The Dowker theorem via discrete Morse theory

Morten Brun*, Darij Grinberg†

20 July 2024

Abstract. The Dowker theorem is a classical result in the topology of finite spaces, claiming that any binary relation between two finite spaces defines two homotopy-equivalent complexes (the Dowker complexes). Recently, Barmak strengthened this to a simple-homotopy-equivalence. We reprove Barmak's result using a combinatorial argument that constructs an explicit acyclic matching in the sense of discrete Morse theory.

Keywords: Dowker duality, simplicial complexes, discrete Morse theory, combinatorial topology, finite topological spaces, acyclic matching.

Mathematics Subject Classification 2020: 57Q10, 55-01, 55U10.

1. Introduction

The classical Dowker theorem (in the form given by Björner [Bjorne95, Theorem 10.9]) has recently been revisited by Brun and Salbu [BruSal22], who reproved it using what they call the *rectangle complex*.

We give a simpler proof using what we call the *biclique complex*. This proof is purely combinatorial, as it relies on discrete Morse theory instead of homotopy theory.

We briefly discuss the relation between the biclique and rectangle complexes.

2. Notation

We recall the notion of a simplicial complex:

*Department of Mathematics, University of Bergen, Bergen, Norway. (`morten.brun@uib.no`)
†Department of Mathematics, Drexel University, Philadelphia, U.S.A. (`darijgrinberg@gmail.com`)

Definition 2.1. Let W be a finite set. Then, a *simplicial complex* (or, for short, *complex*) with ground set W is defined to be a set Δ of subsets of W that is closed under taking subsets (i.e., we have $G \in \Delta$ whenever G is a subset of a set $F \in \Delta$). The elements of Δ are called the *faces* of Δ.

Note that (unlike some authors) we do not require that each $w \in W$ belongs to some face of Δ; thus, one and the same Δ can be a complex with several ground sets. However, the properties of complexes that we will study (homotopical and Morse-theoretical ones) do not depend on the ground set.

Any complex Δ is a set of sets, and thus is partially ordered by inclusion; hence, it becomes a poset. We will use the notations $\succ$ and $\prec$ for the cover relations of this poset. Thus, two faces A and B satisfy $A \prec B$ (or, equivalently, $B \succ A$) if and only if $A \subseteq B$ and $|B \setminus A| = 1$.

3. The Dowker theorem

Let X and Y be two finite sets. Let R be a binary relation from X to Y, that is, a subset of $X \times Y$. Given two elements $x \in X$ and $y \in Y$, we write $x\ R\ y$ for $(x, y) \in R$.

Given two subsets $U \subseteq X$ and $V \subseteq Y$, we write $U\ \mathbf{R}\ V$ if and only if all $u \in U$ and all $v \in V$ satisfy $u\ R\ v$ (that is, if and only if $U \times V \subseteq R$). This defines a binary relation $\mathbf{R}$ from the power set of X to the power set of Y. Note that $\varnothing\ \mathbf{R}\ V$ holds (for vacuous reasons) for each $V \subseteq Y$, and likewise we have $U\ \mathbf{R}\ \varnothing$ for each $U \subseteq X$.

Example 3.1. Let $X = \{1, 2, 3, 4\}$ and $Y = \{5, 6, 7, 8\}$, and let R be the "divides" relation (i.e., we set $x\ R\ y$ if and only if $x \mid y$ in $\mathbb{Z}$). Then, $\{1, 2\}\ \mathbf{R}\ \{6, 8\}$ but not $\{1, 2, 4\}\ \mathbf{R}\ \{6, 8\}$ (since $4 \nmid 6$).

A *Y-neighbor* of a subset $U \subseteq X$ is defined to be an element $y \in Y$ such that $U\ \mathbf{R}\ \{y\}$. In other words, a Y-neighbor of a subset $U \subseteq X$ means a $y \in Y$ such that all $u \in U$ satisfy $u\ R\ y$.

A *Y-conic set* shall mean a subset $U \subseteq X$ that is empty or has a Y-neighbor. We define C_X to be the set of all Y-conic sets. This C_X is a simplicial complex with ground set X (since a subset of a Y-conic set is again a Y-conic set).

Likewise we define a simplicial complex C_Y with ground set Y: An *X-neighbor* of a subset $V \subseteq Y$ is defined to be an element $x \in X$ such that $\{x\}\ \mathbf{R}\ V$. An *X-conic set* shall mean a subset $V \subseteq Y$ that is empty or has an X-neighbor. We define C_Y to be the set of all X-conic sets.

The complexes C_X and C_Y will be called the *left Dowker complex* and the *right Dowker complex* of the relation R. They have been called K and L in [Dowker52, §1], have been called $D(R)$ and $D(R^T)$ in [BruSal22], and have been called Δ_0 and Δ_1 in [Bjorne95, Theorem 10.9]. (Some of these sources are slightly imprecise around the empty set, or use a different notion of simplicial complexes that defines them to consist of nonempty subsets. In substance, all definitions agree when the sets X and Y are nonempty; we believe that ours is best suited for the empty case.)

Example 3.2. Let X, Y and R be as in Example 3.1. Then, it is easy to see that

$$C_X = \{\text{all } U \subseteq X \text{ such that } \{3,4\} \not\subseteq U\} \qquad \text{and} \qquad C_Y = \{\text{all } V \subseteq Y\}.$$

(The latter is because 1 is an X-neighbor of any subset $V \subseteq Y$.)

Remark 3.3. The binary relation R can be re-encoded as a bipartite graph whose black vertices are the elements of X, whose white vertices are the elements of Y, and whose edges are the edges $\{x, y\}$ for all $x \in X$ and $y \in Y$ satisfying $x \; R \; y$. Thus, a Y-conic set means a set of black vertices that is either empty or has a common neighbor, whereas an X-conic set means a set of white vertices that is either empty or has a common neighbor.

Dowker's theorem (in Björner's version [Bjorne95, Theorem 10.9]) says the following:

Theorem 3.4 (Dowker's theorem)**.** The complexes C_X and C_Y are homotopy-equivalent.

In [Barmak11, Theorem 4.4], Barmak strengthened this theorem to a simple-homotopy-equivalence:

Theorem 3.5. The complexes C_X and C_Y are simple-homotopy-equivalent.

In the two following sections, we shall reprove Theorem 3.5 (and thus also Theorem 3.4) using discrete Morse theory, by exhibiting a larger complex that collapses to each of C_X and C_Y. (By [Kozlov20, Proposition 9.28], this suffices to show that C_X and C_Y are simple-homotopy-equivalent.) We note that this claim is purely combinatorial, requiring no topology to state, and our proof will be similarly combinatorial.

4. Reminders on discrete Morse theory

Before we step to the actual proof, let us quickly recall some preliminaries.

First, we recall the notion of a *subcomplex*:

Definition 4.1. Let Δ and Γ be two complexes (possibly with different ground sets). Then, we say that Γ is a *subcomplex* of Δ if and only if $\Gamma \subseteq \Delta$.

Next, and more importantly, let us recall the discrete Morse theory that we will need. In this, we follow [Kozlov20, Chapter 10], adapted slightly to simplify our life (copying several definitions from [GrKaLe21, §7]). We begin with the definition of an elementary simplicial collapse ([Kozlov20, §9.1]):

Definition 4.2. Let Γ be a subcomplex of a simplicial complex Δ. If there exist faces τ and σ in $\Delta \setminus \Gamma$ such that

$$\Delta = \Gamma \cup \{\tau, \sigma\} \qquad \text{and} \qquad \tau \prec \sigma,$$

then we say that Δ *collapses to* Γ by an *elementary simplicial collapse*. (Equivalently, we say in this case that the inclusion $\Gamma \subseteq \Delta$ is an *elementary simplicial expansion*.)

