) = \sum_{o_I} a_I [\mathfrak{s}_{o_I}] \tag{4}$$

where the sum runs over all orientations on L_1 compatible with o and each coefficient a_I is nonzero.

Remark The hypothesis on the absence of closed components is clearly necessary. For example, if we take S to be the union of a product cobordism and a trivially embedded sphere, ϕ'_S is the zero map.

Proof We induct on the number of elementary cobordisms in the composition. For the base case, we must check that Eq. 4 holds when S is an elementary cobordism. For the Reidmeister-type moves, this follows directly from Proposition 2.3. Below, we check that it holds at the chain level for each of the Morse-type moves.

0-*handle move* In this case, we have $\phi(\mathfrak{s}_o) = \mathfrak{s}_o \otimes \frac{1}{2}(\mathbf{a} - \mathbf{b})$, where the second factor in the tensor product refers to the labels on the newly created circle. There are two orientations on L_1 compatible with o, and their corresponding canonical generators are $\mathfrak{s}_o \otimes \mathbf{a}$ and $\mathfrak{s}_o \otimes \mathbf{b}$. Thus Eq. 4 holds.

1-*handle move* Here, we consider several subcases. First, suppose that under S, one component of L_0 splits to form two components of L_1. Then there is a unique orientation o_1 on L_1 compatible with o, and the two strands in the neighborhood where the Morse move takes place must be pointing in opposite directions, as in Fig. 4. We consider the state $\mathfrak{s}_o$. By Lemma 2.4, the two strands in the neighborhood of the move must have the same label. Since

$$\begin{aligned} m'(\mathbf{a} \otimes \mathbf{a}) &= 2\mathbf{a} & \qquad \Delta'(\mathbf{a}) &= \mathbf{a} \otimes \mathbf{a} \\ m'(\mathbf{b} \otimes \mathbf{b}) &= -2\mathbf{b} & \qquad \Delta'(\mathbf{b}) &= \mathbf{b} \otimes \mathbf{b} \end{aligned}$$

the state $\phi(\mathfrak{s}_o)$ is some nonzero multiple of the state obtained by starting with $\mathfrak{s}_o$ and either merging or splitting the component(s) in the neighborhood of the move, keeping all the labels the same. We claim this is precisely the state $\mathfrak{s}_{o_1}$. Indeed, at each crossing the oriented resolution with respect to o is the same as the oriented resolution with respect to o_1. Moreover, the orientations on each component of the resulting diagram are clearly the same, as are the number of circles separating a given circle from infinity. (Choose a path to infinity avoiding the neighborhood in which the move takes place.) Thus $\phi(\mathfrak{s}_o) = a\mathfrak{s}_{o_1}$, where a is either 1 or ± 2.

Next, suppose that S merges two components of L_0 into a single component of L_1. Here, there are two possibilities. First, suppose that the two strands in the neighborhood of the move are oriented oppositely under o. Then there is a unique compatible orientation o_1 on L_1, and the argument proceeds exactly as it did in the previous case. The other possibility is that the two strands in the neighborhood of the move have parallel orientations. In this case, o does not extend to an orientation on S, so there is no compatible orientation on L_1. Applying Lemma 2.4, we see that the two strands must have opposite labels in $\mathfrak{s}_o$. Since $m'(\mathbf{a} \otimes \mathbf{b}) = 0$, we have $\phi(\mathfrak{s}_o) = 0$, and both sides of Eq. 4 vanish.

2-*handle move* In this case, there is a unique orientation o_1 on L_1 compatible with o. Since $\epsilon'(\mathbf{a}) = \epsilon'(\mathbf{b}) = 1$, $\phi(\mathfrak{s}_o) = \mathfrak{s}_{o_1}$. Thus Eq. 4 holds here as well.

In general, if S is a composition of several elementary cobordisms, we decompose it into the composition of a cobordism S_1 from L_0 to $L_{1/2}$ and a cobordism S_2 from $L_{1/2}$ to L_1. By the induction hypothesis, Eq. 4 holds for S_1 and S_2, so

$$\begin{aligned} \phi'_S([\mathfrak{s}_o]) &= \phi'_{S_2}\left(\sum_{o_I} a_I[\mathfrak{s}_{o_I}]\right) \\ &= \sum_{o_I} a_I \sum_{o_J} b_J[\mathfrak{s}_{o_J}] \\ &= \sum_{(o_I,o_J)} a_I b_J[\mathfrak{s}_{o_J}] \end{aligned}$$

where the sum runs over pairs (o_I, o_J) such that o_I is an orientation on $L_{1/2}$ compatible with o and o_J is an orientation on L_1 compatible with o_I. We claim that such pairs are in bijective correspondence with the set of orientations $\{o_K\}$ on L_1 compatible with o. Indeed, if (o_I, o_J) is such a pair, it is easy to see that o_J is compatible with o. Conversely, given o_K on L_1 compatible with o, the fact that S has no closed components implies that there is a unique orientation o_S on S compatible with o and o_K. Restricting o_S to $L_{1/2}$, we see that there is a unique compatible pair (o_I, o_K), and the claim is proved. Finally, since all the a_I's and b_J's are nonzero, it follows that S satisfies Eq. 4. □

Remark Although all our calculations involving the Morse moves were made at the chain level, the proposition itself only holds at the level of homology.

