Mathematisches
Forschungsinstitut
Oberwolfach

Member of the
Leibniz
Association

Oberwolfach Preprints

OWP 2021 - 05

DARIJ GRINBERG

The Elser Nuclei Sum Revisited

Mathematisches Forschungsinstitut Oberwolfach gGmbH
Oberwolfach Preprints (OWP) ISSN 1864-7596

Oberwolfach Preprints (OWP)

The MFO publishes a preprint series **Oberwolfach Preprints (OWP)**, ISSN 1864-7596, which mainly contains research results related to a longer stay in Oberwolfach, as a documentation of the research work done at the MFO. In particular, this concerns the Research in Pairs-Programme (RiP) and the Oberwolfach-Leibniz-Fellows (OWLF), but this can also include an Oberwolfach Lecture, for example.

A preprint can have a size from 1 - 200 pages, and the MFO will publish it on its website as well as by hard copy. Every RiP group or Oberwolfach-Leibniz-Fellow may receive on request 20 free hard copies (DIN A4, black and white copy) by surface mail.

The full copyright is left to the authors. With the submission of a manuscript, the authors warrant that they are the creators of the work, including all graphics. The authors grant the MFO a perpetual, irrevocable, non-exclusive right to publish it on the MFO's institutional repository.

In case of interest, please send a **pdf file** of your preprint by email to *rip@mfo.de* or *owlf@mfo.de*, respectively. The file should be sent to the MFO within 12 months after your stay as RiP or OWLF at the MFO.

There are no requirements for the format of the preprint, except that the introduction should contain a short appreciation and that the paper size (respectively format) should be DIN A4, "letter" or "article".

On the front page of the hard copies, which contains the logo of the MFO, title and authors, we shall add a running number (20XX – XX). Additionally, each preprint will get a Digital Object Identifier (DOI).

We cordially invite the researchers within the RiP or OWLF programme to make use of this offer and would like to thank you in advance for your cooperation.

Imprint:

Mathematisches Forschungsinstitut Oberwolfach gGmbH (MFO)
Schwarzwaldstrasse 9-11
77709 Oberwolfach-Walke
Germany

Tel	+49 7834 979 50
Fax	+49 7834 979 55
Email	admin@mfo.de
URL	www.mfo.de

The Oberwolfach Preprints (OWP, ISSN 1864-7596) are published by the MFO.

DOI	10.14760/OWP-2021-05

The Elser nuclei sum revisited

Version of March 21, 2021

(arXiv:2009.11527v6)

Darij Grinberg

Drexel University

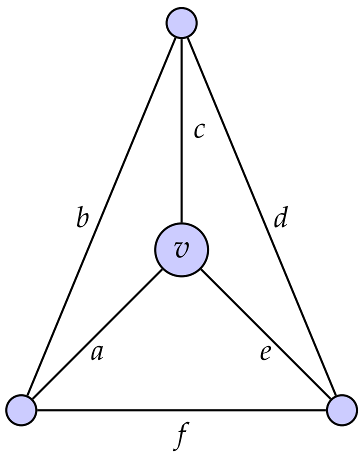

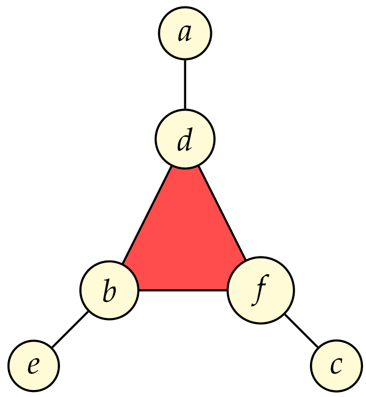

Oberwolfach Preprints, 2021

The Elser nuclei sum revisited.

Version of 21 March 2021,
equivalent to `arXiv:2009.11527v6`.

DARIJ GRINBERG:
Drexel University
Korman Center, Room 263
15 S 33rd Street
Philadelphia PA, 19104
USA
(temporary:)
Mathematisches Forschungsinstitut Oberwolfach,
Schwarzwaldstrasse 9–11
77709 Oberwolfach
Germany
`darijgrinberg@gmail.com`
`https://www.cip.ifi.lmu.de/~grinberg/`

Abstract. Fix a finite undirected graph Γ and a vertex v of Γ. Let E be the set of edges of Γ. We call a subset F of E *pandemic* if each edge of Γ has at least one endpoint that can be connected to v by an F-path (i.e., a path using edges from F only). In 1984, Elser showed that the sum of $(-1)^{|F|}$ over all pandemic subsets F of E is 0 if $E \neq \varnothing$. We give a simple proof of this result via a sign-reversing involution, and discuss variants, generalizations and a refinement using discrete Morse theory.

Keywords: graph theory, nuclei, simplicial complex, discrete Morse theory, alternating sum, enumerative combinatorics, inclusion/exclusion, convexity.

Future versions of this text will be available from the first author's website:
`http://www.cip.ifi.lmu.de/~grinberg/algebra/elsersum.pdf`
`http://www.cip.ifi.lmu.de/~grinberg/algebra/elsersum-long.pdf` (detailed version).

****

In [Elser84], Veit Elser studied the probabilities of clusters forming when n points are sampled randomly in a d-dimensional volume. In the process, he found a purely graph-theoretical lemma [Elser84, Lemma 1], which served a crucial role in his work. For decades, the lemma stayed hidden from the eyes of combinatorialists in a physics journal, until it resurfaced in recent work [DHLetc19] by Dorpalen-Barry, Hettle, Livingston, Martin, Nasr, Vega and Whitlatch. In this note, I will show a simpler proof of the lemma that illustrates the use of sign-reversing involutions and also suggests a generalization. I will then discuss a strengthening of the lemma as well as some open questions.

Note to the reader

The pictures on the title page illustrate the simplicial complex $\mathcal{A}$ from Proposition 5.2 on an example. The left picture is a graph Γ (with the vertex labelled v playing the role of v), whereas the right picture shows the corresponding simplicial complex $\mathcal{A}$ for $G = E$ (that is, the simplicial complex whose faces are the subsets of E that are not pandemic).

Much of this text has been conceived and written during a stay at the Mathematisches Forschungsinstitut Oberwolfach in 2020. This research was supported through the programme "Oberwolfach Leibniz Fellows" by the Mathematisches Forschungsinstitut Oberwolfach in 2020.

Remark on alternative versions

This paper also has a detailed version [Grinbe20], which elaborates on the proofs.

1. Elser's result

Let us first introduce our setting, which is slightly more general (and perhaps also simpler) than that used in [Elser84].

We fix an arbitrary graph Γ with vertex set V and edge set E. Here, "graph" means "finite undirected multigraph" – i.e., it can have self-loops and parallel edges, but it has finitely many vertices and edges, and its edges are undirected.

We fix a vertex $v \in V$.

If $F \subseteq E$, then an *F-path* shall mean a path of Γ such that all edges of the path belong to F.

If $e \in E$ is any edge and $F \subseteq E$ is any subset, then we say that F *infects* e if there exists an F-path from v to some endpoint of e. (The terminology is inspired by the idea of an infectious disease starting in the vertex v and being transmitted along edges.)[1]

A subset $F \subseteq E$ is said to be *pandemic* if it infects each edge $e \in E$.

Example 1.1. Let Γ be the following graph:

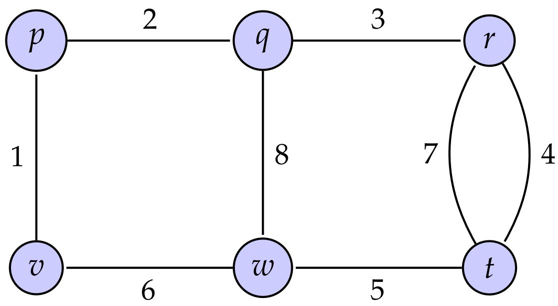

(where the vertex v is the vertex labelled v). Then, for example, the set $\{1,2\} \subseteq E$ infects edges $1,2,3,6,8$ (but none of the other edges). The set $\{1,2,5\}$ infects the same edges as $\{1,2\}$ (indeed, the additional edge 5 does not increase its infectiousness, since it is not on any $\{1,2,5\}$-path from v). The set $\{1,2,3\}$ infects every edge other than 5. The set $\{1,2,3,4\}$ infects each edge, and thus is pandemic.

