Geometric and Electrical Properties of Graphs-with-Boundary

David Jekel

July 8, 2015

Contents

1 Graphs with Boundary

1.1 The Category of Graphs

A *graph* G consists of two sets $V = V(G))$ and $E = E(G)$, a function $\iota : E \to V$, and a function $^-\colon E \to E$ with $\overline{e} \neq e$ and $\overline{\overline{e}} = e$. For $e \in E$, $\iota(e)$ is the *initial vertex* of e and $\tau(e) = \iota(\overline{e})$ is the *terminal vertex*. Let E' be the set of pairs $\{e, \overline{e}\}$. An *edge* is an element of E' and an *directed edge* is an element of E.

We say a directed edge e is *incident* to a vertex p if p is the initial vertex of e, and a non-directed edge is *incident* to p if p is one of its endpoints. For a vertex p, $\iota^{-1}(p)$ is the set of directed edges incident to p, and its cardinality is called the *valence* of p.

We will work with both finite and infinite graphs, but assume that all vertices have finite valence. Since any connected component of an infinite graph must be countable anyway, we will also assume that $V(G)$ and $E(G)$ are countable.

A morphism of graphs $f : G_1 \to G_2$ consists is a pair of functions $f_V : V(G_1) \to V(G_2)$ and $f_E : E(G_1) \to E(G_2)$ such that $f_V(\iota(e)) = \iota(f_E(e))$ and $f_E(\overline{e}) = \overline{f_E(e)}$. Henceforth, we will write f for f_E and f_V because it is simpler and causes no confusion.

1.2 The Category of Graphs with Boundary

A *graph-with-boundary* G is a graph together with a partition of V into two sets B and I. The elements of B are called *boundary vertices* and those of I are *interior vertices*. We will call a graph-with-boundary a "bgraph" for short.

A bgraph morphism $f : G_1 \to G_2$ of is a graph morphism with the following properties:

- f maps interior vertices to interior vertices. That is, if $p \in I(G_1)$, then $f(p) \in I(G_2)$.
- For each interior vertex p, the restriction of f to a map $\iota^{-1}(p) \to \iota^{-1}(f(p))$ has constant fiber size ≥ 1.

The reader may verify that this class of graph morphisms is closed under composition and includes the identity.

A bgraph morphism is loosely analogous to an analytic function on a plane domain or Riemann surface. Mapping interior vertices to interior vertices is analogous to being an open map. A vertex v where the map $\iota^{-1}(p) \to \iota^{-1}(f(p))$ is bijective is analogous to a point z_0 where an analytic function ϕ has nonzero derivative and is thus bijective in a neighborhood of z_0. If the map $\iota^{-1}(p) \to \iota^{-1}(f(p))$ is n-to-1, then v is analogous to a point z_0 where $\phi(z) - \phi(z_0)$ has a zero of order n. With this in mind, we say that a vertex v has *order* n if the fibers of $\iota^{-1}(p) \to \iota^{-1}(f(p))$ have cardinality n.

1.3 Basic Constructions in the Category of Bgraphs

Some special types of morphisms: A bgraph morphism $f : G_1 \to G_2$ is called

- *injective/surjective* if the maps $V(G_1) \to V(G_2)$ and $E(G_1) \to E(G_2)$ are injective/surjective.
- *locally injective/surjective* if the maps $\iota^{-1}(p) \to \iota^{-1}(f(p))$ are injective/surjective for all v.
- *a covering map* if it is surjective and locally bijective, and sends boundary vertices to boundary vertices.

Sub-bgraphs: We say G_1 is a *sub-bgraph* of G_2 if $V(G_1) \subset V(G_2)$, $E(G_1) \subset E(G_2)$, $I(G_1) \subset I(G_2)$, and if $p \in I(G_1)$, then $\iota^{-1}(p)$ (with respect to G_2) is contained in $E(G_1)$. Or equivalently, G_1 is a sub-bgraph of G_2 if it is a subgraph and the inclusion $G_1 \to G_2$ is a bgraph morphism.

If S_1 and S_2 are sub-bgraphs of G, we define $S_1 \cup S_2$ by

$$V(S_1 \cup S_2) = V(S_1) \cup V(S_2), \quad E(S_1 \cup S_2) = E(S_1) \cup E(S_2), \quad I(S_1 \cup S_2) = I(S_1) \cup I(S_2).$$

And $S_1 \cap S_2$ is defined in the same way.

If $f : G_1 \to G_2$ is a bgraph morphism and S_2 is a subgraph of G_2, then we define $f^{-1}(S_2)$ as follows:

$$V(f^{-1}(S_2)) = f^{-1}(V(S_2)), \quad E(f^{-1}(S_2)) = f^{-1}(E(S_2)), \quad I(f^{-1}(S_2)) = f^{-1}(E(S_2)) \cap I(G_1);$$

the reader should verify that this is a subgraph of G_1.

Pull-backs: The category of bgraphs has pull-backs (constructed in the typical way for concrete categories). Suppose $f_1 : G_1 \to G$ and $f_2 : G_2 \to G$ are bgraph morphisms. Then the *pull-back* or *fiber product* $G_1 \times_G G_2$ is constructed as follows:

- $V(G_1 \times_G G_2)$ is the set of pairs $(p_1, p_2) \in V(G_1) \times V(G_2)$ such that $f_1(p_1) = f_2(p_2)$.
- $E(G_1 \times_G G_2)$ is the set of pairs $(e_1, e_2) \in E(G_1) \times E(G_2)$ such that $f_1(e_1) = f_2(e_2)$.
- (p_1, p_2) is interior if and only if both p_1 and p_2 are interior.
- $\iota((e_1, e_2)) = (\iota(e_1), \iota(e_2))$ and $\overline{(e_1, e_2)} = (\overline{e_1}, \overline{e_2})$.