This is because the chain level analog of Proposition 2.3 does not hold, as can be seen from its proof in Sect. 6.

Corollary 4.2 *If S is a connected cobordism between knots K_0 and K_1, then ϕ'_S is an isomorphism.*

Proof Fix an orientation o on K_0. Since S is connected, there is a unique compatible orientation o_1 on K_1. Then $\{\mathfrak{s}_o, \mathfrak{s}_{\overline{o}}\}$ is a basis for $H_{Lee}(K_0)$. Its image under ϕ'_S is $\{k_1\mathfrak{s}_{o_1}, k_2\mathfrak{s}_{\overline{o}_1}\}$ $(k_1, k_2 \neq 0)$, which is a basis for $H_{Lee}(K_1)$. □

4.5 The slice genus

We can now prove the first two theorems from the introduction.

Proof of Theorem 1 Suppose $K \subset S^3$ bounds an oriented surface of genus g in B^4. Then there is an orientable connected cobordism of Euler characteristic $-2g$ between K and the unknot U in $\mathbf{R}^3 \times [0, 1]$. Let $x \in H_{Lee}(K) - \{0\}$ be a class for which $s(x)$ is maximal. Then $\phi'_S(x)$ is a nonzero element of $H_{Lee}(U)$. Now ϕ'_S is a filtered map with filtered degree $-2g$, so

$$s(\phi'_S(x)) \geq s(x) - 2g$$

On the other hand, $s_{\max}(U) = 1$, so

$$s(\phi'_S(x)) \leq 1$$

It follows that $s(x) \leq 2g + 1$, so $s_{\max}(K) \leq 2g + 1$ and $s(K) \leq 2g$. To show that $s(K) \geq -2g$, we apply the same argument to $\overline{K}$ (which bounds a surface $\overline{S}$ of genus g) and use the fact that $s(\overline{K}) = -s(K)$. □

Proof of Theorem 2 If K_1 and K_2 are concordant, then $K_1\#\overline{K_2^r}$ is slice. Then

$$0 = s(K_1\#\overline{K_2^r}) = s(K_1) - s(K_2)$$

so s gives a well-defined map from $\mathrm{Conc}(S^3)$ to $\mathbf{Z}$. That this map is a homomorphism is immediate from Propositions 3.10 and 3.12. □

Corollary 4.3 *Suppose K_+ and K_- are knots that differ by a single crossing change—from a positive crossing in K_+ to a negative one in K_-. Then*

$$s(K_-) \leq s(K_+) \leq s(K_-) + 2$$

Proof In [13], Livingston shows that this skein inequality holds for any knot invariant satisfying the properties of Theorems 1 and 2. □

5 Computations and relations with other invariants

Although the invariant $s(K)$ is algorithmically computable from a diagram of K, it is impossible to compute by hand for all but the smallest knots. In this section, we describe some techniques which enable us to efficiently compute s.

5.1 Using H_{Kh}

For many knots, it is a simple matter to compute $s(K)$ from the ordinary Khovanov homology $H_{Kh}(K)$. Although $H_{Kh}(K)$ is also hard to compute by hand, there are already a number of computer programs available for this purpose, including Bar-Natan's pioneering program [1] and a more recent, faster program written by Shumakovitch [22].

In [1], Bar-Natan made the following observation, based on his computations of H_{Kh} for knots with 10 and fewer crossings.

Conjecture (Bar-Natan) *The graded Poincare polynomial $P_{Kh}(K)$ of $H_{Kh}(K)$ has the form*

$$P_{Kh}(K) = q^{s(K)}(q + q^{-1}) + (1 + tq^4)Q_{Kh}(K)$$

where $Q_{Kh}(K)$ is a polynomial with all positive coefficients.

In [11], Lee observed that this conjecture is related to the convergence of the spectral sequence. To be precise, recall that if (E^i, d_i) is the ith term of spectral sequence, the differential d_i raises the q-grading by i.

Lemma 5.1 $d_i = 0$ *unless* i *is divisible by* 4.

Proof As we remarked in the proof of Lemma 3.5, Lee's differential can be decomposed as $d' = d + \Phi$, where d preserves the q grading and Φ raises it by 4. It is not hard to see that this implies the statement of the lemma. □

Thus after d_0, the first non-vanishing differential in the spectral sequence is d_4. It is now easy to see that Bar-Natan's conjecture holds whenever Lee's spectral sequence converges after the E^4 term (*i.e.* $d_i = 0$ for $i > 4$), and that the invariant $s(K)$ is equal to the exponent $s(K)$ appearing in the conjecture.