This definition is equivalent to [Kozlov20, Definition 9.1]. (Indeed, [Kozlov20, Definition 9.1] requires that the only faces of Δ that contain τ are τ and σ, whereas we instead require that Γ is a subcomplex; but these two requirements imply each other.)

Elementary simplicial collapses can be composed, leading to the notion of (non-elementary) collapses ([Kozlov20, Definition 9.4]):

Definition 4.3. Let Γ and Δ be two simplicial complexes. We say that Δ *collapses to* Γ if Γ can be obtained from Δ by a sequence (possibly empty) of elementary simplicial collapses – i.e., if there exist finitely many simplicial complexes $\Delta_0, \Delta_1, \ldots, \Delta_n$ with $\Delta_0 = \Delta$ and $\Delta_n = \Gamma$ and such that for each $i \in \{0, 1, \ldots, n-1\}$, the complex Δ_i collapses to Δ_{i+1} by an elementary simplicial collapse. (This clearly implies that Γ is a subcomplex of Δ.)

We shall use the notation "$\Delta \searrow \Gamma$" for the statement "Δ collapses to Γ".

Composing elementary simplicial collapses with their inverses leads to the weaker notion of simple-homotopy-equivalence ([Kozlov20, Definition 9.26], where it is called "having the same simple homotopy type"):

Definition 4.4. Two simplicial complexes Γ and Δ are said to be *simple-homotopy-equivalent* if there exist finitely many simplicial complexes $\Delta_0, \Delta_1, \ldots, \Delta_n$ with $\Delta_0 = \Delta$ and $\Delta_n = \Gamma$ and with the following property: For each $i \in \{0, 1, \ldots, n-1\}$, either Δ_{i+1} collapses to Δ_i or Δ_i collapses to Δ_{i+1}.

Clearly, simple-homotopy-equivalence is an equivalence relation.

Next, we recall the definition of a matching ([Kozlov20, Definition 10.6], specialized to subposets of a complex):

Definition 4.5. Let Δ be a simplicial complex. A *partial matching* (or *matching* for short) on Δ shall mean a pair (M, μ), where M is a subset of Δ (that is, a set of faces of Δ), and where $\mu : M \to M$ is an involution (that is, a map satisfying $\mu \circ \mu = \mathrm{id}$) with the property that each $F \in M$ satisfies

$$\text{either } \mu(F) \prec F \text{ or } \mu(F) \succ F.$$

Note that M is uniquely determined by μ (namely, as the domain of μ), so that we will refer to μ alone as a matching.

Example 4.6. Let $W = \{1, 2, 3\}$. Let Δ be the complex with ground set W that contains all the eight subsets of W as faces. Set $M := \Delta \setminus \{\varnothing, W\}$, and define a map $\mu : M \to M$ by

$$\mu(\{1\}) = \{1,2\}, \qquad \mu(\{2\}) = \{2,3\}, \qquad \mu(\{3\}) = \{3,1\},$$
$$\mu(\{1,2\}) = \{1\}, \qquad \mu(\{2,3\}) = \{2\}, \qquad \mu(\{3,1\}) = \{3\}.$$

Then, (M, μ) is a matching on Δ.

Partial matchings on simplicial complexes are useful for counting purposes (being a simple kind of sign-reversing involutions; cf. [Sagan19, Chapter 2]). However, under a certain condition, they become powerful tools for understanding the topology of Δ. This condition is known as "acyclicity", and is defined as follows ([Kozlov20, Definition 10.7]):

Definition 4.7. Let Δ be a simplicial complex. Let (M, μ) be a matching on Δ.

(a) A *cycle* of μ means an n-tuple $(F_1, F_2, \ldots, F_n)$ of distinct faces in M such that $n \geq 2$ and

$$F_1 \succ \mu(F_1) \prec F_2 \succ \mu(F_2) \prec F_3 \succ \cdots \prec F_n \succ \mu(F_n) \prec F_1$$

(that is, such that each $i \in \{1, 2, \ldots, n\}$ satisfies $F_i \succ \mu(F_i) \prec F_{i+1}$, where $F_{n+1} := F_1$).

(b) The matching μ is said to be *acyclic* if it has no cycle.

Example 4.8. Let W, Δ, M and μ be as in Example 4.6. Then, the 3-tuple $(\{1,2\}, \{3,1\}, \{2,3\})$ is a cycle of the matching μ, since

$$\{1,2\} \succ \underbrace{\{1\}}_{=\mu(\{1,2\})} \prec \{3,1\} \succ \underbrace{\{3\}}_{=\mu(\{3,1\})} \prec \{2,3\} \succ \underbrace{\{2\}}_{=\mu(\{2,3\})} \prec \{1,2\}.$$

Thus, the matching μ is not acyclic. However, if we remove the faces $\{1\}$ and $\{1,2\}$ from M (and correspondingly restrict μ to the set of the remaining four faces in M), then the cycle disappears, and we obtain an acyclic matching.

Discrete Morse theory can be seen as the study of acyclic matchings. Roughly speaking, an acyclic matching (M, μ) on a complex Δ allows us to "cancel" the faces in M when computing homotopical or homological data, albeit the precise meaning of "cancelling" here depends on the situation. Different instances of this principle can be found in [Kozlov20, Chapters 10, 11, 12, 13]; we will only use the following one (a part of [Kozlov20, Theorem 10.9]):

Theorem 4.9. Let Γ be a subcomplex of a simplicial complex Δ. Assume that there exists an acyclic matching (M, μ) on Δ with $M = \Delta \setminus \Gamma$ (that is, M consists of all faces of Δ that don't belong to Γ). Then, $\Delta \searrow \Gamma$.

Acyclic matchings are combinatorial objects and can often be constructed by hand. However, there is also a number of "macros" available to construct them from simpler objects, such as the pairing lemma of Linusson and Shareshian ([LinSha03, Lemma 3.4]). We will use one such "macro" – actually a version of [LinSha03, Lemma 3.4] adapted to our situation:

Lemma 4.10. Let W be a finite partially ordered set. Let Δ be a simplicial complex with ground set W. Let M be a subset of Δ. Let $f : M \to W$ be a function. Consider the following two conditions:

(C1) For any $F \in M$, we have $F \cup \{f(F)\} \in M$ and $F \setminus \{f(F)\} \in M$ and

$$f(F \cup \{f(F)\}) = f(F \setminus \{f(F)\}) = f(F). \tag{1}$$

(C2) For any $F \in M$ and any $G \in M$ satisfying $G \subseteq F$, we have $f(F) \leq f(G)$.

Assume that Condition **(C1)** holds. Then:

(a) The map

$$\mu : M \to M,$$
$$F \mapsto \begin{cases} F \setminus \{f(F)\}, & \text{if } f(F) \in F; \\ F \cup \{f(F)\}, & \text{if } f(F) \notin F \end{cases}$$

is well-defined, and the pair (M, μ) is a matching on Δ.

(b) If Condition **(C2)** holds as well, then this matching (M, μ) is acyclic.

Proof. Condition **(C1)** ensures that each $F \in M$ satisfies $F \setminus \{f(F)\} \in M$ and $F \cup \{f(F)\} \in M$ and therefore

$$\begin{cases} F \setminus \{f(F)\}, & \text{if } f(F) \in F; \\ F \cup \{f(F)\}, & \text{if } f(F) \notin F \end{cases} \in M.$$

Hence, the map

$$\mu : M \to M,$$
$$F \mapsto \begin{cases} F \setminus \{f(F)\}, & \text{if } f(F) \in F; \\ F \cup \{f(F)\}, & \text{if } f(F) \notin F \end{cases}$$

is well-defined. We shall now show that this map is an involution:

Claim 1: We have $\mu(\mu(G)) = G$ for each $G \in M$.