The reader may verify that the evident maps $\pi_1 : G_1 \times_G G_2 \to G_1$ and $\pi_2 : G_1 \times_G G_2 \to G_2$ are bgraph morphisms.

The pull-back satisfies the following the universal property: For any bgraph H and any maps $g_1 : H \to G_1$ and $g_2 : H \to G_2$ such that $f_1 \circ g_1 = f_2 \circ g_2$, there exists a unique bgraph morphism $\phi : H \to G_1 \times_G G_2$ such that $g_j = \pi_j \circ \phi$, as depicted here:

$$\begin{tikzcd}
H \arrow[dr, dashed, "\exists!\phi"] \arrow[drr, "g_1"] \arrow[ddr, "g_2"'] & & \\
& G_1 \times_G G_2 \arrow[r, "\pi_1"'] \arrow[d, "\pi_2"] & G_1 \arrow[d, "f_1"] \\
& G_2 \arrow[r, "f_2"'] & G
\end{tikzcd}$$

This universal property characterizes $G_1 \times_G G_2$ up to isomorphism.

Remark. Although the product exists in the category of graphs, it does not generally exist in the category of graphs with boundary (exercise).

Coproducts and Push-Outs: The coproduct or disjoint union $G_1 \coprod G_2$ of two bgraphs G_1 and G_2 is formed by taking the disjoint union of the sets of vertices, edges, and interior vertices of G_1 and G_2. It is only well-defined up to isomorphism. The evident inclusions $i_1 : G_1 \to G_1 \coprod G_2$ and $i_2 : G_2 \to G_1 \coprod G_2$ are bgraph morphisms. $G_1 \coprod G_2$ satisfies the universal property that for bgraph morphisms $f_1 : G_1 \to G$ and $f_2 : G_2 \to G$, there is a unique corresponding bgraph morphism $f : G_1 \coprod G_2 \to G$ such that $f_1 = f \circ i_1$ and $f_2 = f \circ i_2$.

If we have two maps $f_1 : G \to G_1$ and $f_2 : G \to G_2$, we can define a *push-out* $G_1 \coprod_G G_2$ as follows: The vertices of $G_1 \coprod_G G_2$ are the vertices of the disjoint union, modulo the equivalence relation generated by the relations $f_1(p) \sim f_2(p)$ for $p \in V(G)$. We perform the same operations to obtain $E(G_1 \coprod_G G_2)$. A vertex of $G_1 \coprod_G G_2$ is declared to be interior if at least one representative of the equivalence class is interior. If we assume that f_1 and f_2 are locally injective, then the graph morphisms $i_1 : G_1 \to G_1 \coprod_G G_2$ and $i_2 : G_2 \to G_1 \coprod_G G_2$ will be bgraph morphisms, because we can guarantee that the maps $\iota^{-1}(p) \to \iota^{-1}(i_j(p))$ have constant fiber size 1 for each interior p in G_1 or G_2. If we allow points of order > 1 in f_1 and f_2, then this fails in general.

1.4 Paths and Connections

For a graph G, a *path* is a sequence of vertices $p_0, \dots, p_K$ and oriented edges $e_1, \dots, e_K$ such that $\iota(e_k) = p_{k-1}$ and $\tau(e_k) = p_k$. We allow a "trivial" path with one vertex and no edges. A path is an *embedded path* if the vertices $p_0, \dots, p_K$ are distinct and the non-oriented edges in the path are distinct. A *boundary-to-boundary path* is an embedded path such that p_0 and p_K are boundary vertices and the other vertices are interior. A *cycle* is a non-trivial path such that the edges are distinct, and the vertices $p_0, \dots, p_{K-1}$ are distinct with $p_K = p_0$.

We a *path through the graph* is an embedded path in which the vertices $p_1, \dots, p_{K-1}$ are interior. We say two vertices p and q are *connected through the graph* if there exists such a path with $p_0 = p$ and $p_K = q$. We say two directed edges e and e' are connected through the graph if there exists such a path with $e_1 = e$ and $e_K = e'$. We say two edges are connected through the graph if we can choose an orientation for each one such that the resulting directed edges are connected through the graph.

A graph is *connected* if for any two vertices p and q, there exists a path from p to q. For any graph, there exist connected subgraphs $G_1, \dots, G_N$, called *components*, such that $V(G_1), \dots, V(G_N)$ are a partition of $V(G)$, and $E(G_1), \dots, E(G_N)$ are a partition of $E(G)$, and $B(G_1), \dots, B(G_N)$ are a partition of $B(G)$.

2 Networks

To motivate our general definition of electrical networks, we summarize the types of networks we are primarily interested in:

1. It is standard to consider real-valued linear electrical networks. We are given a bgraph G together with a positive real number a_e for each edge (called the conductance). For a "potential" function $u : V \to \mathbb{R}$, the current on edge is given by $c_e = a_e(u(\iota(e)) - u(\tau(e)))$ (Ohm's Law). The potential is *harmonic* if the net current $\sum_{e \in \iota^{-1}(p)} c_e$ is 0 for each interior vertex p (Kirchhoff's Law).

2. In a "nonlinear conductance network," each edge is assigned a *conductance function* $\gamma_e : \mathbb{R} \to \mathbb{R}$ rather than a positive number. We define $c_e = \gamma_e(u(\iota(e)) - u(\tau(e)))$. Again, a potential is harmonic if the net current on interior vertices is zero. This provides a simple model for non-ohmic resistors or semiconductors; in the physically relevant cases, γ_e is weakly increasing and $\gamma_e(0) = 0$.