To see how widely applicable this condition is, we introduce the notion of the homological *width* of a knot.

Definition 5.2 If K is a knot, let $\mu(K) = \{a/2 - b \mid q^a t^b$ is a monomial in $P_{Kh}(K)\}$. The width $W(K)$ is one more than the difference between the maximum and minimum elements of $\mu(K)$.

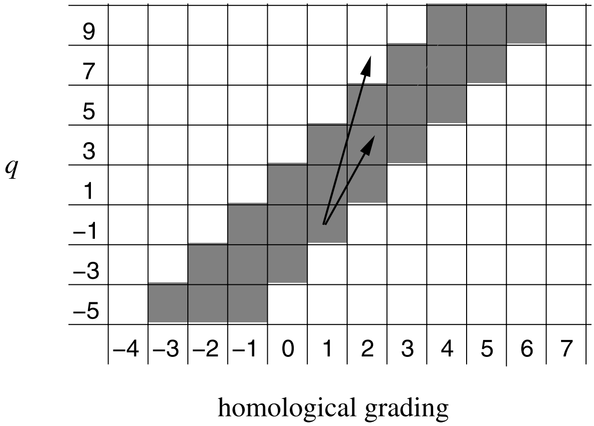

Fig. 5 The E_1 term of Lee's spectral sequence for a knot of width 3. $E_1 = H_{Kh}(K)$ is supported on the shaded diagonals. Possible differentials corresponding to d_4 (the *lower arrow*) and d_8 (the *upper one*) are shown

In other words, if we arrange $H_{Kh}(K)$ along diagonals for which the difference between the q-grading and twice the homological grading is constant, $W(K)$ is the number of diagonals spanned by the support of $H_{Kh}(K)$.

Proposition 5.3 *If $W(K) \leq 3$, then the spectral sequence for $H_{Lee}(K)$ converges after the E^4 term, and our $s(K)$ is the same as Bar-Natan's.*

Proof Suppose K has width ≤ 3. Then as illustrated in Fig. 5, the differential d_4 can (and typically will) be nonvanishing, but d_8 and all higher differentials must vanish for geometrical reasons. □

Theorem 3 follows from this fact, since Lee has shown [12] that if K is an alternating knot, then it has width two and Bar-Natan's s is equal to $\sigma(K)$.

The proposition also applies to many non-alternating knots. Indeed, using Shumakovitch's tables and a computer, it is straightforward to check that there are only four knots with 13 or fewer crossings whose width is greater than three. Inspecting H_{Kh} of these four exceptions, one sees that in each case, the spectral sequence must converge after the E^4 term. Thus for all knots with 13 or fewer crossings, the value of $s(K)$ agrees with the value of Bar-Natan's s tabulated in [1] and [22]. In Table 1 we list those knots of 11 crossings or fewer for which $s(K) \neq \sigma(K)$. There are 22 such knots, and $|s(K)| > |\sigma(K)|$ (and thus provides a better bound on the slice genus) for precisely half of them.

Knots with 10 or fewer crossings are labeled according to their numbering in Rolfsen, while those with 11 crossings use the *Knotscape* numbering. The values of the signature are taken from [2]. All of the knots in Table 1 have a homological width of 3, which raises the following question: if K has homological width 2 (*i.e.* is H-thin in the terminology of [9]), must $s(K) = \sigma(K)$?

Table 1 Knots with ≤ 1 crossings and $s \neq \sigma$

K	$s(K)$	$\sigma(K)$	K	$s(K)$	$\sigma(K)$	K	$s(K)$	$\sigma(K)$
9_{42}	0	2	11_{n9}	6	4	11_{n70}	2	4
10_{132}	−2	0	11_{n12}	2	0	11_{n77}	8	6
10_{136}	0	2	11_{n19}	−2	−4	11_{n79}	0	2
10_{139}	8	6	11_{n20}	0	−2	11_{n92}	0	−2
10_{145}	−4	−2	11_{n24}	0	2	11_{n96}	0	2
10_{152}	−8	−6	11_{n31}	4	2	11_{n138}	0	2
10_{154}	6	4	11_{n38}	0	2	11_{n183}	6	4
10_{161}	−6	−4						