Now, we can state our version of [Elser84, Lemma 1]:

Theorem 1.2. Assume that $E \neq \varnothing$. Then,

$$\sum_{\substack{F \subseteq E \text{ is} \\ \text{pandemic}}} (-1)^{|F|} = 0. \qquad (1)$$

Example 1.3. Let Γ be the following graph:

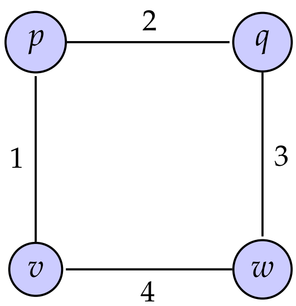

[1]Note that if an edge e contains the vertex v, then any subset F of E (even the empty one) infects e, since there is a trivial (edgeless) F-path from v to v.

(where the vertex v is the vertex labelled v). Then, the pandemic subsets of E are the sets

$$\{1,2\}, \quad \{1,4\}, \quad \{3,4\}, \quad \{1,2,3\}, \quad \{1,3,4\}, \quad \{1,2,4\}, \quad \{2,3,4\}, \quad \{1,2,3,4\}.$$

The sizes of these subsets are $2,2,2,3,3,3,3,4$, respectively. Hence, (1) says that

$$(-1)^2 + (-1)^2 + (-1)^2 + (-1)^3 + (-1)^3 + (-1)^3 + (-1)^3 + (-1)^4 = 0.$$

We note that the equality (1) can be restated as "there are equally many pandemic subsets $F \subseteq E$ of even size and pandemic subsets $F \subseteq E$ of odd size". Thus, in particular, the number of all pandemic subsets F of E is even (when $E \neq \varnothing$).

Remark 1.4. Theorem 1.2 is a bit more general than [Elser84, Lemma 1]. To see why, we assume that the graph Γ is connected and simple (i.e., has no self-loops and parallel edges). Then, a *nucleus* is defined in [Elser84] as a subgraph N of Γ with the properties that

1. the subgraph N is connected, and
2. each edge of Γ has at least one endpoint in N.

Given a subgraph N of Γ, we let $\mathrm{E}(N)$ denote the set of all edges of N. Now, [Elser84, Lemma 1] claims that if $E \neq \varnothing$, then

$$\sum_{\substack{N \text{ is a nucleus} \\ \text{containing } v}} (-1)^{|\mathrm{E}(N)|} = 0.$$

But this is equivalent to (1), because there is a bijection

$$\begin{aligned} \{\text{nuclei containing } v\} &\to \{\text{pandemic subsets } F \subseteq E\}, \\ N &\mapsto \mathrm{E}(N). \end{aligned}$$

We leave it to the reader to check this in detail; what needs to be checked are the following three statements:

- If N is a nucleus containing v, then $\mathrm{E}(N)$ is a pandemic subset of E.
- Every nucleus N containing v is uniquely determined by the set $\mathrm{E}(N)$. (Indeed, since a nucleus has to be connected, each of its vertices must be an endpoint of one of its edges, unless its only vertex is v.)

- If F is a pandemic subset of E, then there is a nucleus N containing v such that $\mathrm{E}(N) = F$. (Indeed, N can be defined as the subgraph of Γ whose vertices are the endpoints of all edges in F as well as the vertex v, and whose edges are the edges in F. To see that this subgraph N is connected, it suffices to argue that each of its vertices has a path to v; but this follows from the definition of "pandemic", since each vertex of N other than v belongs to at least one edge in F.)

Thus, Theorem 1.2 is equivalent to [Elser84, Lemma 1] in the case when Γ is connected and simple.

Remark 1.5. It might appear more natural to talk about a subset $F \subseteq E$ infecting a vertex rather than an edge. (Namely, we can say that F infects a vertex w if there is an F-path from v to w.) However, the analogue of Theorem 1.2 in which pandemicity is defined via infecting all vertices is not true. The graph of Example 1.3 provides a counterexample.

2. The proof

2.1. Shades

Our proof of Theorem 1.2 will rest on a few notions. The first is that of a *shade*:

Definition 2.1. Let F be a subset of E. Then, we define a subset $\operatorname{Shade} F$ of E by

$$\operatorname{Shade} F = \{e \in E \mid F \text{ infects } e\}. \tag{2}$$

We refer to $\operatorname{Shade} F$ as the *shade* of F.

Thus, the shade of a subset $F \subseteq E$ is the set of all edges of Γ that are infected by F.

Example 2.2. In Example 1.1, we have $\operatorname{Shade}\{1,2\} = \{1,2,3,6,8\}$ and $\operatorname{Shade}\{1\} = \{1,2,6\}$ and $\operatorname{Shade}\{8\} = \{1,6\}$.

The following property of shades is rather obvious:

Lemma 2.3. Let A and B be two subsets of E such that $A \subseteq B$. Then, $\operatorname{Shade} A \subseteq \operatorname{Shade} B$.

Proof of Lemma 2.3. We must show that each $q \in \operatorname{Shade} A$ satisfies $q \in \operatorname{Shade} B$. In other words, we must prove that if A infects some edge $q \in E$, then B also infects this edge q. But this is clear, since any A-path is a B-path. $\square$

The major property of shades that we will need is the following:

Lemma 2.4. Let F be a subset of E. Let $u \in E$ be such that $u \notin \operatorname{Shade} F$. Then,

$$\operatorname{Shade}(F \cup \{u\}) = \operatorname{Shade} F \tag{3}$$

and

$$\operatorname{Shade}(F \setminus \{u\}) = \operatorname{Shade} F. \tag{4}$$

Proof of Lemma 2.4. We shall prove (3) and (4) separately:

[*Proof of (3):* Let $q \in \operatorname{Shade}(F \cup \{u\})$. We shall show that $q \in \operatorname{Shade} F$.

We have assumed that $q \in \operatorname{Shade}(F \cup \{u\})$. In other words, q is an edge in E with the property that $F \cup \{u\}$ infects q (by the definition of $\operatorname{Shade}(F \cup \{u\})$).

We shall now show that F infects q. Indeed, assume the contrary. Thus, F does not infect q. In other words, there exists no F-path from v to any endpoint of q (by the definition of "infects").

We know that $F \cup \{u\}$ infects q. In other words, there exists an $(F \cup \{u\})$-path from v to some endpoint of q (by the definition of "infects"). Let π be this path. If this $(F \cup \{u\})$-path π did not contain the edge u, then it would be an F-path, which would contradict the fact that there exists no F-path from v to any endpoint of q. Hence, this $(F \cup \{u\})$-path π must contain the edge u. By removing u, we can thus cut this path π into two segments: The first segment is a path from v to some endpoint of u, while the second segment is a path from the other endpoint of u to some endpoint of q. Both segments are F-paths. Thus, in particular, the first segment is an F-path from v to some endpoint of u. Hence, there exists an F-path from v to some endpoint of u. In other words, F infects u (by the definition of "infects"). Hence, $u \in \operatorname{Shade} F$ (because of (2)). This contradicts $u \notin \operatorname{Shade} F$.

This contradiction shows that our assumption was false. Hence, we have proved that F infects q. In other words $q \in \operatorname{Shade} F$.

Forget that we fixed q. We thus have shown that $q \in \operatorname{Shade} F$ for each $q \in \operatorname{Shade}(F \cup \{u\})$. In other words, $\operatorname{Shade}(F \cup \{u\}) \subseteq \operatorname{Shade} F$. On the other hand, $F \subseteq F \cup \{u\}$; therefore, $\operatorname{Shade} F \subseteq \operatorname{Shade}(F \cup \{u\})$ (by Lemma 2.3). Combining this with $\operatorname{Shade}(F \cup \{u\}) \subseteq \operatorname{Shade} F$, we obtain $\operatorname{Shade}(F \cup \{u\}) = \operatorname{Shade} F$. This proves (3).]

[*Proof of (4):* We must prove that $\operatorname{Shade}(F \setminus \{u\}) = \operatorname{Shade} F$. This is obvious if $F \setminus \{u\} = F$. Thus, for the rest of this proof, we WLOG assume that $F \setminus \{u\} \neq F$. Hence, $u \in F$ and thus $(F \setminus \{u\}) \cup \{u\} = F$.

