Layering Graphs-with-Boundary and Networks

David Jekel

July 20, 2015

Abstract

We analyze geometric structure of graphs-with-boundary (or "B-graphs") relevant to the study of electrical networks, especially the electrical inverse problem. As in [5] and [9], networks are generalized to include signed, nonlinear, and infinite resistor networks. We review fundamental results about linear networks, including the Star-K transformation and grove-determinant formula, as the motivation and groundwork for our results, and sketch how they generalize to arbitrary fields. Following [1], we describe gluing together networks as composition in a category; this formalism sheds light on the relationship between ranks and connections and its generalization to the nonlinear case. A related geometric structure called a scaffold allows us describe harmonic continuation and formalize the layer-stripping approach to the inverse problem in a way which includes infinite networks. A class of "solvable B-graphs" is defined by purely geometric conditions which guarantee the inverse problem has at most one solution for a general class of nonlinear resistors. Critical circular planar graphs are solvable, and so are any subgraphs, covering graphs, and "graph convolutions" of solvable B-graphs.

Acknowledgements

This paper owes its existence to James Morrow. He organized the University of Washington's math REU, which I participated in during summer 2013, and he encouraged me to continue to think about electrical networks after the program was finished. During the program, both he and Ian Zemke spent a lot of time listening to and critiquing my ideas. More recently, Avi Levy has been a continual source of new questions and ideas, and introduced me to the work of John Baez and Brandan Fong.

Above all, the tradition of the UW's REU, its library of techniques and examples invented by undergraduates, graduates, and professors together and preserved by Jim Morrow, has been the catalyst for my creative processes. In particular, this paper continues the work of Curtis and Morrow, Will Johnson, Konrad Schrøder, Ian Zemke, and others. The REU student papers are generally not polished, and it is hard to say for sure who came up with what ideas, but I will do my best to give credit where credit is due. Several times during my research, I discovered that some of my ideas had already been invented, in a different and more sophisticated form, by other mathematicians studying electrical networks (Lam and Pylyavsky, Baez and Fong). This type of thing has happened many times in network theory: Variants of the matrix-tree theorem were rediscovered many times. Besides, Curtis/Morrow and de Verdiere solved the inverse problem for circular planar networks simultaneously and independently. If I have copied anyone else's results, be assured it is unintentional, and I will insert proper citations when I become aware of the need.

Soli Deo Gloria.

Introduction

Network theory touches on electrical engineering, graph theory, combinatorics, and symplectic Lie theory. Much of it is motivated by analogy with PDE and complex analysis. TO BE CONTINUED

Contents

1	Graphs with Boundary								
	1.1 The Category of Graphs	4							
	1.2 The Category of B-graphs	4							
	1.3 Basic Constructions in the Category of B-graphs	5							
	1.4 Other Graph Terminology	6							
2 Networks									
	2.1 Harmonic Functions	8							
	2.2 Boundary Behavior; Forward and Inverse Problems	9							
3	Linear Networks	9							
	3.1 The Kirchhoff Matrix	10							
	3.2 A Grove-Determinant Formula	11							
	3.3 Singular Networks over \mathbb{R}	16							
	3.4 Ranks and Connections	19							
	3.5 Properties of L for Networks over a Field \ldots	19							
	3.6 Local Network Equivalences	20							
4	Layering I: Reduction Operations								
	4.1 Motivation: Layer-stripping for the Inverse Problem	24							
	4.2 Definitions and Basic Properties	25							
5	Layering II: IO-Graphs and Factorization	27							
	5.1 The Category of Input-Output Graphs	27							
	5.2 Elementary IO-graphs	29							
	5.3 Factorization, Layerability, and Regularity	32							
	5.4 Factorization, Rank, and Connections	34							
6	Layering III: Scaffolds	37							
	6.1 Motivation \ldots	37							
	6.2 Definitions	38							
	6.3 Scaffolds, Elementary Factorization, and Layerability	39							
	6.4 Scaffolds, Harmonic Extensions, and Recovery	41							
	6.5 Solvable and Totally Layerable B-graphs	44							
7	Graphs on Surfaces								
	7.1 Embeddings of Graphs, Strand Arrangements, and Medial Strands	45							
	7.2 Producing Scaffolds from the Medial Graph	47							
	7.3 B-graphs on the Disk	48							
	7.4 B-graphs on the Half-Plane	52							
8	A Theorem on Unique Connections	53							

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⁻: $E \to E$ 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 B-graphs

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 "B-graph" for short.