5.2 Positive knots

If K is a positive knot, $s(K)$ can be computed directly from the definition. To see this, consider a canonical generator $\mathfrak{s}_o$ for a positive diagram of K. Since each crossing of K is positive, its oriented resolution is the 0-resolution. Thus the state $\mathfrak{s}_o$ lives in the extreme corner of the cube of resolutions: it has homological grading 0, and there are no generators in $C_{Lee}(K)$ with homological grading -1. It follows that the only class homologous to $\mathfrak{s}_o$ is $\mathfrak{s}_o$ itself, so

$$s_{\min}(K) = s([\mathfrak{s}_o]) = q(\mathfrak{s}_o)$$

To compute $q(\mathfrak{s}_o)$, we change back to the basis $\{\mathbf{v}_-, \mathbf{v}_+\}$. In the expansion of $\mathfrak{s}_o$ with respect to this basis, there is a unique state with minimal q-grading, namely, the state in which every circle of the oriented resolution is labeled with a $\mathbf{v}_-$. If the positive diagram of K has n crossings, and its oriented resolution has k circles, then

$$\begin{aligned} q(\mathfrak{s}_o) &= p(\mathfrak{s}_o) + \mathrm{gr}(\mathfrak{s}_o) + n_+ - n_- \\ &= -k + 0 + n - 0 \end{aligned}$$

so

$$s(K) = -k + n + 1$$

On the other hand, Seifert's algorithm gives a Seifert surface S for K with Euler characteristic $k - n$, so

$$2g(K) \leq 2g(S) = n - k + 1 = s(K) \leq 2g_*(K)$$

Since $g_*(K) \leq g(K)$, the inequalities above must all be equalities. This completes the proof of Theorem 4.

5.3 Comparison with τ

We end this section by commenting on the conjecture relating s and τ which was stated in the introduction. In addition to the fact that the two invariants share the properties of Theorems 1 through 4, there is a good deal of numerical evidence supporting the conjecture. Recently, a fair amount of work has been done on the problem of computing τ for knots with 10 and fewer crossings. Combining the results of [6, 13, 15, 16], and [17] with some unpublished computations of the author, it appears that the value of τ has been determined for all but two knots of 10 crossings and fewer. (The exceptions are 10_{141} and 10_{150}.) For all of these knots, $s = 2\tau$. The equality can also be checked on certain special classes of knots, such as the pretzel knots of [18]. If the conjecture were true, it would make many computations in knot Floer homology easier. (For example, with our current technology, it seems like quite a laborious project to compute τ for all 11-crossing non-alternating knots.) Even if it is not true, we hope that the remarkable similarity between the two theories will have an enlightening explanation.

6 Reidemeister moves

In this section, we prove the results involving Reidemeister moves which were stated in Sects. 2 and 3.

Proof of Theorem 2.1 The proof that the desired spectral sequence exists was sketched in Sect. 2. To prove its invariance, we use the following basic lemma, whose proof may be found in [14], Proposition 3.2.

Lemma 6.1 *Suppose $f: C \to C'$ is a map of filtered complexes which respects the filtrations. Then f induces maps of spectral sequences $f^n: E^n \to E'^n$, and if f^n is an isomorphism, f^m is an isomorphism for all $m \geq n$.*

In Sect. 4 of [11], Lee proves the invariance of H_{Lee} by checking its invariance under the three Reidemeister moves. For each move, she exhibits a chain map between the complexes associated to the link diagram before and after the move. To prove the theorem, it suffices to check that these maps respect the q-filtration, and that they induce isomorphisms on the E^1 terms. The latter claim is straightforward, since in each case the induced maps on the E^0 terms are identical to the maps used in Sect. 5 of [8] to prove invariance of H_{Kh}. Below, we sketch the proof of invariance for each move and explain why the maps in question respect the filtrations. For full details, we refer the reader to [8] and [11].

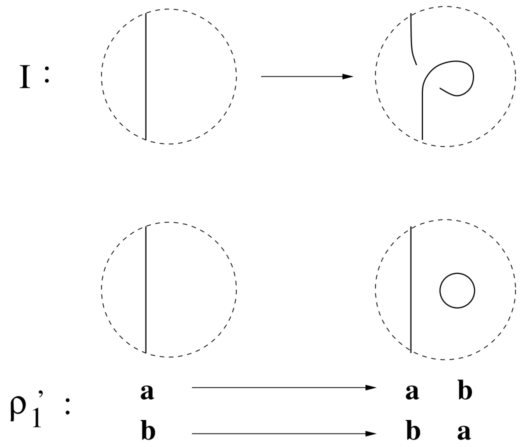

Fig. 6 The Reidemeister I move and the map ρ_1'

Reidemeister I move Let $\tilde{L}$ be the diagram L with an additional left-hand curl added in. Then $C_{Lee}(\tilde{L})$ can be decomposed as a direct sum $X_1 \oplus X_2$, where X_2 is acyclic and X_1 is isomorphic to $C_{Lee}(L)$ via the map $\rho_1' : C_{Lee}(L) \to X_1$ illustrated in Fig. 6. In terms of the basis $\{\mathbf{v}_\pm\}$, we have

$$\rho_1'(\mathbf{v}_-) = \mathbf{v}_- \otimes \mathbf{v}_- - \mathbf{v}_+ \otimes \mathbf{v}_+$$
$$\rho_1'(\mathbf{v}_+) = \mathbf{v}_+ \otimes \mathbf{v}_- - \mathbf{v}_- \otimes \mathbf{v}_+$$