We have $F \setminus \{u\} \subseteq F$ and thus $\operatorname{Shade}(F \setminus \{u\}) \subseteq \operatorname{Shade} F$ (by Lemma 2.3). Hence, from $u \notin \operatorname{Shade} F$, we obtain $u \notin \operatorname{Shade}(F \setminus \{u\})$. Therefore, (3) (applied to $F \setminus \{u\}$ instead of F) yields $\operatorname{Shade}((F \setminus \{u\}) \cup \{u\}) = \operatorname{Shade}(F \setminus \{u\})$. Thus, $\operatorname{Shade}(F \setminus \{u\}) = \operatorname{Shade} \underbrace{((F \setminus \{u\}) \cup \{u\})}_{=F} = \operatorname{Shade} F$. This proves (4).]

We have now proved both (3) and (4). Thus, Lemma 2.4 is proved. $\square$

2.2. A slightly more general claim

Lemma 2.4 might not look very powerful, but it contains all we need to prove Theorem 1.2. Better yet, we shall prove the following slightly more general version of Theorem 1.2:

Theorem 2.5. Let G be any subset of E. Assume that $E \neq \varnothing$. Then,

$$\sum_{\substack{F\subseteq E;\\ G\subseteq \operatorname{Shade} F}} (-1)^{|F|} = 0.$$

We will soon prove Theorem 2.5 and explain how Theorem 1.2 follows from it. First, however, let us give an equivalent (but slightly easier to prove) version of Theorem 2.5:

Theorem 2.6. Let G be any subset of E. Then,

$$\sum_{\substack{F\subseteq E;\\ G\not\subseteq \operatorname{Shade} F}} (-1)^{|F|} = 0.$$

Proof of Theorem 2.6. Let

$$\mathcal{A} = \{P \subseteq E \mid G \not\subseteq \operatorname{Shade} P\}. \tag{5}$$

Thus, $\mathcal{A}$ is a subset of the power set of E, and each $F \in \mathcal{A}$ satisfies $G \not\subseteq \operatorname{Shade} F$.

We equip the finite set E with a total order (chosen arbitrarily, but fixed henceforth). If $F \in \mathcal{A}$, then there exists a unique **smallest** edge $e \in G \setminus \operatorname{Shade} F$ (since $F \in \mathcal{A}$ entails $G \not\subseteq \operatorname{Shade} F$ and thus $G \setminus \operatorname{Shade} F \neq \varnothing$). This unique smallest edge e will be denoted by $\varepsilon(F)$.

We note that the edge $\varepsilon(F)$ (for a set $F \in \mathcal{A}$) depends only on $\operatorname{Shade} F$, but not on F itself (because it was defined as the smallest edge $e \in G \setminus \operatorname{Shade} F$). Thus, if two sets $F_1 \in \mathcal{A}$ and $F_2 \in \mathcal{A}$ satisfy $\operatorname{Shade}(F_1) = \operatorname{Shade}(F_2)$, then

$$\varepsilon(F_1) = \varepsilon(F_2). \tag{6}$$

We also notice the following simple fact: If F and F' are two subsets of E such that $F \in \mathcal{A}$ and $\operatorname{Shade}(F') = \operatorname{Shade} F$, then

$$F' \in \mathcal{A}. \tag{7}$$

(Indeed, $F \in \mathcal{A}$ means that $G \not\subseteq \operatorname{Shade} F$; but because of $\operatorname{Shade}(F') = \operatorname{Shade} F$, this entails $G \not\subseteq \operatorname{Shade}(F')$ as well, and therefore $F' \in \mathcal{A}$.)

We now define two subsets $\mathcal{A}_+$ and $\mathcal{A}_-$ of $\mathcal{A}$ by

$$\mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\} \qquad \text{and} \qquad \mathcal{A}_- = \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\}.$$

Next, we claim the following:

> *Claim 1:* Let $F \in \mathcal{A}_+$. Set $F' = F \setminus \{\varepsilon(F)\}$. Then, $F' \in \mathcal{A}_-$ and $F' \cup \{\varepsilon(F')\} = F$ and $(-1)^{|F'|} = -(-1)^{|F|}$.

[*Proof of Claim 1:* We have $F \in \mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$; in other words, $F \in \mathcal{A}$ and $\varepsilon(F) \in F$. From $F \in \mathcal{A} = \{P \subseteq E \mid G \not\subseteq \operatorname{Shade} P\}$, we obtain $F \subseteq E$ and $G \not\subseteq \operatorname{Shade} F$.

The definition of F' yields $F' \subseteq F \subseteq E$.

Recall that $\varepsilon(F)$ is the smallest edge $e \in G \setminus \operatorname{Shade} F$ (by the definition of $\varepsilon(F)$). Hence, $\varepsilon(F) \in G \setminus \operatorname{Shade} F$. In other words, $\varepsilon(F) \in G$ and $\varepsilon(F) \notin \operatorname{Shade} F$. Thus, $\varepsilon(F) \in G \subseteq E$ and $\varepsilon(F) \notin \operatorname{Shade} F$. Therefore, (4) (applied to $u = \varepsilon(F)$) yields $\operatorname{Shade}(F \setminus \{\varepsilon(F)\}) = \operatorname{Shade} F$. This can be rewritten as $\operatorname{Shade}(F') = \operatorname{Shade} F$ (since $F' = F \setminus \{\varepsilon(F)\}$). Hence, (7) yields $F' \in \mathcal{A}$. In light of the preceding two sentences, (6) (applied to F' and F instead of F_1 and F_2) yields $\varepsilon(F') = \varepsilon(F)$. However, $\varepsilon(F) \notin F'$ (since $F' = F \setminus \{\varepsilon(F)\}$). In other words, $\varepsilon(F') \notin F'$ (since $\varepsilon(F') = \varepsilon(F)$). Hence, $F' \in \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\}$ (since $F' \in \mathcal{A}$). In other words, $F' \in \mathcal{A}_-$ (since $\mathcal{A}_- = \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\}$).

Moreover, from $\varepsilon(F') = \varepsilon(F)$, we obtain $F' \cup \{\varepsilon(F')\} = F' \cup \{\varepsilon(F)\} = F$ (since $F' = F \setminus \{\varepsilon(F)\}$ and $\varepsilon(F) \in F$).

Finally, the set $F' = F \setminus \{\varepsilon(F)\}$ has exactly one less element than the set F (since $\varepsilon(F) \in F$). That is, $|F'| = |F| - 1$. Hence, $(-1)^{|F'|} = -(-1)^{|F|}$. This completes the proof of Claim 1.]

> *Claim 2:* Let $F \in \mathcal{A}_-$. Set $F' = F \cup \{\varepsilon(F)\}$. Then, $F' \in \mathcal{A}_+$ and $F' \setminus \{\varepsilon(F')\} = F$.

[*Proof of Claim 2:* We have $F \in \mathcal{A}_- = \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\}$; in other words, $F \in \mathcal{A}$ and $\varepsilon(F) \notin F$. From $F \in \mathcal{A} = \{P \subseteq E \mid G \not\subseteq \operatorname{Shade} P\}$, we obtain $F \subseteq E$ and $G \not\subseteq \operatorname{Shade} F$.

As in the proof of Claim 1, we can see that $\varepsilon(F) \in G$ and $\varepsilon(F) \notin \operatorname{Shade} F$. Thus, $\varepsilon(F) \in G \subseteq E$ and $\varepsilon(F) \notin \operatorname{Shade} F$. Now, $F' = F \cup \{\varepsilon(F)\} \subseteq E$ (since $F \subseteq E$ and $\varepsilon(F) \in E$). Furthermore, (3) (applied to $u = \varepsilon(F)$) yields $\operatorname{Shade}(F \cup \{\varepsilon(F)\}) = \operatorname{Shade} F$ (since $\varepsilon(F) \notin \operatorname{Shade} F$). This can be rewritten as $\operatorname{Shade}(F') = \operatorname{Shade} F$ (since $F' = F \cup \{\varepsilon(F)\}$). Hence, (7) yields $F' \in \mathcal{A}$. In light of the preceding two sentences, (6) (applied to F' and F instead of F_1 and F_2) yields $\varepsilon(F') = \varepsilon(F) \in \{\varepsilon(F)\} \subseteq F \cup \{\varepsilon(F)\} = F'$. Thus, $F' \in \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$ (since $F' \in \mathcal{A}$). In other words, $F' \in \mathcal{A}_+$ (since $\mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$).

Moreover, from $\varepsilon(F') = \varepsilon(F)$, we obtain $F' \setminus \{\varepsilon(F')\} = F' \setminus \{\varepsilon(F)\} = F$ (since $F' = F \cup \{\varepsilon(F)\}$ and $\varepsilon(F) \notin F$). This completes the proof of Claim 2.]