3. In a "nonlinear resistance network," each edge is assigned a *resistance function* $\rho_e : \mathbb{R} \to \mathbb{R}$ (if we wish, weakly increasing and zero-preserving). We consider potential/current functions such that $u(\iota(e)) - u(\tau(e)) = \rho_e(c(e))$. In this case, c is no longer a function of u, and u is not a function of c.

4. Returning to the linear case, we can generalize in a more algebraic direction: Instead of considering the conductance to be a positive number, we could also allow it to be negative or complex. Or we could replace real numbers by an arbitrary field or even ring.

5. More generally, suppose each γ_e is a unit of a ring R. Let M be an R-module, and consider a potential function $u : V \to M$ and current function $c : E \to M$ with $c(e) = \gamma_e \cdot (u(\iota(e)) - u(\tau(e)))$.

The minimal amount of structure we need to discuss these types of electrical networks is

- the ability to add potentials and currents together (our u and c must take values in some abelian group M),
- a way to specify when a potential drop $u(\iota(e)) - u(\tau(e))$ and current $c(e)$ are compatible for a given edge e (so we have some relation in $M \times M$ specifying the compatible potential drop/current pairs).

Let's formalize this.

For a set S we define the *category of S-labelled graphs* as follows: An object is a pair (G, ℓ) where $G = (V, E)$ is a graph and together ℓ is a function $E \to S$ called a *labelling*. A morphism of S-labelled graphs is a morphism of graphs

$f : G_1 \to G_2$ that preserves the labelling, meaning that $\ell_2(f(e)) = \ell_1(e)$ for any $e \in E(G_1)$. The category of S-labelled bgraphs is defined in the same way.

We can describe each of the above types of networks by labelling the edges with different things. If we wish to include all these networks under one umbrella, we can use the following general definition:

Let M be an abelian group (written additively). A *network* Γ (taking values in M) is a bgraph where each edge is labelled by a relation $\Theta \subset M \times M$, such that $\Theta_{\overline{e}} = -\Theta_e$. Networks form a subcategory of $\mathcal{P}(M \times M)$-labelled bgraphs. We denote the labelling by Θ.

An *M-valued potential function* on a bgraph G is a map $u : V \to M$. To make notation neater, we will write $u_p = u(p)$ for $p \in V$. An *M-valued current function* on a G is a map $c : E \to M$ such that

- $c_{\overline{e}} = -c_e$.
- The net current $\sum_{e \in \iota^{-1}(p)} c_e = 0$ for each interior vertex p.

Simiiarly, we have written $c_e = c(e)$. For a network Γ, a potential u and current c are *compatible* if for each e, $(u_{\iota(e)} - u_{\tau(e)}, c_e) \in \Theta_e$.

The condition $\Theta_{\overline{e}} = -\Theta_e$ guarantees that

$$(u_{\iota(e)} - u_{\tau(e)}, c_e) \in \Theta_e \text{ if and only if } (u_{\iota(\overline{e})} - u_{\tau(\overline{e})}, c_{\overline{e}}) \in \Theta_{\overline{e}}.$$

2.1 Harmonic Functions

Fix an abelian group M. A *harmonic function* on a network Γ is a compatible potential/current pair (u, c). Let $\mathcal{H}_\Gamma$ be the set of harmonic functions on Γ.

Suppose $f : \Gamma_1 \to \Gamma_2$ is a network morphism. If u is a potential function on Γ_2, we can define f^*u (or "$u \circ f$") on Γ_1, by $f^*u_p = u_{f(p)}$. Similarly, we can define $f^*c_e = c_{f(e)}$. If (u, c) is harmonic on Γ_2, then (f^*u, f^*c) is harmonic on Γ_1. To prove this, we first show f^*c is actually a current function–that is, the net current on each interior vertex is zero. Suppose p is an interior vertex of Γ_1 and suppose that the fibers of $\iota^{-1}(p) \to \iota^{-1}(f(p))$ all have cardinality n (as in the definition of bgraph morphisms). Then

$$\sum_{e \in \iota^{-1}(p)} f^*c_e = n \sum_{e \in \iota^{-1}(f(p))} c_e,$$

which is zero since $f(p)$ is interior. Then to show that u and c are compatible, note that for each edge e,

$$(f^*u_{\iota(e)} - f^*u_{\tau(e)}, f^*c_e) = (u_{\iota(e)} - u_{\tau(e)}, c_e) \in \Theta_{f(e)} = \Theta_e.$$

Remark. This is analogous to the result from complex analysis that if ϕ is analytic and u is harmonic, then $u \circ \phi$ is harmonic, and it is what motivated our definition of bgraph morphisms.

Thus, $\Gamma \mapsto \mathcal{H}_\Gamma$ defines a contravariant functor from the category of networks to the category of sets. For the various types of networks described earlier, it could also define a functor into a category with more structure such as Top or R-mod.

2.2 Boundary Data; Forward and Inverse Problems

Let Γ be a network with values in an abelian group M. For a potential function u, the *boundary potential function* is $u|_B \in M^B$. For a current function c, the *boundary net current function* $\psi \in M^B$ is given by $\psi_p = \sum_{e \in \iota^{-1}(p)} c_e$. If ϕ is the boundary voltage of u and ψ is the boundary current of c, then $(\phi, \psi) \in M^B \times M^B$ is called the *boundary data* of (u, c). Define a function $\Phi : \mathcal{H}_\Gamma \to M^B \times M^B$ by mapping a harmonic function to its boundary data. We define the *set of boundary data* L as the image of Φ.