Proof of Claim 1. Let $G \in M$. We must show that $\mu(\mu(G)) = G$. We are in one of the following two cases:

Case 1: We have $f(G) \in G$.

Case 2: We have $f(G) \notin G$.

Let us first consider Case 1. In this case, we have $f(G) \in G$. Hence, the definition of μ yields $\mu(G) = G \setminus \{f(G)\}$. Let $F = \mu(G)$. Thus, $F = \mu(G) = G \setminus \{f(G)\}$. Hence, $f(G) \notin F$.

The second equality sign in the equality (1) (applied to G instead of F) yields $f(G \setminus \{f(G)\}) = f(G)$. In other words, $f(F) = f(G)$ (since $F = G \setminus \{f(G)\}$). Thus, $f(F) = f(G) \notin F$. Hence, the definition of μ yields

$$\mu(F) = \underbrace{F}_{=G\setminus\{f(G)\}} \cup \left\{\underbrace{f(F)}_{=f(G)}\right\} = (G \setminus \{f(G)\}) \cup \{f(G)\} = G$$

(since $f(G) \in G$). In other words, $\mu(\mu(G)) = G$ (since $F = \mu(G)$). This proves Claim 1 in Case 1.

Let us now consider Case 2. In this case, we have $f(G) \notin G$. Hence, the definition of μ yields $\mu(G) = G \cup \{f(G)\}$. Let $F = \mu(G)$. Thus, $F = \mu(G) = G \cup \{f(G)\}$. Hence, $f(G) \in F$.

From (1) (applied to G instead of F), we obtain $f(G \cup \{f(G)\}) = f(G)$. In other words, $f(F) = f(G)$ (since $F = G \cup \{f(G)\}$). Thus, $f(F) = f(G) \in F$. Hence, the definition of μ yields

$$\mu(F) = \underbrace{F}_{=G\cup\{f(G)\}} \setminus \left\{\underbrace{f(F)}_{=f(G)}\right\} = (G \cup \{f(G)\}) \setminus \{f(G)\} = G$$

(since $f(G) \notin G$). In other words, $\mu(\mu(G)) = G$ (since $F = \mu(G)$). This proves Claim 1 in Case 2.

We have now proved Claim 1 in both Cases 1 and 2. Hence, Claim 1 is proved. □

As a byeffect of our proof of Claim 1, we obtain the following:

Claim 2: Let $G \in M$. Then, $f(\mu(G)) = f(G)$.

Proof of Claim 2. Let $F = \mu(G)$. Then, $f(F) = f(G)$ (indeed, this equality has been proved in our above proof of Claim 1, in both Cases 1 and 2). In other words, $f(\mu(G)) = f(G)$ (since $F = \mu(G)$). This proves Claim 2. □

Here is a simple set-theoretical fact that will come useful later:

> *Claim 3:* Let F, G, H be three subsets of W such that $F \succ H \prec G$. Let s be an element such that $s \in F$ and $s \in G$ but $s \notin H$. Then, $F = G$.

Proof of Claim 3. Left to the reader. (Show that both F and G equal $H \cup \{s\}$.) □

Now, we prove one more claim about μ:

> *Claim 4:* Let $F, G \in M$ be such that $F \succ \mu(F) \prec G \succ \mu(G)$ and $f(F) = f(G)$. Then, $F = G$.

Proof of Claim 4. If we had $f(F) \notin F$, then the definition of μ would yield $\mu(F) = F \cup \{f(F)\} \supseteq F$, which would contradict $F \succ \mu(F)$. Thus, we must have $f(F) \in F$. Similarly, $f(G) \in G$ (since $G \succ \mu(G)$).

Let us denote the element $f(F) = f(G)$ by s. Thus, $s = f(F) \in F$ and $s = f(G) \in G$.

We have $f(F) \in F$. Thus, the definition of μ yields $\mu(F) = F \setminus \{f(F)\} = F \setminus \{s\}$ (since $f(F) = s$). Hence, $s \notin \mu(F)$ (since $s \notin F \setminus \{s\}$). Thus, Claim 3 (applied to $H = \mu(F)$) yields $F = G$. This proves Claim 4. □

Claim 1 shows that $\mu \circ \mu = \text{id}$. In other words, the map μ is an involution. Moreover, it has the property that each $F \in M$ satisfies

$$\text{either } \mu(F) \prec F \text{ or } \mu(F) \succ F$$

(since we either have $f(F) \in F$, in which case $\mu(F) = F \setminus \{f(F)\} \prec F$, or we have $f(F) \notin F$, in which case $\mu(F) = F \cup \{f(F)\} \succ F$). Thus, the map μ (or, to be more precise, the pair (M, μ)) is a matching on Δ. This proves Lemma 4.10 **(a)**.

(b) Assume that Condition **(C2)** holds as well. Lemma 4.10 **(a)** shows that the map μ (defined in said lemma) is well-defined, and that the pair (M, μ) is a matching on Δ.

It remains to show that this matching (M, μ) is acyclic. In other words, we must show that it has no cycle.

Assume the contrary. Thus, (M, μ) has a cycle. By the definition of a cycle, this cycle is an n-tuple $(F_1, F_2, \ldots, F_n)$ of distinct faces in M such that $n \geq 2$ and

$$F_1 \succ \mu(F_1) \prec F_2 \succ \mu(F_2) \prec F_3 \succ \cdots \prec F_n \succ \mu(F_n) \prec F_1.$$

Consider this n-tuple $(F_1, F_2, \ldots, F_n)$. Set $F_{n+1} := F_1$. Thus, each $i \in \{1, 2, \ldots, n\}$ satisfies

$$F_i \succ \mu(F_i) \prec F_{i+1}. \tag{2}$$

Recall that the n faces $F_1, F_2, \ldots, F_n$ are distinct; thus, $F_1 \neq F_2$ (since $n \geq 2$).

Now, let $i \in \{1, 2, \ldots, n\}$. Then, Claim 2 (applied to $G = F_i$) yields $f(\mu(F_i)) = f(F_i)$. Furthermore, $\mu(F_i) \subseteq F_{i+1}$ (since (2) yields $\mu(F_i) \prec F_{i+1}$). Thus, Condition **(C2)** (applied to $F = F_{i+1}$ and $G = \mu(F_i)$) yields $f(F_{i+1}) \leq f(\mu(F_i)) = f(F_i)$. In other words, $f(F_i) \geq f(F_{i+1})$.

Forget that we fixed i. We thus have proved that $f(F_i) \geq f(F_{i+1})$ for each $i \in \{1, 2, \ldots, n\}$. In other words,

$$f(F_1) \geq f(F_2) \geq \cdots \geq f(F_{n+1}). \tag{3}$$

Hence, in particular, $f(F_1) \geq f(F_2)$ and $f(F_2) \geq f(F_{n+1})$. Hence, $f(F_2) \geq f(F_{n+1}) = f(F_1)$ (since $F_{n+1} = F_1$). Combining this with $f(F_1) \geq f(F_2)$, we obtain $f(F_1) = f(F_2)$.