Each $F \in \mathcal{A}_+$ satisfies $F \setminus \{\varepsilon(F)\} \in \mathcal{A}_-$ (by Claim 1, applied to $F' = F \setminus \{\varepsilon(F)\}$). Thus, we can define a map

$$\begin{aligned} \Phi : \mathcal{A}_+ &\to \mathcal{A}_-, \\ F &\mapsto F \setminus \{\varepsilon(F)\}. \end{aligned}$$

Each $F \in \mathcal{A}_-$ satisfies $F \cup \{\varepsilon(F)\} \in \mathcal{A}_+$ (by Claim 2, applied to $F' = F \cup \{\varepsilon(F)\}$). Thus, we can define a map

$$\begin{aligned} \Psi : \mathcal{A}_- &\to \mathcal{A}_+, \\ F &\mapsto F \cup \{\varepsilon(F)\}. \end{aligned}$$

We have $\Phi \circ \Psi = \mathrm{id}$ (this follows from the "$F' \setminus \{\varepsilon(F')\} = F$" part of Claim 2) and $\Psi \circ \Phi = \mathrm{id}$ (this follows from the "$F' \cup \{\varepsilon(F')\} = F$" part of Claim 1). Thus, the maps Φ and Ψ are mutually inverse. Hence, the map Φ is invertible, thus a bijection.

Moreover, each $F \in \mathcal{A}_+$ satisfies

$$(-1)^{|\Phi(F)|} = -(-1)^{|F|}. \tag{8}$$

(Indeed, this is just the "$(-1)^{|F'|} = -(-1)^{|F|}$" part of Claim 1.)

Now, the summation sign "$\sum\limits_{\substack{F \subseteq E; \\ G \not\subseteq \operatorname{Shade} F}}$" is equivalent to the summation sign "$\sum\limits_{F \in \mathcal{A}}$" (since the set of all subsets F of E satisfying $G \not\subseteq \operatorname{Shade} F$ is precisely $\mathcal{A}$). Thus,

$$\begin{aligned} &\sum_{\substack{F \subseteq E; \\ G \not\subseteq \operatorname{Shade} F}} (-1)^{|F|} \\ &= \sum_{F \in \mathcal{A}} (-1)^{|F|} = \underbrace{\sum_{\substack{F \in \mathcal{A}; \\ \varepsilon(F) \in F}}}_{\substack{= \sum\limits_{F \in \mathcal{A}_+} \\ (\text{since } \{P \in \mathcal{A} \mid \varepsilon(P) \in P\} = \mathcal{A}_+)}} (-1)^{|F|} + \underbrace{\sum_{\substack{F \in \mathcal{A}; \\ \varepsilon(F) \notin F}}}_{\substack{= \sum\limits_{F \in \mathcal{A}_-} \\ (\text{since } \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\} = \mathcal{A}_-)}} (-1)^{|F|} \\ &= \sum_{F \in \mathcal{A}_+} (-1)^{|F|} + \sum_{F \in \mathcal{A}_-} (-1)^{|F|} = \sum_{F \in \mathcal{A}_+} (-1)^{|F|} + \sum_{F \in \mathcal{A}_+} \underbrace{(-1)^{|\Phi(F)|}}_{\substack{= -(-1)^{|F|} \\ (\text{by } (8))}} \\ &\qquad \left(\begin{array}{c} \text{here, we have substituted } \Phi(F) \text{ for } F \text{ in the second sum,} \\ \text{since the map } \Phi : \mathcal{A}_+ \to \mathcal{A}_- \text{ is a bijection} \end{array} \right) \\ &= \sum_{F \in \mathcal{A}_+} (-1)^{|F|} + \sum_{F \in \mathcal{A}_+} \left(-(-1)^{|F|} \right) = \sum_{F \in \mathcal{A}_+} (-1)^{|F|} - \sum_{F \in \mathcal{A}_+} (-1)^{|F|} = 0. \end{aligned}$$

This proves Theorem 2.6. $\square$

In order to derive Theorem 2.5 from Theorem 2.6, we need the following innocent lemma – which is one of the simplest facts in enumerative combinatorics:

Lemma 2.7. Let U be a finite set with $U \neq \varnothing$. Then,

$$\sum_{F \subseteq U} (-1)^{|F|} = 0.$$

Lemma 2.7 can easily be derived from the fact that $\sum_{k=0}^{n} (-1)^k \binom{n}{k} = 0$ for any positive integer n (as follows readily from the binomial theorem). However, keeping true to the spirit of this paper, let us give a bijective proof for it:

Proof of Lemma 2.7. This is a standard argument that underlies many combinatorial proofs of alternating sum identities (see, for example, [Sagan20, proof of (2.4)] or [BenQui08, proof of (1)]). For the sake of completeness, let us nevertheless recall it.

We have $U \neq \varnothing$; hence, there exists some $u \in U$. Consider this u.

It is easy to see that the maps

$$\Phi : \{F \subseteq U \mid u \in F\} \to \{F \subseteq U \mid u \notin F\}, \qquad F \mapsto F \setminus \{u\}$$

and

$$\Psi : \{F \subseteq U \mid u \notin F\} \to \{F \subseteq U \mid u \in F\}, \qquad F \mapsto F \cup \{u\}$$

are well-defined.

The maps Φ and Ψ we just defined are clearly mutually inverse. Thus, they are invertible, i.e., are bijections. Hence, in particular, Φ is a bijection. Thus, we can substitute $\Phi(F)$ for F in the sum $\sum_{\substack{F \subseteq U;\\ u \notin F}} (-1)^{|F|}$. We thus obtain

$$\begin{aligned}
\sum_{\substack{F \subseteq U;\\ u \notin F}} (-1)^{|F|} &= \sum_{\substack{F \subseteq U;\\ u \in F}} \underbrace{(-1)^{|\Phi(F)|}}_{\substack{=(-1)^{|F\setminus\{u\}|} \\ \text{(since } \Phi(F)=F\setminus\{u\} \\ \text{(by the definition of } \Phi))}} = \sum_{\substack{F \subseteq U;\\ u \in F}} \underbrace{(-1)^{|F\setminus\{u\}|}}_{\substack{=(-1)^{|F|-1} \\ \text{(since } |F\setminus\{u\}|=|F|-1 \\ \text{(because } u\in F))}} \\
&= \sum_{\substack{F \subseteq U;\\ u \in F}} \underbrace{(-1)^{|F|-1}}_{=-(-1)^{|F|}} = -\sum_{\substack{F \subseteq U;\\ u \in F}} (-1)^{|F|}.
\end{aligned}$$

However, each $F \subseteq U$ satisfies either $u \in F$ or $u \notin F$ (but not both). Hence,

$$\sum_{F \subseteq U} (-1)^{|F|} = \sum_{\substack{F \subseteq U;\\ u \in F}} (-1)^{|F|} + \underbrace{\sum_{\substack{F \subseteq U;\\ u \notin F}} (-1)^{|F|}}_{=-\sum_{\substack{F \subseteq U;\\ u \in F}} (-1)^{|F|}} = \sum_{\substack{F \subseteq U;\\ u \in F}} (-1)^{|F|} - \sum_{\substack{F \subseteq U;\\ u \in F}} (-1)^{|F|} = 0.$$

This proves Lemma 2.7. $\square$

We can now easily derive Theorem 2.5 from Theorem 2.6:

Proof of Theorem 2.5. Each subset F of E satisfies either $G \subseteq \operatorname{Shade} F$ or $G \not\subseteq \operatorname{Shade} F$ (but not both at the same time). Hence,

$$\sum_{F \subseteq E} (-1)^{|F|} = \sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|} + \underbrace{\sum_{\substack{F \subseteq E; \\ G \not\subseteq \operatorname{Shade} F}} (-1)^{|F|}}_{\substack{=0 \\ \text{(by Theorem 2.6)}}} = \sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|}.$$

Therefore,

$$\sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|} = \sum_{F \subseteq E} (-1)^{|F|} = 0$$

(by Lemma 2.7, applied to $U = E$). This proves Theorem 2.5. □

2.3. Proving Theorem 1.2

Theorem 1.2 is now a simple particular case of Theorem 2.5:

Proof of Theorem 1.2. Let G be the set E. Thus, $G = E$. Hence, for each subset F of E, we have the following chain of logical equivalences:

$$\begin{aligned} (G \subseteq \operatorname{Shade} F) &\iff (E \subseteq \operatorname{Shade} F) \\ &\iff (\text{each } u \in E \text{ satisfies } u \in \operatorname{Shade} F) \\ &\iff (\text{each } u \in E \text{ has the property that } F \text{ infects } u) \\ &\qquad\qquad (\text{since } \operatorname{Shade} F = \{e \in E \mid F \text{ infects } e\}) \\ &\iff (F \text{ infects each } u \in E) \\ &\iff (F \text{ is pandemic}) \qquad (\text{by the definition of "pandemic"}). \end{aligned}$$

Thus, the summation sign "$\sum\limits_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}}$" can be rewritten as "$\sum\limits_{\substack{F \subseteq E \text{ is} \\ \text{pandemic}}}$". Hence,

$$\sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|} = \sum_{\substack{F \subseteq E \text{ is} \\ \text{pandemic}}} (-1)^{|F|}.$$

But the left hand side of this equality is 0 (by Theorem 2.5). Hence, its right hand side is 0 as well. This proves Theorem 1.2. □

3. Vertex infection and other variants

In our study of graphs so far, we have barely ever mentioned vertices (even though they are, of course, implicit in the notion of a path). It may appear somewhat strange to talk about a subset infecting an edge, when the infection is spread from vertex to vertex. One might thus wonder if there is also a vertex counterpart of Theorem 1.2. So let us define analogues of our notions for vertices:

If $F \subseteq V$, then an *F-vertex-path* shall mean a path of Γ such that all vertices of the path except (possibly) for its two endpoints belong to F. (Thus, if a path has only one edge or none, then it automatically is an F-vertex-path.)

If $w \in V \setminus \{v\}$ is any vertex and $F \subseteq V \setminus \{v\}$ is any subset, then we say that *F vertex-infects w* if there exists an F-vertex-path from v to w. (This is always true when w is v or a neighbor of v.)

A subset $F \subseteq V \setminus \{v\}$ is said to be *vertex-pandemic* if it vertex-infects each vertex $w \in V \setminus \{v\}$.

Example 3.1. Let Γ be as in Example 1.3. Then, the path $v \xrightarrow{1} p \xrightarrow{2} q$ is an F-vertex-path for any subset $F \subseteq V$ that satisfies $p \in F$. The subset $\{p\}$ of $V \setminus \{v\}$ vertex-infects each vertex (for example, $v \xrightarrow{1} p \xrightarrow{2} q$ is a $\{p\}$-vertex-path from v to q, and $v \xrightarrow{4} w$ is a $\{p\}$-vertex-path from v to w), and thus is vertex-pandemic. The vertex-pandemic subsets of $V \setminus \{v\}$ are the sets

$$\{p\}, \quad \{w\}, \quad \{p,q\}, \quad \{p,w\}, \quad \{q,w\}, \quad \{p,q,w\}.$$

We now have the following analogue of Theorem 1.2:

Theorem 3.2. Assume that $V \setminus \{v\} \neq \varnothing$. Then,

$$\sum_{\substack{F \subseteq V \setminus \{v\} \text{ is} \\ \text{vertex-pandemic}}} (-1)^{|F|} = 0.$$

Proof of Theorem 3.2. With just a few easy modifications, our above proof of Theorem 1.2 can be repurposed as a proof of Theorem 3.2. Namely:

- We need to replace "edge" by "vertex" throughout the argument (including Definition 2.1, Lemma 2.3, Lemma 2.4, Theorem 2.5 and Theorem 2.6), as well as replace E by $V \setminus \{v\}$.
- The words "F-path", "infects" and "pandemic" have to be replaced by "F-vertex-path", "vertex-infects" and "vertex-pandemic", respectively.
- In the proofs of Lemma 2.3 and Lemma 2.4, the words "an endpoint of" (as well as "any endpoint of" and "some endpoint of") need to be removed (since the notion of "vertex-infects" is defined not in terms of paths to an endpoint of a given edge, but in terms of paths to a given vertex).

- In the proof of Lemma 2.4, specifically in the proof of (3), the path π is now cut not by removing the edge u, but by splitting the path π at the vertex u.

The reader may check that these changes result in a valid proof of Theorem 3.2. □

Another variant of Theorem 1.2 (and Theorem 2.5 and Theorem 2.6) is obtained by replacing the undirected graph Γ with a directed graph (while, of course, replacing paths by directed paths). More generally, we can replace Γ by a "hybrid" graph with some directed and some undirected edges.[2] No changes are required to the above proofs. Yet another variation can be obtained by replacing "endpoint" by "source" (for directed edges). We cannot, however, replace "endpoint" by "target".

4. An abstract perspective

Seeing how little graph theory we have used in proving Theorem 1.2, and how easily the same argument adapted to Theorem 3.2, we get the impression that there might be some general theory lurking behind it. What follows is an attempt at building this theory.

Let $\mathcal{P}(E)$ denote the power set of E. In Definition 2.1, we have encoded the "infects" relation as a map $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ defined by

$$\operatorname{Shade} F = \{e \in E \mid F \text{ infects } e\}.$$

As we recall, Theorem 2.5 (a generalization of Theorem 1.2) states that

$$\sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|} = 0 \tag{9}$$

for any $G \subseteq E$, under the assumption that $E \neq \varnothing$.

To generalize this, we forget about the graph Γ and the map Shade, and instead start with an **arbitrary** finite set E. (This corresponds to the set E in Theorem 1.2 and to the set $V \setminus \{v\}$ in Theorem 3.2.) Let $\mathcal{P}(E)$ be the power set of E. Let $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ be an arbitrary map (meant to generalize the map Shade from the previous paragraph). We may now ask:

> **Question 4.1.** What (combinatorial) properties must Shade satisfy in order for (9) to hold for any $G \subseteq E$ under the assumption that $E \neq \varnothing$?

A partial answer to this question can be given by analyzing our above proof of Theorem 2.5 and extracting what was used:

[2] We understand that a directed edge still has two endpoints: its source and its target.

Theorem 4.2. Let E be a finite set. Let $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ be an arbitrary map that satisfies the following axioms:

Axiom 1: If $F \in \mathcal{P}(E)$ and $u \in E \setminus \operatorname{Shade} F$, then $\operatorname{Shade}(F \cup \{u\}) = \operatorname{Shade} F$.

Axiom 2: If $F \in \mathcal{P}(E)$ and $u \in E \setminus \operatorname{Shade} F$, then $\operatorname{Shade}(F \setminus \{u\}) = \operatorname{Shade} F$.

Assume that $E \neq \varnothing$. Let G be any subset of E. Then,

$$\sum_{\substack{F \subseteq E; \\ G \subseteq \operatorname{Shade} F}} (-1)^{|F|} = 0.$$

Proof sketch. Again, analogous to our above proof of Theorem 2.5. (This time, in the proof of Lemma 2.4, the equalities (3) and (4) follow directly from Axiom 1 and Axiom 2, respectively.) $\square$

Question 4.3. Are the axioms in Theorem 4.2 related to some known concepts in the combinatorics of set families (such as topologies, clutters, matroids, or submodular functions)?

Note that Axioms 1 and 2 can be weakened to the following statements:

Axiom 1′: If $F \in \mathcal{P}(E)$ and $u \in E \setminus \operatorname{Shade} F$, then $\operatorname{Shade}(F \cup \{u\}) \subseteq \operatorname{Shade} F$.

Axiom 2′: If $F \in \mathcal{P}(E)$ and $u \in E \setminus \operatorname{Shade} F$, then $\operatorname{Shade}(F \setminus \{u\}) \subseteq \operatorname{Shade} F$.

Axiom 1′ is weaker than Axiom 1, and likewise Axiom 2′ is weaker than Axiom 2. However, Axioms 1′ and 2′ combined are equivalent to Axioms 1 and 2 combined (exercise!).

Here is one restatement of both Axioms 1′ and 2′: For any $F \subseteq E$ and any two elements $u, v \in E \setminus \operatorname{Shade} F$, we have $v \notin \operatorname{Shade}(F \cup \{u\})$ (Axiom 1′) and $v \notin \operatorname{Shade}(F \setminus \{u\})$ (Axiom 2′).