We will consider the following problems:

- **The Dirichlet Problem:** For $\phi \in M^B$, does there exist a harmonic function with boundary potential ϕ? Is it unique?
- **The Neumann Problem:** For $\psi \in M^B$, does there exist a harmonic function with boundary current ψ? It is unique?
- **Regularity:** If the solutions to the Dirichlet and Neumann problems exist, do they depend "nicely" on ϕ or ψ? More generally, how "nice" is L? (The notion of "nice" depends on the type of network.)
- **The Inverse Problem:** Let $\Gamma = (G, \Theta)$ be a network. Is Θ uniquely determined by G and L?

3 Layer-stripping to Solve the Inverse Problem

3.1 General Strategy

A boundary edge on a bgraph G is a directed edge (or simply an edge) such that both endpoints are boundary vertices.

A boundary spike on a bgraph G is a oriented edge e such that $\iota(e)$ is a boundary vertex of degree 1. We allow $\tau(e)$ to be interior or boundary. If $\tau(e)$ is boundary, we say the spike is degenerate. We say that an unoriented edge is a boundary spike if at least one orientation of it is a boundary spike.

Given an oriented boundary spike e, we can define a new bgraph G' by $V(G') = V(G) \setminus \{\iota(e)\}$, $E(G') = E(G) \setminus \{e, \overline{e}\}$, $I(G') = I(G) \setminus \{\tau(e)\}$. Then we say *$G'$ is obtained from G by contracting the boundary spike e*. If we have a collection of boundary spikes with distinct endpoints, then we can contract multiple spikes at the same time, even infinitely many spikes.

Given an oriented boundary edge e of a bgraph G, we can define a new bgraph G' by *deleting the boundary edge*, that is $V(G') = V(G)$, $I(G') = I(G)$, and $E(G') = E(G) \setminus \{e, \overline{e}\}$. We can also delete multiple boundary edges at the same time.

The strategy for solving the inverse problem employed by Curtis/Morrow and Will Johnson was to recover the conductance functions on boundary spikes and boundary edges first, then work one's way inward, as it were, stripping layers off the graph until nothing is left. We describe the process roughly as follows:

- Given a graph G with a boundary spike or boundary edge, figure out how to recover its conductance function from L.
- Remove the boundary spike or boundary edge from the graph to obtain a new graph G'. Find the set of boundary data L' for the new graph.
- Repeat.

This does not work for all graphs; for instance, some graphs do not have any boundary spikes or boundary edges. To formalize the process and state conditions when it works, we need to address several questions:

1. How do we recover the conductances of boundary spikes and boundary edges?
2. How can we find L' from L?
3. What sorts of graphs can be "layer-stripped" so as to remove all the edges? How do we know there is a boundary spike or boundary edge at each step of the process? For infinite graphs, how can we make sure that our layer-stripping exhausts all the edges?

This chapter will develop the formal machinery to describe the layer-stripping process, and state purely geometric conditions involving "scaffolds" to guarantee recoverability over BZCF. These geometric conditions may at first appear unmotivated and hard to verify unless the reader is familiar with enough examples, but in later chapters, we will describe various general ways of constructing scaffolds.

One key consequence will be that if $f : G \to G'$ is a locally injective bgraph morphism, then the information propagation and layer-stripping structures on G' can be pulled back to G by taking preimages. Hence, if this method of solving the inverse problem works on G', it also works on G, which is potentially a much larger graph (such as the universal cover). We will be careful to formulate our definitions in such a way that this pulling-back is straightforward and functorial, and works equally well for finite or infinite graphs.

3.2 Scaffolds

A *scaffold* $\mathcal{S}$ on G consists of

- A (strict) partial order $\prec$ on E',
- A partition of E' into two sets $\operatorname{Vert}\mathcal{S}$ and $\operatorname{Hor}\mathcal{S}$, whose elements are called respectively *vertical* and *horizontal edges*.
- Two functions $t, b : \operatorname{Vert}\mathcal{S} \to V$ which assign a "top" endpoint $t(e)$ and a "bottom" endpoint $b(e)$ to each $e \in \operatorname{Vert}\mathcal{S}$, which are distinct endpoints of e.

satisfying the following conditions:

1. Every subset of E' has a minimal element.

2. If $e \in \operatorname{Vert} \mathcal{S}$ and e' are incident at $t(e)$, then $e \prec e'$.

3. If $e \in \operatorname{Vert} \mathcal{S}$ and e' are incident at $b(e)$, then $e' \prec e$.

4. If p_1 and p_2 are interior vertices incident to e_1 and e_2 respectively, with $e_1 \preceq e_2$, then either $p_1 \in b(\operatorname{Vert} \mathcal{S})$ or $p_2 \in t(\operatorname{Vert} \mathcal{S})$.

Remark. Condition (1) is only necessary for infinite graphs.

Some consequences of the definition help to clarify the geometric picture: Because of the comparison conditions, there are at most two vertical edges incident to a given vertex. Thus, if we start at a given vertex p, we can form a unique increasing path of vertical edges, which will either terminate or continue infinitely. And it could terminate at an interior vertex or boundary vertex. Similarly, we can form a decreasing path of vertical edges. This path must terminate by (1). So our vertex p is on a unique increasing path in which all the edges are vertical beginning at a vertex q.

Let $\mathcal{S}$ be a scaffold on a bgraph G. Define $\operatorname{Top} \mathcal{S}$ as the set of edges e such that $e \succeq e'$ for some e' with an endpoint in $I \setminus t(\operatorname{Vert} \mathcal{S})$. Let $\operatorname{Bot} \mathcal{S}$ be the set of edges e such that $e \preceq e'$ for some e' with an endpoint in $I \setminus b(\operatorname{Vert} \mathcal{S})$. Define $\operatorname{Mid} \mathcal{S} = E' \setminus (\operatorname{Top} \mathcal{S} \cup \operatorname{Bot} \mathcal{S})$. Note condition (4) implies that $\operatorname{Top} \mathcal{S}$ and $\operatorname{Bot} \mathcal{S}$ are disjoint.