The corresponding map ρ_1 in [8] is given by

$$\rho_1(\mathbf{v}_-) = \mathbf{v}_- \otimes \mathbf{v}_-$$
$$\rho_1(\mathbf{v}_+) = \mathbf{v}_+ \otimes \mathbf{v}_- - \mathbf{v}_- \otimes \mathbf{v}_+$$

so ρ_1' is filtration non-decreasing, and its induced map on E^0 terms is ρ_1.

Remark There is another version of the first Reidemeister move, corresponding to the addition of a right-hand curl. Although it is not difficult to define an appropriate map $\rho_{1'}'$ for this move directly, for the sake of brevity we adopt the solution of [1] and [11] and define it to be the composition of maps induced by an appropriate Reidemeister II move followed by a Reidemeister I move.

Reidemeister II move Let L and $\tilde{L}$ be as shown in Fig. 7. In this case, $C_{Lee}(\tilde{L})$ can be decomposed as a direct sum $X_1 \oplus X_2 \oplus X_3$, where X_2 and X_3 are acyclic and there is an isomorphism $\rho_2' : C_{Lee}(L) \to X_1$, which is given by

$$\rho_2'(z) = (-1)^{\mathrm{gr}(z)}(z + \iota(d_{01\to 11}'(z)))$$

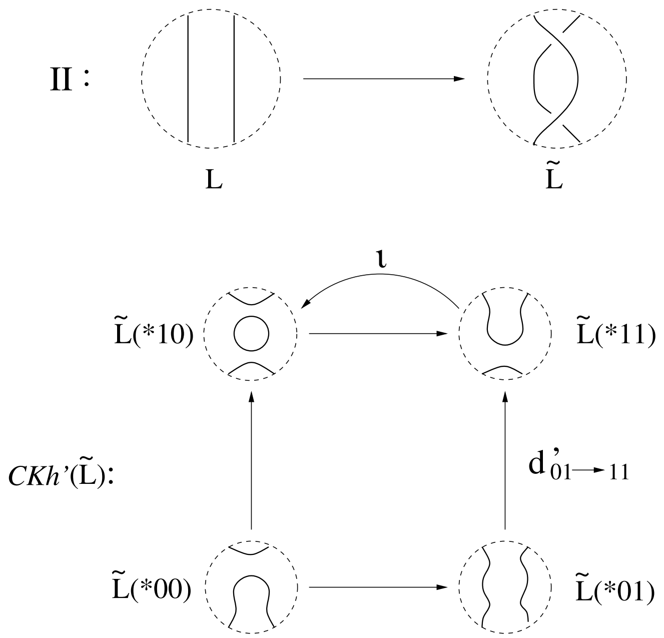

Fig. 7 The Reidemeister II move and the maps ι and $d'_{01\to 11}$

The maps ι and $d'_{01\to 11}$ are shown in the figure. The isomorphism ρ_2 in [8] has the same form, but with $d_{01\to 11}$ in place of $d'_{01\to 11}$. Since $d-d'$ is strictly filtration increasing, it follows that ρ'_2 is filtration non-decreasing, and its induced map on E^0 terms is ρ_2.

Reidemeister III move Let L and $\tilde{L}$ be as shown in Fig. 8. Then there are direct sum decompositions

$$C_{Lee}(L) \cong X_1 \oplus X_2 \oplus X_3$$
$$C_{Lee}(\tilde{L}) \cong \tilde{X}_1 \oplus \tilde{X}_2 \oplus \tilde{X}_3$$

where $X_2, X_3, \tilde{X}_2$, and $\tilde{X}_3$ are acyclic and there is an isomorphism $\rho'_3 : X_1 \to \tilde{X}_1$. To describe X_1 and $\tilde{X}_1$, we first define maps

$$\beta' : C_{Lee}(L(*100)) \to C_{Lee}(L(*010))$$
$$\tilde{\beta}' : C_{Lee}(\tilde{L}(*010)) \to C_{Lee}(\tilde{L}(*100))$$