Thus, Claim 4 (applied to $F = F_1$ and $G = F_2$) yields $F_1 = F_2$ (since $F_1 \succ \mu(F_1) \prec F_2 \succ \mu(F_2)$). But this contradicts $F_1 \neq F_2$. This contradiction shows that our assumption was wrong. Hence, we have shown that the matching (M, μ) has no cycle, i.e., is acyclic. This proves Lemma 4.10 **(b)**. □

5. The biclique complex of a bipartite relation

Let R be relation from a set X to a set Y. We say that R is a *bipartite relation* if the sets X and Y are disjoint. In this section we assume that R is a bipartite relation.

A *biclique* of the relation R shall mean a set of the form $U \cup V$, where $U \subseteq X$ and $V \subseteq Y$ are two nonempty subsets satisfying $U \mathbf{R} V$. (Note the "nonempty" requirement! Thus, a biclique must intersect both X and Y.) In this section, we will only deal with bicliques of R, so we will just refer to them as "bicliques".

Clearly, any biclique is a subset of $X \cup Y$. The following facts are near-trivial:

Lemma 5.1. Let $U \subseteq X$ and $V \subseteq Y$ be two subsets. Let $F = U \cup V$. Then, $F \cap X = U$ and $F \cap Y = V$.

Proof. This is because X and Y are disjoint. □

Lemma 5.2. Let G be a subset of a biclique. Then:

(a) If $G \subseteq X$, then $G \in C_X$.

(b) If $G \subseteq Y$, then $G \in C_Y$.

Proof. **(a)** Assume that $G \subseteq X$. We know that G is a subset of a biclique. Let F be this biclique. Then, $G \subseteq F$.

But F is a biclique. In other words, $F = U \cup V$ for some nonempty subsets $U \subseteq X$ and $V \subseteq Y$ satisfying $U \mathbf{R} V$ (by the definition of a biclique). Consider these U and V.

Combining $G \subseteq F$ and $G \subseteq X$, we obtain $G \subseteq F \cap X = U$ (by Lemma 5.1).

There exists some $v \in V$ (since V is nonempty). This v is then a Y-neighbor of U (since $U \mathbf{R} V$), and thus a Y-neighbor of G (since $G \subseteq U$). Thus, the set G has a Y-neighbor, and hence is Y-conic, i.e., belongs to C_X. This proves Lemma 5.2 **(a)**.

(b) This is analogous to part **(a)**. □

Proposition 5.3. Set
$$B := \{\text{bicliques}\} \cup C_X \cup C_Y.$$
Then, B is a simplicial complex with ground set $X \cup Y$.

Proof. Let $F \in B$. We must prove that every subset of F belongs to B as well. If $F \in C_X$, then this is clear (since C_X is a complex); likewise if $F \in C_Y$. Thus, we only need to consider the case when $F \in \{\text{bicliques}\}$. So let us assume that $F \in \{\text{bicliques}\}$.

Thus, F is a biclique. A subset of F can be either fully contained in X, or fully contained in Y, or intersect both X and Y. In the first of these three cases, it belongs to C_X (by Lemma 5.2 **(a)**). In the second, it belongs to C_Y (by Lemma 5.2 **(b)**). In the third, it is itself a biclique (since $U \mathbf{R} V$ entails $U' \mathbf{R} V'$ whenever $U' \subseteq U$ and $V' \subseteq V$). In either case, it thus belongs to B (by the definition of B). Proposition 5.3 is thus proven. □

The biclique complex B has been independently introduced by Yoon in [Yoon24, Definition 3.13], where she named it the "Dowker join complex". She, too, used it as an intermediary to prove the Dowker theorem (Theorem 3.4), although her proof differs from ours and relies on Quillen's poset fiber lemma.

6. Collapses of the biclique complex

In this section, we still assume that R is a bipartite relation. Consider the simplicial complex B defined in Proposition 5.3. We shall call B the *biclique complex* of the relation R. Both C_X and C_Y are subcomplexes of B. In this section, we show that the biclique complex B collapses to each of the complexes C_X and C_Y. This argument will then lead us to what is perhaps the simplest proofs of Theorem 3.4 and Theorem 3.5.

Our main claim is the following:

Theorem 6.1. Let R be a bipartite relation from a set X to a set Y. Then the biclique complex B of R collapses to both C_X and C_Y. That is, $B \searrow C_X$ and $B \searrow C_Y$.

Proof of Theorem 6.1. We shall only show that $B \searrow C_X$, since the proof of $B \searrow C_Y$ is analogous.

Let $M := B \setminus C_X$. This is a subset of B, and consists of those faces of B that are not faces of C_X. We shall prove some easy claims:

> *Claim 1:* Let $F \in M$. Then, the set $F \cap Y$ is nonempty and has an X-neighbor.

Proof of Claim 1. We have $F \in M = B \setminus C_X$, so that $F \in B$ but $F \notin C_X$. From $F \notin C_X$, we obtain $F \neq \varnothing$ (since the definition of C_X yields $\varnothing \in C_X$), so that F is nonempty. Moreover,

$$F \in B = \{\text{bicliques}\} \cup C_X \cup C_Y.$$

In other words, $F \in \{\text{bicliques}\}$ or $F \in C_X$ or $F \in C_Y$. Since $F \notin C_X$, we can rule out the second possibility, so we conclude that $F \in \{\text{bicliques}\}$ or $F \in C_Y$. Thus, we are in one of the following two cases:

Case 1: We have $F \in \{\text{bicliques}\}$.

Case 2: We have $F \in C_Y$.

Let us consider Case 1. In this case, F is a biclique. In other words, $F = U \cup V$ for some nonempty subsets $U \subseteq X$ and $V \subseteq Y$ satisfying $U \mathbf{R} V$ (by the definition of a biclique). Consider these U and V. Any element of U is an X-neighbor of V (since $U \mathbf{R} V$). Hence, the set V has an X-neighbor (since U is nonempty).

So we have shown that the set V is nonempty and has an X-neighbor. In other words, the set $F \cap Y$ is nonempty and has an X-neighbor (since Lemma 5.1 yields $F \cap Y = V$). This proves Claim 1 in Case 1.

Let us now consider Case 2. In this case, we have $F \in C_Y$. In other words, F is an X-conic set. Since F is nonempty, this entails that F has an X-neighbor. Moreover, $F \subseteq Y$ (since F is an X-conic set), so that $F \cap Y = F$. Hence, the set $F \cap Y$ is nonempty and has an X-neighbor (since F is nonempty and has an X-neighbor). This proves Claim 1 in Case 2.

We have now proved Claim 1 in both Cases 1 and 2. Hence, Claim 1 always holds. $\square$

Let us now fix a total order on the set $X \cup Y$ (chosen arbitrarily). We define a map $f : M \to X \cup Y$ as follows: For each face $F \in M$,

$$\text{we let } f(F) \text{ be the largest } X\text{-neighbor of the set } F \cap Y$$

(this is well-defined, since Claim 1 shows that $F \cap Y$ has an X-neighbor). Note that $f(F)$ actually belongs to X (by the definition of an X-neighbor), but we prefer to use $X \cup Y$ as the target in order to agree with Lemma 4.10 notationally.

We now claim the following:

> *Claim 2:* For any $F \in M$, we have $F \cup \{f(F)\} \in M$ and $F \setminus \{f(F)\} \in M$ and
> $$f(F \cup \{f(F)\}) = f(F \setminus \{f(F)\}) = f(F).$$

Proof of Claim 2. Let $F \in M$. Thus, $F \in M = B \setminus C_X$, so that $F \in B$ and $F \notin C_X$.

By the definition of f, the element $f(F)$ is the largest X-neighbor of the set $F \cap Y$. Thus, $f(F)$ is an X-neighbor of $F \cap Y$, so that $f(F) \in X$ (by the definition of an X-neighbor), and thus $f(F) \notin Y$ (since the sets X and Y are disjoint).