Scaffolds behave nicely with respect to locally injective bgraph morphisms. Let $\operatorname{Scaf} G$ be the set of scaffolds on a bgraph G. Suppose $f : G_1 \to G_2$ is a locally injective bgraph morphism. Suppose $\mathcal{S} \in \operatorname{Scaf} G_2$. Then define $f^*\mathcal{S}$ as follows:

- Set $e_1 \prec e_2$ in $f^*\mathcal{S}$ if and only if $f(e_1) \prec f(e_1)$ in $\mathcal{S}$.
- Let $\operatorname{Vert} f^*\mathcal{S} = f^{-1}(\operatorname{Vert} \mathcal{S})$ and $\operatorname{Hor} f^*\mathcal{S} = f^{-1}(\operatorname{Hor} \mathcal{S})$.
- Since the map is locally injective, we can define b, t for $f^*(\mathcal{S})$ such that $f(b(e)) = b(f(e))$ and $f(t(e)) = t(f(e))$.

The reader may verify that $f^*\mathcal{S}$ satisfies properties (1) through (4). Thus, we have

Proposition 3.1. *$G \mapsto \operatorname{Scaf} G$ defines a contravariant functor from the category of bgraphs and locally injective bgraph morphisms to Set. It also satisfies $f^{-1}(\operatorname{Mid} \mathcal{S}) \subset \operatorname{Mid} f^*\mathcal{S}$.*

Proof. Straightforward and left to the reader. □

3.3 Recovery of Boundary Spikes/Edges

Lemma 3.2. *Let G be a bgraph. Suppose that either*

1. e_0 is a boundary spike and there is a scaffold $\mathcal{S}$ with $e_0 \in \operatorname{Hor} \mathcal{S} \cap \operatorname{Mid} \mathcal{S}$, or

2. e_0 is a boundary edge and there is a scaffold $\mathcal{S}$ with $e_0 \in \operatorname{Vert}\mathcal{S} \cap \operatorname{Mid}\mathcal{S}$.

For any BZCF network Γ on G, Θ_{e_0} is uniquely determined by L over BZCF.

To prove the lemma, we need the following definitions:
Let $T \subset E'$. The *subgraph G_T induced by T* is defined as follows:

- $E'(G_T) = T$.
- $V(G_T)$ is the set of vertices incident to edges in T.
- A vertex is interior in G_T if and only if it is interior in G and all the edges incident to it are in T.

A sub-bgraph $G' \subset G$ is induced if and only if any vertex $p \in V(G') \cap I(G)$ with all edges incident to it contained in $E(G')$ must be interior in G'.

Let $\mathcal{S}$ be a scaffold on G. We say that $G' \subset G$ is a *lower sub-bgraph* if $e \prec e' \in E(G')$ implies $e \in E(G')$. We say that $G' \subset G$ is an *upper sub-bgraph* if $e \succ e' \in E(G')$ implies $e \in E(G')$.

Proof. Consider the case of a boundary spike first. Let p be the boundary vertex of the spike, q the interior vertex. Choose $t \in M$. Let

- Γ_0 be the subnetwork induced by $\{e \prec e_0\}$.
- Γ_1 be the subnetwork induced by $\{e \not\succeq e_0\}$.
- Γ_2 be the subnetwork induced by $\{e \not\succ e_0\}$.

Note $\Gamma_0 \subset \Gamma_1 \subset \Gamma_2 \subset \Gamma$.

Claim 1: *There exists a harmonic function on Γ which is identically zero on $V(\Gamma_1)$ with $u_p = t$.*

Because $e_0 \in \operatorname{Mid}\mathcal{S}$, we know q is the bottom vertex of some vertical edge, hence not all edges of Γ incident to q are in Γ_2, so that q is a boundary vertex of Γ_1. By assumption q is the only neighbor of p. Thus, the potential u_2 which is t at p and 0 everywhere else is harmonic, that is, it has a compatible current function c_2.

Consider the set $\mathcal{Z}$ of pairs (Σ, v), where Σ is an induced lower subnetwork of Γ, $\Gamma_1 \subset \Sigma$, and v is a harmonic potential on Σ which equals u_1 on Γ_1. Let $\mathcal{Z}$ be partially ordered by setting $(\Sigma_1, v_1) \leq (\Sigma_2, v_2)$ if $\Sigma_1 \subset \Sigma_2$ and $v_2|_{V(\Sigma_1)} = v_1$. Note that $(\Gamma_1, u_1) \in \mathcal{Z}$. To apply Zorn's lemma, note that every totally ordered subset $\mathcal{C}$ of $\mathcal{Z}$ has an upper bound. Indeed, for two networks (Σ, v) and $(\Sigma', v) \in \mathcal{C}$, the corresponding harmonic functions agree on the overlap, and hence they produce a well-defined harmonic function v^* on $\Sigma^* = \bigcup_{(\Sigma,v)\in\mathcal{C}} \Sigma$, and (Σ^*, v^*) is an upper bound for $\mathcal{C}$. Hence, $\mathcal{Z}$ has a maximal element (Σ^*, v^*).