by

$$\beta' = \iota \circ d'_{100\to 110}$$

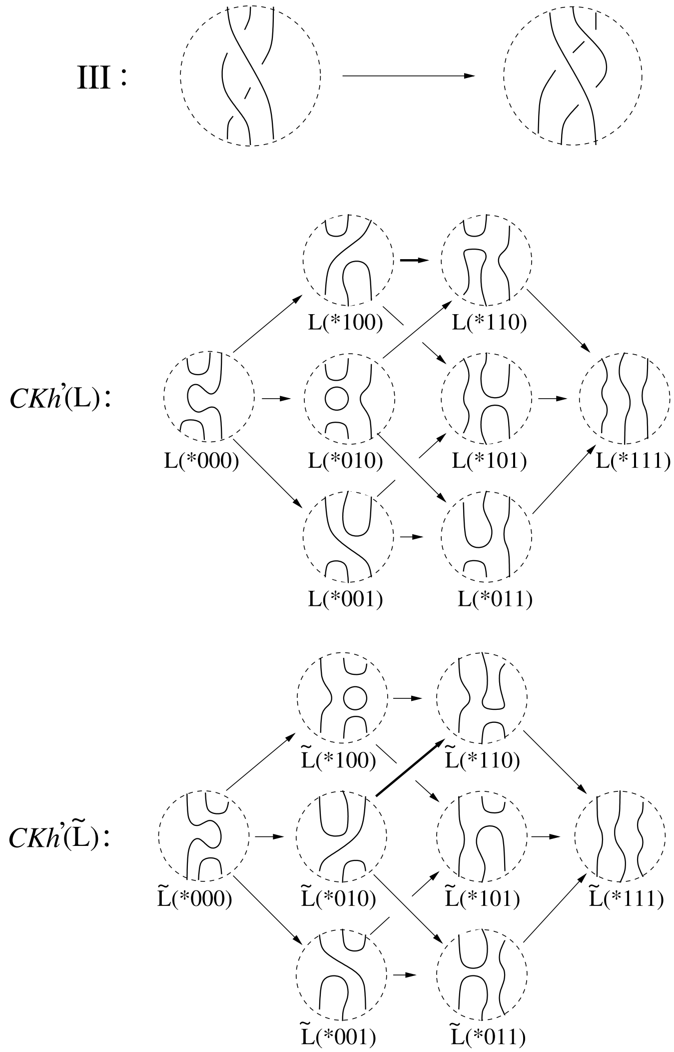

Fig. 8 The Reidemeister III move. The relevant components of the differentials ($d'_{100\to 110}$ and $d'_{010\to 110}$) are marked in *bold*

$$\tilde{\beta}' = \iota \circ d'_{010\to 110}$$

Then

$$X_1 = \{x + \beta'(x) + y \mid x \in C_{Lee}(L(*100)), y \in C'_{Kh}(L(*1))\}$$
$$\tilde{X}_1 = \{x + \tilde{\beta}'(x) + y \mid x \in C_{Lee}(\tilde{L}(*010)), y \in C_{Lee}(\tilde{L}(*1))\}$$

and

$$\rho_3'(x+\beta'(x)+y) = x+\tilde{\beta}'(x)+y.$$

The isomorphism ρ_3 in [8] is defined similarly, except that it uses d instead of d' to define maps β and β'. Since d' does not increase the q-grading, we clearly have $q(\beta'(x)) \geq q(x)$. From this, it follows that ρ_3' does not decrease the q-grading. Since $d-d'$ strictly increases the q-grading, the map induced on E^0 terms by ρ_3' is equal to ρ_3. To finish the proof, we apply Lemma 6.1 three times: first to the inclusions $X_1 \hookrightarrow C_{Lee}(L)$ and $\tilde{X}_1 \hookrightarrow C_{Lee}(\tilde{L})$, and then to the map ρ_3'. □

Proof of Proposition 2.3 We check the claim directly for each Reidemeister move:

Reidemeister I move In this case, it is easy to see that $\rho_1'(\mathfrak{s}_o) = \mathfrak{s}_{\tilde{o}}$.

Reidemeister II move Suppose that the two strands in L point in the same direction. Then by Lemma 2.4, they have different labels, so $d'_{01\to 11}(\mathfrak{s}_o) = 0$. The oriented resolution of $\tilde{L}$ is contained in $C_{Lee}(\tilde{L}(*01)) \cong C_{Lee}(L)$, so $\rho_2'(\mathfrak{s}_o) = (-1)^0(\mathfrak{s}_{\tilde{o}}) = \mathfrak{s}_{\tilde{o}}$.