Set

$$F^+ := F \cup \{f(F)\} \qquad \text{and} \qquad F^- := F \setminus \{f(F)\}.$$

Thus, the sets F^+ and F^- differ from F only in the single element $f(F)$ (if they differ from F at all). Thus, they contain the same elements of Y as F (since $f(F) \notin Y$). In other words,

$$F^+ \cap Y = F \cap Y \qquad \text{and} \qquad F^- \cap Y = F \cap Y.$$

Claim 1 shows that the set $F \cap Y$ is nonempty. Hence, the set $F^- \cap Y$ is nonempty (since $F^- \cap Y = F \cap Y$). In other words, the set F^- intersects Y. Thus, F^- is not a subset of X (since the sets X and Y are disjoint). Thus, $F^- \notin C_X$ (since any face of C_X is a subset of X). Moreover, $F^- = F \setminus \{f(F)\} \subseteq F$. Thus, from $F \in B$, we obtain $F^- \in B$ (since B is a simplicial complex). Combining this with $F^- \notin C_X$, we obtain

$$F^- \in B \setminus C_X = M.$$

We have $F \subseteq F \cup \{f(F)\} = F^+$. Thus, if we had $F^+ \in C_X$, then we would obtain $F \in C_X$ (since C_X is a simplicial complex), which would contradict $F \notin C_X$. Hence, $F^+ \notin C_X$.

Let us now show that $F^+ \in B$. As in the proof of Claim 1, we can show that $F \in \{\text{bicliques}\}$ or $F \in C_Y$. Thus, we are in one of the following two cases:

Case 1: We have $F \in \{\text{bicliques}\}$.

Case 2: We have $F \in C_Y$.

Let us consider Case 1. In this case, F is a biclique. In other words, $F = U \cup V$ for some nonempty subsets $U \subseteq X$ and $V \subseteq Y$ satisfying $U \mathbf{R} V$ (by the definition of a biclique). Consider these U and V. From Lemma 5.1, we obtain $F \cap X = U$ and $F \cap Y = V$. Moreover, the set $U \cup \{f(F)\}$ is clearly nonempty and satisfies $U \cup \{f(F)\} \subseteq X$ (since $U \subseteq X$ and $f(F) \in X$).

But $f(F)$ is an X-neighbor of $F \cap Y$. In other words, $f(F)$ is an X-neighbor of V (since $F \cap Y = V$). Moreover, each element of U is an X-neighbor of V (since $U \mathbf{R} V$). Combining the preceding two sentences, we conclude that each element of $U \cup \{f(F)\}$ is an X-neighbor of V. In other words, $(U \cup \{f(F)\}) \mathbf{R} V$. Hence, the set $(U \cup \{f(F)\}) \cup V$ is a biclique (by the definition of a biclique, since $U \cup \{f(F)\} \subseteq X$ and $V \subseteq Y$ are nonempty subsets satisfying $(U \cup \{f(F)\}) \mathbf{R} V$). Now,

$$\begin{aligned} F^+ &= \underbrace{F}_{=U \cup V} \cup \{f(F)\} = U \cup V \cup \{f(F)\} = (U \cup \{f(F)\}) \cup V \\ &\in \{\text{bicliques}\} \qquad (\text{since } (U \cup \{f(F)\}) \cup V \text{ is a biclique}) \\ &\subseteq B \qquad (\text{since } B = \{\text{bicliques}\} \cup C_X \cup C_Y). \end{aligned}$$

Thus, we have proved $F^+ \in B$ in Case 1.

Let us now consider Case 2. In this case, we have $F \in C_Y$. Hence, $F \subseteq Y$ (since C_Y is a complex with ground set Y). Thus, $F \cap Y = F$. Moreover, F is nonempty (as we have already shown in the proof of Claim 1). But $f(F)$ is an X-neighbor of $F \cap Y$. In other words, $f(F)$ is an X-neighbor of F (since $F \cap Y = F$). In other words, $f(F) \in X$ and $\{f(F)\} \mathbf{R} F$ (by the definition of an X-neighbor). From $f(F) \in X$, we obtain $\{f(F)\} \subseteq X$. Hence, $\{f(F)\} \cup F$ is a biclique (by the definition of "biclique", since $\{f(F)\} \subseteq X$ and $F \subseteq Y$ are nonempty subsets satisfying $\{f(F)\}\ \mathbf{R}\ F$). Now,

$$\begin{aligned} F^{+} &= F \cup \{f(F)\} = \{f(F)\} \cup F \\ &\in \{\text{bicliques}\} \qquad (\text{since } \{f(F)\} \cup F \text{ is a biclique}) \\ &\subseteq B \qquad (\text{since } B = \{\text{bicliques}\} \cup C_X \cup C_Y). \end{aligned}$$

Thus, we have proved $F^{+} \in B$ in Case 2.

Now we have proved $F^{+} \in B$ in both Cases 1 and 2. Thus, $F^{+} \in B$ always holds. Combined with $F^{+} \notin C_X$, this yields

$$F^{+} \in B \setminus C_X = M.$$

The definition of f shows that $f(F^{+})$ is the largest X-neighbor of the set $F^{+} \cap Y$, whereas $f(F)$ is the largest X-neighbor of the set $F \cap Y$. Since $F^{+} \cap Y = F \cap Y$, these two descriptions of $f(F^{+})$ and $f(F)$ are identical, so we conclude that $f(F^{+}) = f(F)$. Similarly, from $F^{-} \cap Y = F \cap Y$, we obtain $f(F^{-}) = f(F)$. Combining this with $f(F^{+}) = f(F)$, we obtain $f(F^{+}) = f(F^{-}) = f(F)$.

Altogether, we have now shown that $F^{+} \in M$ and $F^{-} \in M$ and $f(F^{+}) = f(F^{-}) = f(F)$. In view of $F^{+} = F \cup \{f(F)\}$ and $F^{-} = F \setminus \{f(F)\}$, we can rewrite this as follows: We have $F \cup \{f(F)\} \in M$ and $F \setminus \{f(F)\} \in M$ and

$$f(F \cup \{f(F)\}) = f(F \setminus \{f(F)\}) = f(F).$$

Thus, Claim 2 is proved. □

Claim 3: For any $F \in M$ and any $G \in M$ satisfying $G \subseteq F$, we have $f(F) \leq f(G)$.

Proof of Claim 3. Let $F \in M$ and $G \in M$ satisfy $G \subseteq F$. From $G \subseteq F$, we obtain $G \cap Y \subseteq F \cap Y$.

Recall that $f(F)$ is the largest X-neighbor of the set $F \cap Y$. Thus, $f(F)$ is an X-neighbor of $F \cap Y$. Hence, $f(F)$ is an X-neighbor of $G \cap Y$ as well (since $G \cap Y \subseteq F \cap Y$).