Suppose for the sake of contradiction that $\Sigma^* \subsetneq \Gamma$. Then by condition (1) in the scaffold definition, $E(\Gamma) \setminus E(\Sigma^*)$ has a minimal element e_1. Let Σ^+ be the subnetwork of Γ induced by $E'(\Sigma^*) \cup \{e_1\}$. Now we consider two cases:

- Suppose e_1 is vertical. Then $t(e_1)$ cannot be in Σ^* since Σ^* is a lower subgraph. Because e_1 is not in Γ_0, we must have $e_1 \succ e_0$, and hence $e_1 \notin \operatorname{Bot}\mathcal{S}$. This implies by definition of bottom that $t(e_1)$ is either a boundary vertex of Γ or $t(e_1) \in b(\operatorname{Vert}\mathcal{S})$; hence, $t(e_1)$ is a boundary vertex of Σ^+. Since Θ_{e_1} is BZCF we can extend v^* harmonically to Σ^+ by choosing a potential on $t(e_1)$ so as to make the net current on $b(e_1)$ zero (whether it's interior in Γ or not). This contradicts the maximality of (Σ^*, v^*).
- Suppose e_1 is horizontal. As before, since $e_1 \notin \operatorname{Bot}\mathcal{S}$, its endpoints must be boundary vertices of Σ^+. Hence, if v^* is harmonic on Σ^*, it also defines a harmonic function on Σ^+, and this contradicts maximality of Σ^*.

Hence, $\Sigma^* = \Gamma$ and Claim 1 is proved.

Claim 2: *Any harmonic function with potential zero on $B(\Gamma) \cap V(\Gamma_0)$ and net current zero on $B(\Gamma) \cap b(\operatorname{Vert}\mathcal{S} \cap E'(\Gamma_0))$ must be identically zero on Γ_0.*

Suppose for the sake of contradiction that (u, c) is a harmonic function that violates this. Then there exists a minimal edge e_1 in Γ_0 with nonzero potential on some endpoint.

- Suppose e_1 is vertical. Then all edges incident to $b(e_1)$ other than e_1 itself are less than e_1 by (3), and so they have potential zero on both endpoints by minimality of e_1. In particular, $u_{b(e_1)} = 0$. If $b(e_1) \in I(\Gamma)$, it must have net current zero. Otherwise, we have $b(e_1) \in B(\Gamma) \cap b(\operatorname{Vert}\mathcal{S} \cap E'(\Gamma_0))$, so it has net current zero by assumption. Since there is zero current on the other edges incident to $b(e_1)$, there must be zero current on e_1, and so there is zero potential on $t(e_1)$, a contradiction.
- Suppose e_1 is horizontal. If an endpoint p of e_1 is a boundary vertex of Γ, then $u_p = 0$ by assumption. If it is interior in Γ, then it must be in $t(\operatorname{Vert}\mathcal{S})$ because $e_1 \prec e_0 \in \operatorname{Mid}\mathcal{S}$. If $p = t(e_2)$, that implies $u_p = 0$ by minimality of e_1. In either case, the endpoints of e_1 must have potential zero, which contradicts our choice of e_1.

Conclusion: Choose any harmonic function (u, c) satisfying Claim 2 which has potential $u_p = t$. We know such a function exists by Claim 1. By Claim 2, u_q must be zero since $q \in V(\Gamma_0)$. Since q is the only neighbor of p, this implies the net current on p is $-\gamma_{e_0}(t)$ (if e_0 is oriented with $\iota(e_0) = p$). Thus, by imposing the boundary conditions of Claim 2 and potential $u_p = t$ (no interior conditions), we obtain a unique net current on p which is $-\gamma_{e_0}(t)$. Since this holds for all BZCF networks, $\gamma_{e_0}(t)$ is uniquely determined by L over BZCF. Since t is arbitrary, γ_{e_0} is determined.

This concludes the case for a boundary spike. In the case of a boundary edge, the argument is the same with the following changes:

- Let $q = b(e_0)$, $p = t(e_0)$.
- To define u_1 on Γ_1, note $p, q \in B(\Gamma_1)$ and e_0 is the only edge incident to p in Γ_1. Define u_1 to be zero on Γ_0 and t at vertex p.

- We recover $\gamma_{e_0}(t)$ by noting that it is the net current on q. This is because by (3) all edges incident to q except e_0 are in Γ_0 and hence have current zero. $\square$

3.4 Reduction Operations and Layerable Graphs

In this section, we discuss systematically the process of removing all the edges in the graph through a sequence of boundary spike contractions and boundary edge deletions. A *reduction operation* is a transformation of a bgraph G into a subgraph G' such that

1. The edges removed are all boundary spikes or boundary edges of G.

2. The vertices removed are all boundary vertices of valence 0 or 1.

3. The only boundary vertices of G' that are interior in G are the endpoints of boundary spikes that were removed.

In other words, a reduction operation is some combination of contracting boundary spikes, deleting boundary edges, and deleting disconnected boundary vertices, such that each of the smaller operations affects at most one vertex. If there is exactly one boundary spike/ boundary edge / disconnected boundary vertex removed overall, then the reduction operation is called *simple*.

If $f : G \to G'$ is a locally injective bgraph morphism, and S is obtained from G' by a reduction operation, then $f^{-1}(S)$ is obtained from G by a reduction operation (easy casework left to the reader). However, a boundary spike contraction in G' may produce a disconnected boundary vertex deletion in G or some combination of boundary spike contraction and disconnected boundary vertex deletion in G. This is why the definition was phrased so as to allow mixing boundary spike contraction, boundary edge deletion, and disconnected boundary vertex deletion in one reduction operation.

A *(decreasing) filtration* of a graph G is a sequence of subgraphs $G = G_0 \supset G_1 \supset G_2 \supset \dots$ such that $\bigcap_{n=0}^{\infty} G_n = \varnothing$. If G_{n+1} is obtained from G_n by a reduction operation, then we the filtration is called a *layerable filtration* and the bgraph is said to be *layerable*. A *partial filtration* is a sequence of subgraphs $G = G_0 \supset G_1 \supset \dots$, and it is a partial layerable filtration if each subgraph is obtained from the previous one by a reduction operation.