A B-graph 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 B-graph 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 B-graphs

Some special types of morphisms: A B-graph 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-B-graphs: We say G_1 is a sub-B-graph 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-B-graph of G_2 if it is a sub-graph and the inclusion $G_1 \to G_2$ is a B-graph morphism.

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

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

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

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

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

the reader should verify that this is a suB-graph of G_1 .

Pull-backs: The category of B-graphs 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 B-graph 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 B-graph morphisms.

The pull-back satisfies the following the universal property: For any B-graph 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 B-graph morphism $\phi: H \to G_1 \times_G G_2$ such that $g_j = \pi_j \circ \phi$, as depicted here:



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 B-graphs 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 B-graph 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 B-graph 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 B-graph 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 Other Graph Terminology

For a graph G, a *path* is a sequence of vertices p_0, \ldots, p_K and oriented edges e_1, \ldots, 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, \ldots, 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, \ldots, 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, \ldots, 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 suB-graphs G_1, \ldots, G_N , called *components*, such that $V(G_1), \ldots, V(G_N)$ are a partition of V(G), and $E(G_1), \ldots, E(G_N)$ are a partition of E(G), and $B(G_1), \ldots, B(G_N)$ 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 (and will consider in other papers):

- 1. It is standard to consider real-valued linear electrical networks. We are given a B-graph 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" [5], 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" [5], 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 B-graphs 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 B-graph 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 B-graphs. We denote the labelling by Θ .

An *M*-valued potential function on a B-graph *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.

Similarly, 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

 $(u_{\iota(e)} - u_{\tau(e)}, c_e) \in \Theta_e$ 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}_{Γ} 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 B-graph 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 B-graph 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 Behavior; 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 boundary behavior or set of boundary data L as the image of Φ . (The term "behavior" is borrowed from [1].)

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 Linear Networks

In this section, we shall review / prove a hodgepodge of results about linear networks, which serve as tools or motivations for the results of later sections. We will also describe when familiar results for positive real conductances generalize to nonzero conductances in any field.

3.1 The Kirchhoff Matrix

Let G be a finite B-graph and assume it has no self-looping edges. Let Γ be a network on G with Θ_e given by $\{(t, a_e t)\}$ for some nonzero $a_e = a_{\overline{e}}$ in a field \mathbb{F} . Let \mathcal{H}_{Γ} be the set of F-valued harmonic functions (u, c) on Γ , with $c_e = a_e \cdot (u_{\iota(e)} - u_{\tau(e)})$. Note \mathcal{H} is a linear subspace of $F^V \times F^E$. Since c is determined by u, \mathcal{H}_{Γ} is linearly isomorphic to \mathcal{U}_{Γ} , the set of harmonic potentials, so we might as well work with \mathcal{U}_{Γ} .

Recall \mathcal{U}_{Γ} is the set of all $u \in F^V$ such that

$$\sum_{e \in \iota^{-1}(v)} a_e(u_{\iota(e)} - u_{\tau(e)}) = 0 \text{ for each } v \in I.$$

These equations can be compactly expressed in terms of the Kirchhoff matrix. Define the *Kirchhoff matrix* $K \in M_V \mathbb{F}$ by

$$K_{p,q} = \begin{cases} \sum_{\substack{e:\iota(e)=p, a_e, \\ \tau(e)=q}} a_e, & p \neq q \\ -\sum_{e:\iota(e)=v} a_e, & p = q. \end{cases}$$