Now suppose the two strands point in different directions, so that they have the same label. Let us assume for the moment that this label is **a**. Then we define $\mathfrak{s}_{\widetilde{ij}} \in H_{Lee}(\widetilde{L}(*ij))$ be the state which is identical to $\mathfrak{s}_o$ outside the area where the move takes place and has all components inside the area of the move labeled with an **a**. Then a direct computation shows that either

$$\begin{aligned}\rho_2'(\mathfrak{s}_o) &= \mathfrak{s}_{\widetilde{01}} + \frac{1}{2}(\mathfrak{s}_{\widetilde{10}} - \mathfrak{s}_{\tilde{o}}) \\ &= -\frac{1}{2}(\mathfrak{s}_{\tilde{o}} + d'(\mathfrak{s}_{\widetilde{00}}))\end{aligned}$$

if the two strands belong to the same component, or

$$\begin{aligned}\rho_2'(\mathfrak{s}_o) &= \mathfrak{s}_{\widetilde{01}} + (\mathfrak{s}_{\widetilde{10}} - \mathfrak{s}_{\tilde{o}}) \\ &= -(\mathfrak{s}_{\tilde{o}} + d'(\mathfrak{s}_{\widetilde{00}}))\end{aligned}$$

if they belong to different components. This proves the claim in the case where both strands are labeled with an **a**. We leave it to the reader to check that a similar argument applies when they are both labeled with a **b**.

Reidemeister III move Here there are three cases to consider. First, suppose that the two overlying strands in L are oriented as shown in Fig. 9a. Then $\mathfrak{s}_o \in C_{Lee}(L(*1))$, and it is easy to see that $\rho_3'(\mathfrak{s}_o) = \mathfrak{s}_{\tilde{o}}$.

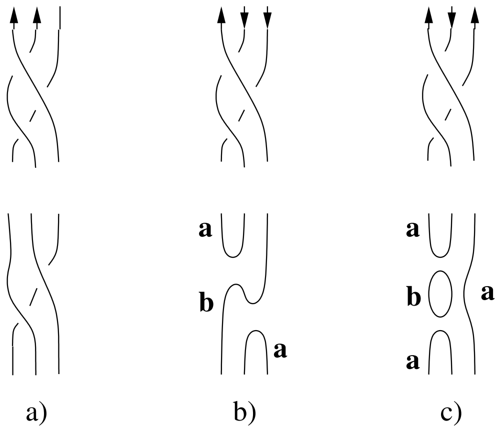

Fig. 9 Possible orientations for L and their respective canonical generators

Next, suppose that the three strands are oriented as shown in Fig. 9b. Then $\mathfrak{s}_o \in C_{Lee}(L(*100))$ and $\mathfrak{s}_{\tilde{o}} \in C_{Lee}(\tilde{L}(*010))$. Clearly $\beta'(\mathfrak{s}_o) = \tilde{\beta}'(\mathfrak{s}_{\tilde{o}}) = 0$, so $\mathfrak{s}_o \in X_1$ and $\mathfrak{s}_{\tilde{o}} \in \tilde{X}_1$. Again, it follows that $\rho_3'(\mathfrak{s}_o) = \mathfrak{s}_{\tilde{o}}$.

Finally, suppose the strands are oriented as shown in Fig. 9c. In this case, the oriented resolution of L is in $L(*010)$, and the oriented resolution of $\tilde{L}$ is in $\tilde{L}(*100)$. Inside the region under consideration, $\mathfrak{s}_o$ looks like the state of Fig. 9c (perhaps with **a**'s and **b**'s reversed.) Our first step is to exhibit some $\mathfrak{t} \in X_1$ which is homologous to $\mathfrak{s}_o$. As before, we denote by $\mathfrak{s}_{ijk}$ the unique state of $L(*ijk)$ which is the same as $\mathfrak{s}_o$ outside the area of the Reidemeister move and has all its components inside this area labeled by **a**'s.

Assume for the moment that all three strands shown in $L(*000)$ belong to different components. In this case, we can take

$$\mathfrak{t} = \mathfrak{s}_o - 2\mathfrak{s}_{100} - \mathfrak{s}_{010} - 2\mathfrak{s}_{001} = \mathfrak{s}_o - d'(\mathfrak{s}_{000}).$$

Indeed, $\beta'(-2\mathfrak{s}_{100}) = \mathfrak{s}_o - \mathfrak{s}_{010}$ and $\mathfrak{s}_{001} \in C_{Lee}(L(*1))$, so $\mathfrak{t} \in X_1$. Then

$$\begin{aligned}\rho_3'(\mathfrak{t}) &= -2\mathfrak{s}_{\widetilde{010}} - 2\tilde{\beta}'(\mathfrak{s}_{\widetilde{010}}) - 2\mathfrak{s}_{\widetilde{001}}\\ &= -2\mathfrak{s}_{\widetilde{010}} - 2\mathfrak{s}_{\widetilde{100}} + 2\mathfrak{s}_{\tilde{o}} - 2\mathfrak{s}_{\widetilde{001}}\\ &= 2\mathfrak{s}_{\tilde{o}} - d'(\mathfrak{s}_{\widetilde{000}})\end{aligned}$$

which proves the claim.