If $f : G \to G'$ is a locally injective bgraph morphism and $G'_0, G'_1, \dots$ is a layerable filtration of G', then $f^{-1}(G_0)$, $f^{-1}(G_1)$, $\dots$ is a layerable filtration of G. Hence, layerability of G' implies layerability of G.

Layerability is related to scaffolds through the following lemma:

Lemma 3.3. *Let G be a bgraph with countably many edges. The following are equivalent:*

a. G admits a layerable filtration in which the reduction operations are simple.

b. G admits a layerable filtration (that is, G is layerable).

c. *There exists a scaffold* $\mathcal{S}$ *on* G *with* $\operatorname{Top} \mathcal{S} = \varnothing$.

d. *For any* $e \in E'(G)$*, there is a scaffold* $\mathcal{S}$ *on* G *with* $e \notin \operatorname{Top} \mathcal{S}$.

e. *For any* $e \in E'(G)$*, there is a finite partial layerable filtration* $G = G_0 \supset \cdots \supset G_n$ *with* $e \notin E(G_n)$.

Proof. (a) $\implies$ (b) is immediate.

(b) $\implies$ (c). Let $G = G_0 \supset G_1 \supset \dots$ be a layerable filtration. Then each edge e is in $E(G_{n_e}) \setminus E(G_{n_e+1})$ for some n_e. Define $\mathcal{S}$ as follows:

- $e \prec e'$ if and only if $n_e < n_{e'}$.
- e is vertical if it is a boundary spike of G_{n_e} and it is horizontal if it is a boundary edge of G_{n_e}.
- If e is a boundary spike in G_{n_e}, then $b(e)$ is the endpoint removed in the spike contraction and $t(e)$ is the other endpoint.

The reader may verify that all the conditions in the definition of scaffold are satisfied.

(c) $\implies$ (d) is immediate.

(d) $\implies$ (e). Define a new scaffold $\mathcal{S}'$ with the same vertical edges and t and b functions as in $\mathcal{S}$, but define the new partial order by taking the transitive closure of the relations defined by conditions (2) and (3) of the scaffold definition. (Thus, we are making as few edges comparable to each other as possible given our choice of vertical edges.) Every subset of E has a minimal element with respect to $\mathcal{S}$, which will automatically be minimal with respect to $\mathcal{S}'$.

I claim that for any $e \in E'(G)$, there are only finitely many edges $e \preceq e_0$ in $\mathcal{S}'$. If we suppose not, then there is a minimal edge e_0 for which the claim does not hold. There are only finitely many edges $e_1, \dots, e_n$ which incident to and less than e_0, and $\{e \preceq e_0\} = \bigcup_{j=1}^n \{e \preceq e_j\} \cup \{e_0\}$ since the relations (2) and (3) used to define our partial order only compare edges which are incident to each other. By minimality of e_0, $\{e \preceq e_j\}$ is finite, which implies $\{e \preceq e_0\}$ is finite, which is a contradiction.

Now choose e. Let $e_1, \dots, e_k = e$ be the edges $\preceq e$ in $\mathcal{S}'$. We can assume they are listed in some nondecreasing order. Let $G_0 = G$. Then e_1 is a minimal edge in G_0. The conditions in the definition of a scaffold force e_1 to be a boundary spike if it is vertical and a boundary edge if it is horizontal (similar reasoning to the lemma about recovery). Let G_1 be the graph formed by deleting/contracting this edge as appropriate. Then e_2 is a minimal edge in G_1, hence a boundary spike or boundary edge. So (e) follows by induction.

(e) $\implies$ (a). Observe: If S is a subgraph of G and e is a boundary spike of G, then it is also a boundary spike of S if it is actually contained in S. Hence, if G' is obtained from G by contracting the boundary spike, then either $S \cap G' = G'$ or else $S \cap G'$ is obtained from S by contracting the boundary spike. The same observation holds for boundary edges.

We assumed in §1 that our graphs have countably many edges, so we can write them in a sequence $e_1, e_2, \dots$. For each e_n, choose a k_n and a sequence of subgraphs $G = G_{n,1} \supset \cdots \supset G_{n,k_n}$ as in (e). Then consider the following filtration:

$$\begin{array}{llll} G = G_{1,1}, & G_{1,2}, & \ldots & G_{1,k_1}, \\ G_{1,k_1} \cap G_{2,1}, & G_{1,k_1} \cap G_{2,2}, & \ldots & G_{1,k_1} \cap G_{2,k_2} \\ G_{1,k_1} \cap G_{2,k_2} \cap G_{3,1}, & \ldots & G_{1,k_1} \cap G_{2,k_2} \cap G_{3,k_3} & \\ \cdots\cdots & & & \end{array}$$

The consecutive elements of this sequence, if they are not equal, are obtained by removing a boundary spike or boundary edge as a result of our earlier observation. Thus, we have a partial layerable filtration which removes all the edges in the graph. We can obtain a new filtration by replacing each reduction operation with two reduction operations–first remove the boundary spike or boundary edge according to our original partial filtration, then remove any disconnected boundary vertices. ☐

3.5 Electrical Properties of Reduction Operations

Suppose that G' is obtained from G by a reduction operation. Our goal is to show that the boundary data L' is uniquely determined by L and the conductance functions of the edges removed in the reduction.

Lemma 3.4. *Suppose that G' is obtained from G by contracting some non-degenerate boundary spikes. Let Γ be a BZCF network on G and let Γ' be the corresponding network on G'. Let L and L' be the sets of boundary data. Then*

- *The inclusion $\Gamma' \to \Gamma$ induces a bijection $\mathcal{U}_\Gamma \to \mathcal{U}_{\Gamma'}$.*
- *L is determined by L' and the conductance functions of the spikes, and L' is determined by L and the conductance functions of the spikes.*

The same holds if we replace "contracting boundary spikes" by "deleting boundary edges."