Axioms 1 and 2 can also be combined into one common axiom:

Axiom 3: If $F \in \mathcal{P}(E)$ and $u \in E \setminus F$, then we have $\operatorname{Shade} F = \operatorname{Shade}(F \cup \{u\})$ or $u \in (\operatorname{Shade} F) \cap \operatorname{Shade}(F \cup \{u\})$.

Question 4.4. What are examples of maps $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ satisfying the two axioms in Theorem 4.2?

We note that the map Shade does not have to be monotonic (i.e., it is not necessary that $\operatorname{Shade} A \subseteq \operatorname{Shade} B$ whenever $A \subseteq B$). Examples of non-monotonic maps Shade that satisfy Axioms 1 and 2 are easily constructed. (Indeed, if $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ is any map satisfying Axioms 1 and 2, then the map $\operatorname{Shade}' : \mathcal{P}(E) \to \mathcal{P}(E)$ that sends each $F \in \mathcal{P}(E)$ to $\operatorname{Shade}(E \setminus F) \in \mathcal{P}(E)$ also satisfies Axioms 1 and 2; but it is rare for both Shade and Shade$'$ to be monotonic.)

Another example of a map $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ satisfying the two axioms in Theorem 4.2 comes from discrete geometry:

Example 4.5. Let A be an affine space over $\mathbb{R}$. If S is a finite subset of A, then a *nontrivial convex combination* of S will mean a point of the form $\sum_{s \in S} \lambda_s s \in A$, where the coefficients λ_s are nonnegative reals smaller than 1 and satisfying $\sum_{s \in S} \lambda_s = 1$.

Fix a finite subset E of A. For any $F \subseteq E$, we define

$$\operatorname{Shade} F = \{e \in E \mid e \text{ is } \textbf{not} \text{ a nontrivial convex combination of } F\}.$$

Then, this map $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ satisfies the two axioms in Theorem 4.2.

For a (not very difficult) proof of Example 4.5, see the detailed version [Grinbe20] of this paper.

As a contrast to Example 4.5, let us mention a not-quite-example (satisfying only one of the two axioms in Theorem 4.2):

Example 4.6. Let V be a vector space over $\mathbb{R}$. If S is a finite subset of V, then a *nontrivial conic combination* of S will mean a vector of the form $\sum_{s \in S} \lambda_s s \in V$, where the coefficients λ_s are nonnegative reals with the property that at least two elements $s \in S$ satisfy $\lambda_s > 0$.

Fix a finite subset E of V. For any $F \subseteq E$, we define

$$\operatorname{Shade} F = \{e \in E \mid e \text{ is } \textbf{not} \text{ a nontrivial conic combination of } F\}.$$

Then, it can be shown that this map $\operatorname{Shade} : \mathcal{P}(E) \to \mathcal{P}(E)$ satisfies Axiom 1 in Theorem 4.2. In general, it does not satisfy Axiom 2.

5. The topological viewpoint

We shall now reinterpret Theorem 2.6 in the terms of combinatorial topology (specifically, finite simplicial complexes) and strengthen it. We recall the definition of a *simplicial complex*:[3]

[3] We forget all the conventions we have introduced so far. (Thus, for example, E no longer means the edge set of a graph Γ.)

Definition 5.1. Let E be a finite set. A *simplicial complex* on ground set E means a subset $\mathcal{A}$ of the power set of E with the following property:

If $P \in \mathcal{A}$ and $Q \subseteq P$, then $Q \in \mathcal{A}$.

Thus, in terms of posets, a simplicial complex on ground set E means a down-closed subset of the Boolean lattice on E. Note that a simplicial complex contains the empty set $\varnothing$ unless it is empty itself.

We refer to [Kozlov20] for context and theory about simplicial complexes. We shall restrict ourselves to the few definitions relevant to what we will prove. The following is fairly simple:

Proposition 5.2. Let us use the notations from Section 1 as well as Definition 2.1. Let G be any subset of E. Let

$$\mathcal{A} = \{F \subseteq E \mid G \not\subseteq \operatorname{Shade} F\}. \tag{10}$$

Then, $\mathcal{A}$ is a simplicial complex on ground set E.

Proof of Proposition 5.2. Clearly, $\mathcal{A}$ is a subset of the power set of E. Thus, we only need to verify the following claim:

Claim 1: If $P \in \mathcal{A}$ and $Q \subseteq P$, then $Q \in \mathcal{A}$.

[*Proof of Claim 1:* Let $P \in \mathcal{A}$ and let $Q \subseteq P$. We must show that $Q \in \mathcal{A}$.

We have $P \in \mathcal{A} = \{F \subseteq E \mid G \not\subseteq \operatorname{Shade} F\}$. In other words, $P \subseteq E$ and $G \not\subseteq \operatorname{Shade} P$. But $Q \subseteq P$ and thus $\operatorname{Shade} Q \subseteq \operatorname{Shade} P$ (by Lemma 2.3). Hence, from $G \not\subseteq \operatorname{Shade} P$, we obtain $G \not\subseteq \operatorname{Shade} Q$. Thus, $Q \in \{F \subseteq E \mid G \not\subseteq \operatorname{Shade} F\}$. This can be rewritten as $Q \in \mathcal{A}$ (by (10)). Thus, Claim 1 is proved.] □

To state the main result of this section, we need a few more notions:

Definition 5.3. Let A and B be two sets. Then, we say that $A \prec B$ if we have $B = A \cup \{b\}$ for some $b \in B \setminus A$.

Equivalently, two sets A and B satisfy $A \prec B$ if and only if $A \subseteq B$ and $|B \setminus A| = 1$.

Definition 5.4. Let E be a finite set. Let $\mathcal{A}$ be a simplicial complex on ground set E.

(a) A *complete matching* of $\mathcal{A}$ means a triple $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$, where $\mathcal{A}_-$ and $\mathcal{A}_+$ are two disjoint subsets of $\mathcal{A}$ satisfying $\mathcal{A}_- \cup \mathcal{A}_+ = \mathcal{A}$, and where $\Phi : \mathcal{A}_+ \to \mathcal{A}_-$ is a bijection with the property that

$$\text{each } F \in \mathcal{A}_+ \text{ satisfies } \Phi(F) \prec F. \tag{11}$$

(b) A complete matching $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$ of $\mathcal{A}$ is said to be *acyclic* if there exists no tuple $(B_1, B_2, \ldots, B_n)$ of distinct sets $B_1, B_2, \ldots, B_n \in \mathcal{A}_+$ with the property that $n \geqslant 2$ and that

$$\Phi(B_i) \prec B_{i+1} \qquad \text{for each } i \in \{1, 2, \ldots, n-1\}$$

and

$$\Phi(B_n) \prec B_1.$$

(c) The simplicial complex $\mathcal{A}$ is said to be *collapsible* if it has an acyclic complete matching.

These definitions are essentially equivalent to the definitions in [Kozlov20], although it takes a bit of work to match them up precisely. Our notion of a "complete matching" as defined in Definition 5.4 **(a)** is a particular case of the notion introduced in [Kozlov20, Chapter 10], as we restrict ourselves to simplicial complexes (i.e., down-closed sets of Boolean lattices) instead of arbitrary posets. To be fully precise, our complete matchings are triples $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$, whereas the complete matchings of [Kozlov20, Chapter 10] are certain fixed-point-free involutions[4] $\mu : \mathcal{A} \to \mathcal{A}$; the equivalence between these two objects is fairly easy to see (in particular, if $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$ is a complete matching in our sense, then the corresponding complete matching $\mu : \mathcal{A} \to \mathcal{A}$ in the sense of [Kozlov20, Chapter 10] is the map that sends each $B \in \mathcal{A}_+$ to $\Phi(B) \in \mathcal{A}_-$ and sends each $B \in \mathcal{A}_-$ to $\Phi^{-1}(B) \in \mathcal{A}_+$). Our notion of "collapsible" as defined in Definition 5.4 **(c)** is equivalent to the classical notion of "collapsible" (even though the latter is usually defined differently) because of [Kozlov20, Theorem 10.9].

We now claim:

Theorem 5.5. Let us use the notations from Section 1 as well as Definition 2.1. Let G be any subset of E. Define $\mathcal{A}$ as in (10). Then, the simplicial complex $\mathcal{A}$ is collapsible.

Collapsible simplicial complexes are well-behaved in various ways – in particular, they are contractible ([Kozlov20, Corollary 9.19]), and thus have trivial homotopy and homology groups (in positive degrees). Moreover, the reduced Euler characteristic of any collapsible simplicial complex is 0 (for obvious reasons: having a complete matching suffices, even if it is not acyclic); thus, Theorem 2.6 follows from Theorem 5.5.