We leave it to the reader to check that a similar argument applies to each of the four other ways in which the segments outside the area of the move can be connected, as well as when the roles of **a** and **b** are reversed. In each case, it is not difficult to verify that $\rho_{3*}'([\mathfrak{s}_0])$ is one of $\pm[\mathfrak{s}_{\tilde{o}}]$, $\pm 2[\mathfrak{s}_{\tilde{o}}]$, or $\pm\frac{1}{2}[\mathfrak{s}_{\tilde{o}}]$. $\square$

Proof of Proposition 3.2 In the case of ρ'_{1*} and ρ'_{2*}, the claim is immediate, since these maps are induced by filtered chain maps. For the others, we use the following

Lemma 6.2 *Suppose $f : C \to C'$ is a map of filtered chain complexes with the property that the induced map of spectral sequences $f^1 : E^1 \to E'^1$ is an isomorphism. Then f_*^{-1} is a filtered map with respect to the induced filtrations on $H_*(C)$ and $H_*(C')$.*

Proof Since f^1 is an isomorphism, f^∞ (the induced map on filtered gradeds) is as well. It follows that f_* is an isomorphism. Suppose f_*^{-1} does not respect the filtration. Then there must be some $\mathbf{v} \in H_*(C)$ whose filtration is strictly increased by f_*. But this contradicts the fact that f^∞ is an isomorphism. □

The remaining cases now follow easily from the results used in the proof of Theorem 2.1. Indeed, ρ'_1 and ρ'_2 both induce isomorphisms of E^1 terms, and $\rho'_{3*} = \iota_{1*} \circ \psi_* \circ \iota_{2*}^{-1}$, where ι_1, ι_2, and ψ all induce isomorphisms of E^1 terms. □

Acknowledgements The author would like to thank Peter Kronheimer, Peter Ozsváth, and Zoltan Szabó for many helpful conversations and for encouraging him to pursue this problem and the referee for helpful suggestions regarding the exposition.

References

1. Bar-Natan, D.: On Khovanov's categorification of the Jones polynomial. Algebr. Geom. Topol. **2**, 337–370 (2002)
2. Bar-Natan, D.: The knot atlas. www.math.toronto.edu/~drorbn/KAtlas/index.html (2003)
3. Bar-Natan, D.: Khovanov's homology for tangles and cobordisms. math.GT/0410495 (2004)
4. Carter, J.S., Saito, M.: Reidemeister moves for surface isotopies and their interpretation as moves to movies. J. Knot Theory Ramif. **2**, 251–284 (1993)
5. Freedman, M.: A surgery sequence in dimension four; the relations with knot concordance. Invent. Math. **68**, 195–226 (1982)
6. Goda, H., Matsuda, H., Morifuji, T.: Knot Floer homology of (1,1)-knots. math.GT/0311084 (2003)
7. Gompf, R.E., Stipsicz, A.I.: 4-Manifolds and Kirby Calculus. Graduate Studies in Mathematics, vol. 20. Am. Math. Soc., Providence (1999)
8. Khovanov, M.: A categorification of the Jones polynomial. Duke Math. J. **101**, 359–426 (2000)
9. Khovanov, M.: Patterns in knot cohomology I. math.QA/0201306 (2002)
10. Kronheimer, P.B., Mrowka, T.S.: Gauge theory for embedded surfaces. I. Topology **32**, 773–826 (1993)
11. Lee, E.S.: Khovanov's invariants for alternating links. math.GT/0210213 (2002)
12. Lee, E.S.: The support of Khovanov's invariants for alternating knots. math.GT/0201105 (2002)

13. Livingston, C.: Computations of the Ozsvath-Szabo knot concordance invariant. math.GT/0311036 (2003)
14. McCleary, J.: User's Guide to Spectral Sequences. Mathematics Lecture Series, vol. 12. Publish or Perish, Wilmington (1985)
15. Ozsváth, P., Szabó, Z.: Heegaard Floer homology and alternating knots. Geom. Topol. **7**, 225–254 (2002). math.GT/0209149
16. Ozsváth, P., Szabó, Z.: Holomorphic disks and knot invariants. math.GT/0209056 (2002)
17. Ozsváth, P., Szabó, Z.: Knot Floer homology and the four-ball genus. Geom. Topol. **7**, 615–639 (2003). math.GT/0301026
18. Ozsváth, P., Szabó, Z.: Knot Floer homology, genus bounds, and mutation. math.GT/0303225 (2003)
19. Plamenevskaya, O.: Transverse knots and Khovanov homology. math.GT/0412184 (2004)
20. Rasmussen, J.: Floer homology and knot complements. math.GT/0306378 (2003)
21. Rudolph, L.: Positive links are strongly quasipositive. Geom. Topol. Monogr. **2**, 555–562 (1999)
22. Shumakovitch, A.: KhoHo pari package. www.geometrie.ch/KhoHo/ (2003)
23. Shumakovitch, A.: Rasmussen invariant, Slice-Bennequin inequality, and sliceness of knots. math.GT/0411643 (2004)