Proof. Let's consider the case of contracting one boundary spikes (the proof for multiple boundary spikes is the same but with more complicated notation). Let e be the oriented boundary spike, $\rho_e = \gamma_e^{-1}$ the resistance function. We want to show that any harmonic potential u' on Γ' extends to a unique harmonic potential u on Γ; we only have to choose the potential on $\iota(e)$ since all the other vertices are in Γ'. Since $\tau(e)$ is boundary in Γ' but interior in Γ, there is only one possible choice for c_e that would yield net current zero on $\tau(e)$ for the function on Γ. We then set $u_{\iota(e)} = u_{\tau(e)} + \rho_e(c_e)$.

Note that the boundary data of u is uniquely determined by ρ_e and the boundary data of u'. Indeed, the net current of u on $\iota(e)$ equals the net current of u' on $\tau(e)$ equals c_e, and the $u_{\iota(e)} = u_{\tau(e)} + \rho_e(c_e)$. Also, $B(\Gamma') \setminus \{\tau(e)\} = B(\Gamma) \setminus \{\iota(e)\}$, and the potential / net current on these vertices is the same for u as it is for u'. Similarly, the boundary of u' is uniquely determined by ρ_e and the boundary data of u. Hence, we can find L from L' and vice versa.

For boundary edges, we can make a similar argument: Any harmonic potential on $V(\Gamma)$ is harmonic on $V(\Gamma')$ as well. To find the boundary data of u from u' or u' from u, we keep the potentials the same, and adjust the net currents on the boundary vertices according to the boundary potentials together with conductance functions γ_e of the boundary edges removed. $\square$

Lemma 3.5. *Suppose that G' is obtained from G by deleting some disconnected boundary vertices. Let Γ, Γ', L, L' be as above. Then*

- *The inclusion $\Gamma' \to \Gamma$ induces a surjection $\mathcal{U}_\Gamma \to \mathcal{U}_{\Gamma'}$.*
- *L' is uniquely determined by L and L is uniquely determined by L'.*

Proof. Easy exercise. $\square$

Lemma 3.6. *Suppose G' is obtained from G by reduction operation. Let Γ, Γ', L, L' be as above. Then*

- *The inclusion $\Gamma' \to \Gamma$ induces a surjection $\mathcal{U}_\Gamma \to \mathcal{U}_{\Gamma'}$.*
- *L' is uniquely determined by L and L is uniquely determined by L'.*

Proof. Any reduction operation can be expressed in three steps as a contraction of non-degenerate spikes, deletion of boundary edges, and deletion of disconnected boundary vertices. $\square$

3.6 Solvable and Totally Layerable Bgraphs

Let $G_0, G_1, \ldots, G_n$ be a layerable filtration of a bgraph G. We say that it is a *solvable filtration* if it satisfies the following:

- For each spike e removed from G_n, there is a scaffold on G_n in which e is a horizontal, middle edge.
- For each boundary edge e removed from G_n, there is a scaffold on G_n in which e is a vertical, middle edge.

A bgraph which admits a solvable filtration is called *solvable*. This name is appropriate because these are precisely the graphs for which the inverse problem can be solved through layer-stripping with repeated application of information propagation:

Theorem 3.7. *Any solvable bgraph is recoverable over BZCF.*

Proof. Let Γ be a BZCF network on G. Let $G_0, G_1, \ldots$ be a solvable filtration, and let $\Gamma_0, \Gamma_1, \ldots$ be the corresponding subnetworks with sets of boundary data $L_0, L_1, \ldots$. By Lemma 3.2, the conductance functions of the edges removed from G_n are uniquely determined by L_n over BZCF. By Lemma 3.6, L_{n+1} is determined by L_n and these conductance functions. Hence, by induction each conductance function and each L_n is uniquely determined by L over BZCF. $\square$

Proposition 3.8. *Let $f : G \to G'$ be a locally injective bgraph morphism. If $G'_0, G'_1, \ldots$ is a solvable filtration of G', then $f^{-1}(G'_0), f^{-1}(G'_1), \ldots$ is a solvable filtration of G. Hence, if G' is solvable, then so is G.*

Proof. We already know that a layerable filtration pulls back to a layerable filtration. To see that $f^{-1}(G'_0), f^{-1}(G'_1), \ldots$ is a solvable filtration, we just pull back the scaffolds used for each edge, using Proposition 3.1. □

A more symmetrical (and it turns out stronger) condition than solvability is total layerability. We say that a bgraph G is *totally layerable* if for any edge e, there exists a scaffold $\mathcal{S}$ with $e \in \mathrm{Hor}\,\mathcal{S} \cap \mathrm{Mid}\,\mathcal{S}$ and a scaffold $\mathcal{S}'$ with $e \in \mathrm{Vert}\,\mathcal{S}' \cap \mathrm{Hor}\,\mathcal{S}'$.

Proposition 3.9.

1. *If $f : G \to G'$ is a locally injective bgraph morphism and G' is totally layerable, then so is G.*
2. *Any totally layerable bgraph is layerable.*
3. *If G is a totally layerable, then it is solvable. In fact, any layerable filtration of G is a solvable filtration.*

Proof. The first claim follows immediately from Proposition 3.1. (2) follows from Lemma 3.3. To prove (3), let $G_0, G_1, \ldots$ be a layerable filtration. If e is a boundary spike / boundary edge of G_{n-1}, then there exists a scaffold on G in which e is a horizontal / vertical middle edge, and this induces a scaffold on G_n as well. □