Let us now sketch the proof of Theorem 5.5:

Proof of Theorem 5.5. We know from Proposition 5.2 that $\mathcal{A}$ is a simplicial complex. It remains to show that $\mathcal{A}$ is collapsible.

[4] A *fixed-point-free involution* means an involution (i.e., a map that is inverse to itself) that has no fixed point.

We have

$$\mathcal{A} = \{F \subseteq E \mid G \not\subseteq \operatorname{Shade} F\} = \{P \subseteq E \mid G \not\subseteq \operatorname{Shade} P\}$$

(here, we have renamed the index F as P). Thus, our set $\mathcal{A}$ is precisely the set $\mathcal{A}$ defined in the proof of Theorem 2.6 above.

We equip the finite set E with a total order (chosen arbitrarily, but fixed henceforth).

If $F \in \mathcal{A}$, then we define the edge $\varepsilon(F) \in G \setminus \operatorname{Shade} F$ as in the proof of Theorem 2.6. That is, we define $\varepsilon(F)$ as the **smallest** edge $e \in G \setminus \operatorname{Shade} F$.

If two sets $F_1 \in \mathcal{A}$ and $F_2 \in \mathcal{A}$ satisfy $\operatorname{Shade}(F_1) = \operatorname{Shade}(F_2)$, then

$$\varepsilon(F_1) = \varepsilon(F_2). \tag{12}$$

(Indeed, this is precisely the equality (6) from the above proof of Theorem 2.6.)

We define two subsets $\mathcal{A}_+$ and $\mathcal{A}_-$ of $\mathcal{A}$ as in the proof of Theorem 2.6. That is, we set

$$\mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\} \qquad \text{and} \qquad \mathcal{A}_- = \{P \in \mathcal{A} \mid \varepsilon(P) \notin P\}.$$

Thus, each $P \in \mathcal{A}$ satisfies either $P \in \mathcal{A}_-$ or $P \in \mathcal{A}_+$ but not both at the same time (since it satisfies either $\varepsilon(P) \notin P$ or $\varepsilon(P) \in P$ but not both at the same time). Hence, $\mathcal{A}_-$ and $\mathcal{A}_+$ are two disjoint subsets of $\mathcal{A}$ satisfying $\mathcal{A}_- \cup \mathcal{A}_+ = \mathcal{A}$.

We define a map $\Phi : \mathcal{A}_+ \to \mathcal{A}_-$ as in the proof of Theorem 2.6. That is, we set

$$\Phi(F) = F \setminus \{\varepsilon(F)\} \qquad \text{for each } F \in \mathcal{A}_+.$$

We know (from the proof of Theorem 2.6) that the map Φ is a bijection. Moreover, it is clear that each $F \in \mathcal{A}_+$ satisfies $\Phi(F) \prec F$ [5]. Hence, the triple $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$ is a complete matching of $\mathcal{A}$.

We shall now prove that this complete matching $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$ is acyclic. Indeed, let $(B_1, B_2, \ldots, B_n)$ be a tuple of distinct sets $B_1, B_2, \ldots, B_n \in \mathcal{A}_+$ with the property that $n \geqslant 2$ and that

$$\Phi(B_i) \prec B_{i+1} \qquad \text{for each } i \in \{1, 2, \ldots, n-1\} \tag{13}$$

and

$$\Phi(B_n) \prec B_1. \tag{14}$$

We shall derive a contradiction.

Set $B_{n+1} = B_1$. Then, (14) can be rewritten as $\Phi(B_n) \prec B_{n+1}$. In other words, we have $\Phi(B_i) \prec B_{i+1}$ for $i = n$. Combining this with (13), we conclude that

$$\Phi(B_i) \prec B_{i+1} \qquad \text{for each } i \in \{1, 2, \ldots, n\}. \tag{15}$$

[5] *Proof.* Let $F \in \mathcal{A}_+$. Thus, $F \in \mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$. In other words, F is a $P \in \mathcal{A}$ satisfying $\varepsilon(P) \in P$. In other words, $F \in \mathcal{A}$ and $\varepsilon(F) \in F$. From $\varepsilon(F) \in F$, we obtain $F = (F \setminus \{\varepsilon(F)\}) \cup \{\varepsilon(F)\}$ and $\varepsilon(F) \in F \setminus (F \setminus \{\varepsilon(F)\})$. Hence, $F \setminus \{\varepsilon(F)\} \prec F$ (by Definition 5.3). In other words, $\Phi(F) \prec F$ (since the definition of Φ yields $\Phi(F) = F \setminus \{\varepsilon(F)\}$). Qed.

Now, set $A_i = \operatorname{Shade}(B_i)$ for each $i \in \{1, 2, \ldots, n+1\}$. Then, $A_{n+1} = A_1$ (since $B_{n+1} = B_1$).

We now claim the following:

Claim 1: We have $A_i \subseteq A_{i+1}$ for each $i \in \{1, 2, \ldots, n\}$.

[*Proof of Claim 1:* Let $i \in \{1, 2, \ldots, n\}$. Then, the definition of A_i yields $A_i = \operatorname{Shade}(B_i)$. Likewise, $A_{i+1} = \operatorname{Shade}(B_{i+1})$.

We have $B_i \in \mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$. In other words, B_i is a $P \in \mathcal{A}$ satisfying $\varepsilon(P) \in P$. In other words, B_i is an element of $\mathcal{A}$ and satisfies $\varepsilon(B_i) \in B_i$.

We set $u = \varepsilon(B_i)$. The definition of Φ yields $\Phi(B_i) = B_i \setminus \{\varepsilon(B_i)\} = B_i \setminus \{u\}$ (since $\varepsilon(B_i) = u$).

Recall that $\varepsilon(B_i)$ is the smallest edge $e \in G \setminus \operatorname{Shade}(B_i)$ (by the definition of $\varepsilon(B_i)$). Hence, $\varepsilon(B_i) \in G \setminus \operatorname{Shade}(B_i)$. In other words, $u \in G \setminus \operatorname{Shade}(B_i)$ (since $u = \varepsilon(B_i)$). In other words, $u \in G$ and $u \notin \operatorname{Shade}(B_i)$. Thus, $u \in G \subseteq E$ and $u \notin \operatorname{Shade}(B_i)$. Therefore, (4) (applied to $F = B_i$) yields $\operatorname{Shade}(B_i \setminus \{u\}) = \operatorname{Shade}(B_i)$. This can be rewritten as $\operatorname{Shade}(\Phi(B_i)) = \operatorname{Shade}(B_i)$ (since $\Phi(B_i) = B_i \setminus \{u\}$).

But (15) yields $\Phi(B_i) \prec B_{i+1}$, so that $\Phi(B_i) \subseteq B_{i+1}$ and thus $\operatorname{Shade}(\Phi(B_i)) \subseteq \operatorname{Shade}(B_{i+1})$ (by Lemma 2.3, applied to $A = \Phi(B_i)$ and $B = B_{i+1}$). In view of $\operatorname{Shade}(\Phi(B_i)) = \operatorname{Shade}(B_i)$, this can be rewritten as $\operatorname{Shade}(B_i) \subseteq \operatorname{Shade}(B_{i+1})$. This proves Claim 1.]

Claim 1 shows that $A_i \subseteq A_{i+1}$ for each $i \in \{1, 2, \ldots, n\}$. In other words,

$$A_1 \subseteq A_2 \subseteq \cdots \subseteq A_n \subseteq A_{n+1}.$$

This is a chain of inclusions, but its last entry equals the first: indeed, $A_{n+1} = A_1$. Thus, all inclusions in this chain must be equalities. That is, we have

$$A_1 = A_2 = \cdots = A_n = A_{n+1}. \tag{16}$$

Hence, in particular, $A_n = A_1$. However, $A_n = \operatorname{Shade}(B_n)$ (by the definition of A_n) and $A_1 = \operatorname{Shade}(B_1)$ (by the definition of A_1). Hence, $\operatorname{Shade}(B_n) = A_n = A_1 = \operatorname{Shade}(B_1)$. Thus, (12) (applied to $F_1 = B_n$ and $F_2 = B_1$) yields $\varepsilon(B_n) = \varepsilon(B_1)$ (since $B_n \in \mathcal{A}_+ \subseteq \mathcal{A}$ and $B_1 \in \mathcal{A}_+ \subseteq \mathcal{A}$).