But $f(G)$ is the **largest** X-neighbor of the set $G \cap Y$ (by the definition of f). Hence, $f(G) \geq x$ whenever x is any X-neighbor of $G \cap Y$. Applying this to $x = f(F)$, we obtain $f(G) \geq f(F)$ (since $f(F)$ is an X-neighbor of $G \cap Y$). In other words, $f(F) \leq f(G)$. This proves Claim 3. □

Now, we can apply both parts **(a)** and **(b)** of Lemma 4.10 to $\Delta = B$ and $W = X \cup Y$ (since Claim 2 shows that condition **(C1)** in Lemma 4.10 holds, whereas Claim 3 shows that condition **(C2)** holds). Thus, we conclude that the map

$$\begin{aligned} \mu : M &\to M, \\ F &\mapsto \begin{cases} F \setminus \{f(F)\}, & \text{if } f(F) \in F; \\ F \cup \{f(F)\}, & \text{if } f(F) \notin F \end{cases} \end{aligned}$$

is well-defined (by Lemma 4.10 **(a)**), and the pair (M, μ) is a matching on B (by Lemma 4.10 **(a)** again), and that this matching (M, μ) is acyclic (by Lemma 4.10 **(b)**). Hence, there exists an acyclic matching (M, μ) on B with $M = B \setminus C_X$. Thus, Theorem 4.9 (applied to $\Delta = B$ and $\Gamma = C_X$) shows that $B \searrow C_X$. As we said, this completes the proof of Theorem 6.1. □

7. Functoriality and Dowker's theorem

In this section, we will prove Dowker's and Barmak's theorems (Theorems 3.4 and 3.5). The hard work has already been done in the preceding sections: In fact, Theorem 3.5 follows immediately from Theorem 6.1 if the sets X and Y are disjoint. The only thing that remains to be done is reducing the general case to this disjoint case. Roughly speaking, this can be done by "renaming" the elements of X and Y, after one shows (Lemma 7.1 below) that isomorphic simplicial complexes are simple-homotopy-equivalent. (This basic fact does not seem to appear explicitly in the literature, so we prove it below – using Theorem 6.1 in fact!)

We shall now present this argument in some detail. We begin by introducing some categories:

1. The *category of simplicial complexes*: The objects of this category are the simplicial complexes. Its morphisms are the *simplicial maps*, defined as follows: The *minimal ground set* of a complex Δ is defined to be the union of all faces of Δ. A *simplicial map* from a complex Δ to a complex Γ means a map f from the minimal ground set of Δ to the minimal ground set of Γ such that every face F of Δ satisfies $f(F) \in \Gamma$ (where $f(F)$ denotes $\{f(x) \mid x \in F\}$ as usual). Composition of simplicial maps is just regular composition of maps.

2. The *category of relations*: Given two relations $R \subseteq X \times Y$ and $S \subseteq Z \times W$, a *morphism of relations* $\phi : R \to S$ is a pair $\phi = (\phi_l, \phi_r)$ of maps $\phi_l : X \to Z$ and $\phi_r : Y \to W$ such that each $(x, y) \in R$ satisfies $(\phi_l(x), \phi_r(y)) \in S$. Thus, we can define the *category of relations*: a category whose objects are binary relations (more precisely, triples (X, Y, R) where $R \subseteq X \times Y$ is a relation), and whose morphisms are morphisms of relations. Composition of morphisms is defined componentwise (i.e., by setting $(\phi_l, \phi_r) \circ (\psi_l, \psi_r) = (\phi_l \circ \psi_l, \phi_r \circ \psi_r)$), and the identity morphism of a given relation $R \subseteq X \times Y$ is the pair $(\mathrm{id}_X, \mathrm{id}_Y)$.

3. The *category of bipartite relations*: Recall that a relation R from a set X to a set Y is called a bipartite relation if the sets X and Y are disjoint. The *category of bipartite relations* is the full subcategory of the category of relations whose objects are the bipartite relations.

Given a relation $R \subseteq X \times Y$, we let $C_l(R) = C_X$ and $C_r(R) = C_Y$ be the left and right Dowker complexes of R.

If R is a bipartite relation, then we furthermore let $B(R)$ be the biclique complex of R (formerly denoted B, defined in Proposition 5.3).

We have seen that $C_l(R) = C_X$ and $C_r(R) = C_Y$ are subcomplexes of $B(R) = B$. We will write $i_l^R \colon C_l(R) \to B(R)$ and $i_r^R \colon C_r(R) \to B(R)$ for the respective inclusion maps.

Given a morphism $\phi = (\phi_l, \phi_r) \colon R \to S$ of relations, the maps ϕ_l and ϕ_r induce simplicial maps $C_l(\phi) \colon C_l(R) \to C_l(S)$ and $C_r(\phi) \colon C_r(R) \to C_r(S)$ of simplicial complexes. (For instance, $C_l(\phi)$ replaces each vertex of $C_l(R)$ by its image under ϕ_l.) The assignments $R \mapsto C_l(R)$ and $R \mapsto C_r(R)$ thus define functors from the category of relations to the category of simplicial complexes. If R and S are bipartite relations, then the morphism $\phi_l \cup \phi_r$ induces a simplicial map $B(\phi) \colon B(R) \to B(S)$ of simplicial complexes. This makes $R \mapsto B(R)$ a functor from the category of bipartite relations to the category of simplicial complexes. The inclusion maps $i_l^R \colon C_l(R) \to B(R)$ and $i_r^R \colon C_r(R) \to B(R)$ then form natural transformations of functors.

As a last preparation for the proof of Theorem 3.5, we now show that isomorphic simplicial complexes are simple-homotopy-equivalent.

Lemma 7.1. Let Δ and Δ' be isomorphic simplicial complexes. Then Δ and Δ' are simple-homotopy-equivalent.

Proof. We WLOG assume that Δ and Δ' are nonempty (since otherwise, $\Delta = \Delta'$, and the claim is obvious). Hence, the empty set $\varnothing$ is both a face of Δ and a face of Δ'.

Let X and X' be the minimal ground sets of Δ and Δ'. Let $\alpha : X \to X'$ be the isomorphism from Δ to Δ'. Let $\alpha_* : \Delta \to \Delta'$ be the map that sends each face F of Δ to the face $\alpha(F) = \{\alpha(f) \mid f \in F\}$ of Δ'. This map α_* is a bijection (since α is an isomorphism of complexes).

We refer to the elements of X as the *vertices* of Δ. Likewise for X' and Δ'.

Pick a set Y that is disjoint from both X and X' and has the same size as Δ. Pick a bijection $f \colon Y \to \Delta$. Thus, f assigns a face $f(y)$ of Δ to each $y \in Y$. Let $f' : Y \to \Delta'$ be the composition $\alpha_* \circ f$ of the bijections $f : Y \to \Delta$ and $\alpha_* : \Delta \to \Delta'$. Thus, f' is itself a bijection (since a composition of two bijections is a bijection), and sends each $y \in Y$ to $\alpha_*(f(y)) = \alpha(f(y))$ (by the definition of α_*).

Now, define a relation R from X to Y by

$$R = \{(x, y) \in X \times Y \mid x \in f(y)\}.$$

Thus, a pair $(x, y) \in X \times Y$ belongs to R if and only if $x \in f(y)$. In other words, R is the containment relation between the vertices and the faces of Δ, except that the faces have been relabelled using f.

The definition of the left Dowker complex $C_l(R)$ shows that

$$\begin{aligned}
C_l(R) &= \{Y\text{-conic sets with respect to the relation } R\} \\
&= \{U \subseteq X \mid U \text{ is empty or has a } Y\text{-neighbor}\} \\
&= \{U \subseteq X \mid U \text{ is empty or there exists some } y \in Y \\
&\qquad \text{such that all } u \in U \text{ satisfy } u \; R \; y\} \\
&\qquad \text{(by the definition of a } Y\text{-neighbor)} \\
&= \{U \subseteq X \mid U \text{ is empty or there exists some } y \in Y \\
&\qquad \text{such that all } u \in U \text{ satisfy } u \in f(y)\} \\
&\qquad \text{(by the definition of } R) \\
&= \{U \subseteq X \mid U \text{ is empty or there exists some } y \in Y \\
&\qquad \text{such that } U \subseteq f(y)\} \\
&= \{U \subseteq X \mid U \text{ is empty or there exists some } F \in \Delta \\
&\qquad \text{such that } U \subseteq F\} \qquad (\text{since } f : Y \to \Delta \text{ is a bijection}) \\
&= \{U \subseteq X \mid U \text{ is empty or a subset of some face } F \in \Delta\} \\
&= \Delta
\end{aligned}$$

(since $\varnothing$ is a face of Δ, and since any subset of a face of Δ is itself a face of Δ).

Next, define a relation R' from X' to Y by

$$R' = \left\{ \left(x', y\right) \in X' \times Y \mid x' \in f'(y) \right\}.$$

Thus, a pair $(x', y) \in X' \times Y$ belongs to R' if and only if $x' \in f'(y)$. In other words, R' is the containment relation between the vertices and the faces of Δ', except that the faces have been relabelled using f'.

We have already proved that $C_l(R) = \Delta$. An analogous argument (using X', Δ', R' and f' instead of X, Δ, R and f) shows that $C_l(R') = \Delta'$.

Next, we shall show that $C_r(R) = C_r(R')$. Indeed, this follows easily from abstract nonsense: The pair $\phi := (\alpha, \mathrm{id}_Y)$ is an isomorphism of relations from R to R' (this follows easily from the definitions of R and R', since $f' = \alpha_* \circ f$ and since α is a bijection[1]). Hence, it induces an isomorphism of complexes $C_r(\phi) : C_r(R) \to C_r(R')$. Recalling the definition of $C_r(\phi)$, we see that this isomorphism $C_r(\phi)$ is simply the identity map (since it replaces each vertex of $C_r(R)$ by its image under

[1] In more detail: We need to show that two elements $x \in X$ and $y \in Y$ satisfy $(x, y) \in R$ if and only if they satisfy $(\alpha(x), \mathrm{id}_Y(y)) \in R'$. But this follows from the equivalences

$$\begin{aligned}
((x,y) \in R) &\iff (x \in f(y)) \qquad \text{(by the definition of } R) \\
&\iff (\alpha(x) \in \alpha(f(y))) \qquad (\text{since } \alpha \text{ is a bijection}) \\
&\iff \left(\alpha(x) \in f'(y)\right) \qquad \left(\text{since } \alpha(f(y)) = \alpha_*(f(y)) = f'(y)\right) \\
&\iff \left((\alpha(x), y) \in R'\right) \qquad \text{(by the definition of } R') \\
&\iff \left((\alpha(x), \mathrm{id}_Y(y)) \in R'\right) \qquad (\text{since } y = \mathrm{id}_Y(y)).
\end{aligned}$$

id_Y, which is exactly the same vertex). Thus, the identity map is an isomorphism of complexes from $C_r(R)$ to $C_r(R')$. This shows that

$$C_r(R) = C_r(R').$$

But R is a bipartite relation (since the sets X and Y are disjoint). Thus, Theorem 6.1 shows that its biclique complex $B(R)$ collapses to both $C_l(R)$ and $C_r(R)$. Therefore, the simplicial complexes $C_l(R)$ and $C_r(R)$ are simple-homotopy-equivalent. In other words, $C_l(R) \sim C_r(R)$, where the symbol $\sim$ means simple-homotopy-equivalence. Likewise, we obtain $C_l(R') \sim C_r(R')$ (since R' is also a bipartite relation). Since simple-homotopy-equivalence is an equivalence relation, we can combine these results to

$$\Delta = C_l(R) \sim C_r(R) = C_r(R') \sim C_l(R') = \Delta'.$$

In other words, Δ and Δ' are simple-homotopy-equivalent. □

Another proof of Lemma 7.1 can be obtained along the lines of [BarMin12, Remark 2.3] (which proves the analogous claim for **strong** homotopy equivalence). Indeed, each of the elementary strong collapses as defined in [BarMin12, Definition 2.1] can be decomposed into a sequence of elementary simplicial collapses, and thus preserves the simple-homotopy-type.

Now we can easily obtain Theorem 3.5:

Proof of Theorem 3.5. If the sets X and Y are disjoint, then Theorem 6.1 shows that the biclique complex B of R collapses to both C_X and C_Y, and therefore the complexes C_X and C_Y are simple-homotopy-equivalent. This proves Theorem 3.5 in the case when the sets X and Y are disjoint.

Now let us consider the general case. We shall reduce this case to the disjoint case by replacing X and Y with two disjoint copies.

Namely, define the two finite sets $\widetilde{X} = X \times \{0\}$ and $\widetilde{Y} = Y \times \{1\}$, which are clearly disjoint. Let $\widetilde{R}$ be the subset of $\widetilde{X} \times \widetilde{Y}$ consisting of all pairs of the form $((x,0), (y,1))$ with $(x,y) \in R$. In other words, $\widetilde{R}$ is the binary relation from $\widetilde{X}$ to $\widetilde{Y}$ for which two elements $(x,0) \in \widetilde{X}$ and $(y,1) \in \widetilde{Y}$ satisfy $(x,0)\ \widetilde{R}\ (y,1)$ if and only if $x\ R\ y$. Informally, this relation $\widetilde{R}$ is simply "the relation R after each $x \in X$ has been renamed as $(x,0)$ and each $y \in Y$ has been renamed as $(y,1)$". Rigorously, this means that the relations R and $\widetilde{R}$ are isomorphic; an explicit isomorphism from R to $\widetilde{R}$ is the pair $\phi := (\phi_l, \phi_r)$, where $\phi_l : X \to \widetilde{X}$ is the map sending each x to $(x,0)$, and where $\phi_r : Y \to \widetilde{Y}$ is the map sending each y to $(y,1)$. This isomorphism induces isomorphisms $C_l(\phi)\colon C_l(R) \to C_l(\widetilde{R})$ and $C_r(\phi)\colon C_r(R) \to C_r(\widetilde{R})$. Since isomorphic simplicial complexes are simple-homotopy-equivalent (by Lemma 7.1), this entails that $C_l(R) \sim C_l(\widetilde{R})$ and $C_r(R) \sim C_r(\widetilde{R})$, where the symbol $\sim$ means simple-homotopy-equivalence. But the sets $\widetilde{X}$ and $\widetilde{Y}$ are disjoint, and thus $\widetilde{R}$ is a bipartite relation. Hence, we can apply Theorem 3.5 to $\widetilde{X}$, $\widetilde{Y}$ and $\widetilde{R}$ instead of X, Y and R (since we have already proved Theorem 3.5 in the case when the sets X and Y

are disjoint). We conclude that $C_l(\widetilde{R}) \sim C_r(\widetilde{R})$. Since simple-homotopy-equivalence is an equivalence relation, we can now combine our results:

$$C_l(R) \sim C_l(\widetilde{R}) \sim C_r(\widetilde{R}) \sim C_r(R).$$

This proves Theorem 3.5 in the general case. □

8. Further directions

Dowker's paper [Dowker52] has motivated much research in combinatorial and applied topology over the 70 years since its publication (with 292 citations as of 2024). Some of it generalizes and extends its results to other settings, or approaches their proofs in different ways. In this final section, we shall connect two such developments with our Morse-theoretic approach.

8.1. On the Brun–Salbu rectangle complex

Brun and Salbu [BruSal22] have recently given a new proof of Theorem 3.4, using projections instead of inclusions. Their proof relies on the following notions (where R is a bipartite relation from X to Y):

- A *rectangle* of R means a subset of R having the form $U \times V$ for some $U \subseteq X$ and some $V \subseteq Y$.
- The *rectangle complex* E (called $E(R)$ in [BruSal22]) is defined to be the complex with ground set $X \times Y$ (not $X \cup Y$) whose faces are the subsets of all rectangles of R.