Set $u = \varepsilon(B_n)$. Thus, $u = \varepsilon(B_n) = \varepsilon(B_1)$.

Recall that the sets $B_1, B_2, \ldots, B_n$ are distinct. Hence, $B_n \neq B_1$ (since $n \geqslant 2$).

The definition of Φ yields $\Phi(B_n) = B_n \setminus \{\varepsilon(B_n)\} = B_n \setminus \{u\}$ (since $\varepsilon(B_n) = u$). Thus, $u \notin \Phi(B_n)$ and $B_n \setminus \{u\} = \Phi(B_n)$.

However, $B_1 \in \mathcal{A}_+ = \{P \in \mathcal{A} \mid \varepsilon(P) \in P\}$. In other words, B_1 is a $P \in \mathcal{A}$ satisfying $\varepsilon(P) \in P$. In other words, B_1 is an element of $\mathcal{A}$ and satisfies $\varepsilon(B_1) \in B_1$. Now, $u = \varepsilon(B_1) \in B_1$. The same argument (applied to B_n instead of B_1) yields $u \in B_n$ (since $u = \varepsilon(B_n)$). Hence,

$$B_n = \underbrace{(B_n \setminus \{u\})}_{=\Phi(B_n)} \cup \{u\} = \Phi(B_n) \cup \{u\}. \tag{17}$$

However, (14) says that $\Phi(B_n) \prec B_1$. In other words, we have

$$B_1 = \Phi(B_n) \cup \{b\} \tag{18}$$

for some $b \in B_1 \setminus \Phi(B_n)$ (by Definition 5.3). Consider this b. Combining $u \in B_1 = \Phi(B_n) \cup \{b\}$ with $u \notin \Phi(B_n)$, we obtain

$$u \in (\Phi(B_n) \cup \{b\}) \setminus \Phi(B_n) \subseteq \{b\}.$$

In other words, $u = b$. Thus, (17) can be rewritten as $B_n = \Phi(B_n) \cup \{b\}$. Comparing this with (18), we obtain $B_n = B_1$. This contradicts $B_n \neq B_1$.

Forget that we fixed $(B_1, B_2, \ldots, B_n)$. We thus have found a contradiction whenever $(B_1, B_2, \ldots, B_n)$ is a tuple of distinct sets $B_1, B_2, \ldots, B_n \in \mathcal{A}_+$ with the property that $n \geqslant 2$ and that (13) and (14). Hence, there exists no such tuple. In other words, the complete matching $(\mathcal{A}_-, \mathcal{A}_+, \Phi)$ is acyclic. Therefore, the simplicial complex $\mathcal{A}$ has an acyclic complete matching, and thus is collapsible (by Definition 5.4 **(c)**). This finishes the proof of Theorem 5.5. □

The analogue of Theorem 5.5 for vertex-infection (instead of usual infection) also holds (with the same proof). However, Theorem 5.5 cannot be lifted to the generality of Theorem 4.2, since $\mathcal{A}$ will generally not be a simplicial complex unless Lemma 2.3 holds.

6. Open questions

I shall now comment on two natural directions of research so far unexplored.

6.1. The Alexander dual

Any simplicial complex has an *Alexander dual*, which is defined as follows:

Definition 6.1. Let E be a finite set. Let $\mathcal{A}$ be a simplicial complex on ground set E. Then, we define a new simplicial complex $\mathcal{A}^\vee$ on ground set E by

$$\mathcal{A}^\vee = \{F \subseteq E \mid E \setminus F \notin \mathcal{A}\}.$$

(That is, $\mathcal{A}^\vee$ consists of those subsets of E whose complements don't belong to $\mathcal{A}$.) This simplicial complex $\mathcal{A}^\vee$ is called the *Alexander dual* of $\mathcal{A}$.

It is well-known that a simplicial complex $\mathcal{A}$ and its Alexander dual $\mathcal{A}^\vee$ share many properties; in particular, the reduced homology of $\mathcal{A}$ is isomorphic to the reduced cohomology of $\mathcal{A}^\vee$ (see, e.g., [BjoTan09, Theorem 1.1]). However, the collapsibility and the homotopy types of $\mathcal{A}$ and $\mathcal{A}^\vee$ are not always related. Thus, the following question is suggested but not answered by Theorem 5.5:

Question 6.2. Let us use the notations from Section 1 as well as Definition 2.1. Let G be any subset of E. Define $\mathcal{A}$ as in (10). Is the simplicial complex

$$\mathcal{A}^{\vee} = \{F \subseteq E \ \mid \ G \subseteq \operatorname{Shade}(E \setminus F)\}$$

collapsible? Is it contractible?

6.2. Several vertices v

Elser's nuclei-based viewpoint in [Elser84] (and [DHLetc19, Conjecture 9.1]) suggests yet another question.

Our definition of Shade F (Definition 2.1), and the underlying notion of "infecting" an edge, implicitly relied on the choice of vertex v. It thus is advisable to rename the set Shade F as $\operatorname{Shade}_v F$ and combine such sets for different values of v. In particular, we can define

$$\mathcal{A}_U = \{F \subseteq E \ \mid \ G \not\subseteq \operatorname{Shade}_v F \text{ for some } v \in U\}$$

for any subset U of V. This $\mathcal{A}_U$ is a simplicial complex (being the union of a family of simplicial complexes), and thus we can ask the same questions about it as we did about $\mathcal{A}$:

Question 6.3. What can we say about the homotopy and discrete Morse theory of $\mathcal{A}_U$? What about its Alexander dual?

For $G = E$ and $|U| > 0$, this simplicial complex $\mathcal{A}_U$ is the Alexander dual of the "U-nucleus complex" Δ_U^G from [DHLetc19, Definition 3.2] (when G is connected). If [DHLetc19, Conjecture 9.1 for $|U| > 1$] is correct, then the homology of $\mathcal{A}_U$ with real coefficients should be concentrated in a single degree; this suggests the possible existence of an acyclic partial matching with all critical faces in one degree.

Acknowledgments

I thank Anders Björner, Galen Dorpalen-Barry, Dmitry Feichtner-Kozlov, Patricia Hersh, Vic Reiner, Tom Roby and Richard Stanley for insightful conversations. A referee has greatly simplified the proof of Theorem 5.5. This research was supported through the programme "Oberwolfach Leibniz Fellows" by the Mathematisches Forschungsinstitut Oberwolfach in 2020. I am deeply grateful to the Institute for its hospitality during a none-too-hospitable time.

References

[BenQui08] Arthur T. Benjamin and Jennifer J. Quinn, *An Alternate Approach to Alternating Sums: A Method to DIE for*, The College Mathematics Journal, Volume 39, Number 3, May 2008, pp. 191-202(12).

[BjoTan09] Anders Björner, Martin Tancer, *Combinatorial Alexander Duality – A Short and Elementary Proof*, Discrete Comput Geom (2009) 42, pp. 586–593.

[DHLetc19] Galen Dorpalen-Barry, Cyrus Hettle, David C. Livingston, Jeremy L. Martin, George Nasr, Julianne Vega, Hays Whitlatch, *A positivity phenomenon in Elser's Gaussian-cluster percolation model*, arXiv:1905.11330v5, Journal of Combinatorial Theory, Series A, **179**:105364, April 2021, doi:10.1016/j.jcta.2020.105364.

[Elser84] Veit Elser, *Gaussian-cluster models of percolation and self-avoiding walks*, J. Phys. A: Math. Gen. **17** (1984), pp. 1515–1523.

[Grinbe20] Darij Grinberg, *The Elser nuclei sum revisited*, detailed version of the present paper. Also available as an ancillary file at arXiv:2009.11527v6.

[Kozlov20] Dmitry N. Kozlov, *Organized Collapse: An Introduction to Discrete Morse Theory*, Graduate Studies in Mathematics **207**, AMS 2020.

[Sagan20] Bruce Sagan, *Combinatorics: The Art of Counting*, Graduate Studies in Mathematics **210**, AMS 2020.
See `https://users.math.msu.edu/users/bsagan/Books/Aoc/final.pdf` for a preprint.

[Stanle11] Richard P. Stanley, *Enumerative Combinatorics, volume 1*, 2nd edition, version of 15 July 2011.
`http://math.mit.edu/~rstan/ec/` .
See `http://math.mit.edu/~rstan/ec/` for errata.