Brun and Salbu show ([BruSal22, Theorem 4.3]) that this rectangle complex E is homotopy-equivalent to both C_X and C_Y, but not via simplicial embeddings like for B, but rather via simplicial projections. Namely, the simplicial projection maps (defined by their action on the ground sets)

$$\pi_X : E \to C_X,$$
$$(x, y) \mapsto x$$

and

$$\pi_Y : E \to C_Y,$$
$$(x, y) \mapsto y$$

are shown to be homotopy equivalences (using Quillen's fiber lemma), so that $C_X \simeq E \simeq C_Y$.

This result, too, can be proved using discrete Morse theory, thus strengthening it to a simple-homotopy equivalence:

Theorem 8.1. The complexes E, B, C_X and C_Y are all simple-homotopy-equivalent.

However, the only proof of this theorem we have found so far is neither as elementary nor as simple as the above proof of Theorem 3.5. It relies on the equivalence between simplicial complexes and posets (via the barycentric subdivision – see [Kozlov20, proof of Proposition 9.29]) and on simple-homotopy analogues of Quillen's fiber lemma and the nerve theorem (cf. [Barmak11]). We intend to elaborate on it in forthcoming work.

8.2. Strong homotopy

In [DocSin21, Proposition 4.4], Dochtermann and Singh construct a strong homotopy equivalence (see [BarMin12] for the definition) between the face posets of C_X and C_Y. This entails that the order complexes of these face posets are simple-homotopy-equivalent. As these order complexes are the barycentric subdivisions of C_X and C_Y, this leads once again to the simple-homotopy-equivalence of C_X and C_Y, that is, to Barmak's Theorem 3.5.

However, the complexes C_X and C_Y themselves do not have the same strong homotopy type in general. Indeed, if they did, then [BarMin12, Theorem 2.11] would yield that their cores are isomorphic; but this cannot be the case for most relations R. (For a specific example, let $X = \{1, 2, 3, 4\}$ and $Y = \{12, 13, 14, 23, 24, 34\}$, and let R be given by $i\ R\ j$ if i is one of the two digits of j. Then, both complexes C_X and C_Y are their own cores, and clearly not isomorphic.)

8.3. The relative case

While Theorem 3.4 is commonly known as Dowker's theorem nowadays, Dowker actually proved a different (if related) result in [Dowker52, Theorems 1 and 1a]. Indeed, his result is simultaneously weaker (making only homological rather than homotopical statements) and stronger than Theorem 3.4. It is stronger, as it concerns itself with a more general situation: a "pair of relations" (a relation R_1 with a subrelation R_2, or, to use our functorial language, an injective morphism of relations $i : R_2 \to R_1$) rather than a single relation. The claim is then an isomorphism of relative (co)homology groups between the respective Dowker complexes.

Question 8.2. Can this generalized claim be proved using Morse-theoretical methods?

This question is not quite straightforward, as discrete Morse theory in relative homotopy is much less well-trodden than in the absolute setting, and our acyclic matching constructed implicitly in the proof of Theorem 6.1 is not adapted to a subrelation.

8.4. Generalized Dowker Duality

The idea of the Dowker nerve has been applied to concepts of relations between objects with more structure than discrete sets.

For example in [BrFoSa23], the authors define a Dowker nerve for relations between categories, and prove a Dowker duality theorem for these. The paper [FerMin20] introduces the cylinder of a relation between partially ordered sets.

Another direction is introduced in the paper [FerMin20]. Given a relation R beween two sets X and Y, we obtain a relation $D(R)$ between the left Dowker nerve $C_l(R) = C_X$ and the right Dowker nerve $C_r(R) = C_Y$ consisting of the pairs $(\sigma, \tau) \in C_l(R) \times C_r(R)$ of simplices such that $\sigma \times \tau \subseteq R$. By [FerMin20, Theorem 2.6], the cylinder of $D(R)$ is simple-homotopy-equivalent to the face posets of C_X and C_Y. Since the barycentric subdivision of a simplicial complex is simple-homotopy-equivalent to the complex itself, this implies that order complex of $D(R)$ is simple-homotopy-equivalent to C_X and C_Y.

Question 8.3. Can the categorical version of Dowker's theorem [BrFoSa23] be proved using discrete Morse theory for simplicial sets?

This would require a discrete Morse theory to be introduced for simplicial sets in the first place. Such theories have been proposed for **semi**simplicial sets previously, including in Brown's original paper [Brown92] that laid the ground for discrete Morse theory.

Acknowledgments

The second author would like to thank Dmitry Kozlov and Tom Roby for helpful and interesting conversations.

References

[Barmak11] Jonathan Ariel Barmak, *On Quillen's Theorem A for posets*, Journal of Combinatorial Theory, Series A **118** (2011), pp. 2445–2453.

[BarMin12] Jonathan Ariel Barmak, Elias Gabriel Minian, *Strong Homotopy Types, Nerves and Collapses*, Discrete & Computational Geometry **47** (2012), pp. 301–328.

[Bjorne95] Anders Björner, *Topological Methods*, Chapter 34 of: R. Graham, M. Grötschel, L. Lovász (eds), *Handbook of Combinatorics*, Elsevier 1995.

[BrFoSa23] Morten Brun, Marius Gårdsmann Fosse, Lars M. Salbu, *Dowker Duality for Relations of Categories*, arXiv:2303.16032v1.

[Brown92] Kenneth S. Brown, *The Geometry of Rewriting Systems: A Proof of the Anick-Groves-Squier Theorem*. In: Algorithms and Classification in Combinatorial Group Theory, Springer, 1992, pp. 137–163. `https://pi.math.cornell.edu/~kbrown/scan/1992.0000.0137.pdf`

[BruSal22] Morten Brun, Lars M. Salbu, *The Rectangle Complex of a Relation*, Mediterranean Journal of Mathematics **20** (2023), article 7. See arXiv:2207.02018v2 for a preprint.

[DocSin21] Anton Dochtermann, Anurag Singh, *Homomorphism complexes, reconfiguration, and homotopy for directed graphs*, arXiv:2108.10948v2, published in: European J. Combin. 110 (2023).

[Dowker52] C. H. Dowker, *Homology Groups of Relations*, Annals of Mathematics, Second Series **56**, No. 1 (Jul., 1952), pp. 84–95.

[FerMin20] Ximena Fernández and Elías Gabriel Minian, *The cylinder of a relation and generalized versions of the nerve theorem*, Discrete & Computational Geometry **63**, (2020).

[GrKaLe21] Darij Grinberg, Lukas Katthän, Joel Brewster Lewis, *The path-missing and path-free complexes of a directed graph*, arXiv:2102.07894v2.

[Kozlov20] Dmitry N. Kozlov, *Organized Collapse: An Introduction to Discrete Morse Theory*, Graduate Studies in Mathematics **207**, AMS 2020. Available at `https://dfklab.com/books/` .

[LinSha03] Svante Linusson, John Shareshian, *Complexes of t-colorable graphs*, SIAM J. Discrete Math. **16** (2003), no. 3, pp. 371–389.

[Sagan19] Bruce Sagan, *Combinatorics: The Art of Counting*, Graduate Studies in Mathematics **210**, AMS 2020. `https://users.math.msu.edu/users/bsagan/Books/Aoc/final.pdf` See `https://users.math.msu.edu/users/bsagan/Books/Aoc/errata.pdf` for errata.

[Yoon24] Iris H. R. Yoon, *Dowker duality, profunctors, and spectral sequences*, arXiv:2408.13136v2.