The Kirchhoff matrix K defines a linear transformation $\mathbb{F}^V \to \mathbb{F}^V$, and the component indexed by p is

$$(Ku)_p = -\sum_{e \in \iota^{-1}(p)} a_e(u_{\iota(e)} - u_{\tau(e)}).$$

If $\pi_I : \mathbb{F}^V \to \mathbb{F}^I$ is the projection map, then u is harmonic if and only if $\pi_I K u = 0$, hence

$$\mathcal{U}_{\Gamma} = \ker(\pi_I K) \subset F^V.$$

The Dirichlet and Neumann problems have an interpretation in terms of linear algebra. In the following, we will assume G is connected. There is no real loss of generality, since a harmonic function on G restricts to a harmonic function on any connected component, and harmonic functions on the connected components combine to form a harmonic function on G. And components with no boundary vertices are of little interest. If G has multiple connected components G_1, \ldots, G_N and we reorder the vertices of G so that the vertices of G_1 are first, then $V(G_2)$, and so on, then the Kirchhoff matrix will decompose into blocks

$$K = \begin{pmatrix} K_1 & 0 & \cdots & 0 \\ 0 & K_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & K_N \end{pmatrix},$$

so the behavior of the whole can easily be understood from the behavior of the smaller blocks.

Consider the Dirichlet problem. For $\phi \in \mathbb{R}^B$, we want to find a harmonic potential u with $u|_B = \phi$. This is the same as letting $u = (\phi, w)$, where w satisfies

$$K_{I,B}\phi + K_{I,I}w = 0.$$

This will have a unique solution if and only if $K_{I,I}$ is invertible. As we will see, this does not always happen. But suppose $K_{I,I}$ is invertible. Then $w = -K_{I,I}^{-1}K_{I,B}\phi$. The current on each edge can be computed from the conductance functions. The net current on the boundary vertices is

$$\psi = K_{B,B}\phi + K_{B,I}w = (K_{B,B} - K_{B,I}K_{I,I}^{-1}K_{I,B})\phi.$$

The matrix $\Lambda = K_{B,B} - K_{B,I}K_{I,I}^{-1}K_{I,B}$ is called the *response matrix* and it acts as a *Dirichlet-to-Neumann map* $\mathbb{F}^B \to \mathbb{F}^B$ sending boundary potentials to the boundary net currents of the corresponding harmonic function. Then $L = \{(\phi, \Lambda \phi) : \phi \in \mathbb{F}^B\}.$

The matrix Λ is an example of a Schur complement. Suppose we have a square matrix M partitioned into blocks

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

and D is square and invertible. Then $A-BD^{-1}C$ is called the Schur complement of D in M and is denoted M/D. In this notation, $\Lambda = K/K_{I,I}$.

The Neumann problem has a similar interpretation. For $\psi \in \mathbb{F}^B$, we want to find a potential u such that

$$Ku = \begin{pmatrix} \psi \\ 0 \end{pmatrix}.$$

Of course, if ψ came from a valid current function, its entries must sum to zero. We cannot expect the solution to the Neumann problem to be unique either. Indeed, if we take a harmonic function and raise the potentials on all the vertices by some constant, then the new function will be harmonic and have the same boundary currents.

So we revise the Neumann problem as follows: Let $A \subset \mathbb{R}^V$ be the set of functions whose entries sum to zero. For $(\psi, 0)^T \in A$, does there exist a unique harmonic (u, c) with $u \in A$ and $Ku = (\psi, 0)^T$? The answer is yes if and only if $K|_A$ is invertible. Since the image of K is contained in A, this happens if and only if rank $K = \dim A$, which is |V| - 1.

3.2 A Grove-Determinant Formula

Our main tool to determine when certain submatrices of K are invertible is the following combinatorial result, which generalizes the matrix-tree theorem attributed to Kirchhoff. The earliest source known to the author for the general version of the formula is Robin Forman [3].

Let G be a graph. A spanning tree T is a subgraph (without boundary) such that T is connected, every vertex is in T, and T has no cycles. A forest F is a subgraph F with no cycles; the components of F have no cycles, and are therefore *trees*. A grove is a forest such that each component contains a boundary vertex.

Let P and Q be disjoint subsets of B with |P| = |Q| = n. Let $\mathcal{F}(P,Q)$ be the set of groves F such that each connected component either contains exactly one vertex from P and one from Q or it contains exactly one vertex from $B \setminus (P \cup Q)$. Let $K_{P \cup I,Q \cup I}$ be the submatrix of K with rows indexed by $P \cup I$ and columns by $Q \cup I$, ordered according to a given indexing of vertices by the integers $1, \ldots, |V|$. Let p_1, \ldots, p_n be the vertices of P and q_1, \ldots, q_n the vertices of Q ordered according to the same indexing. For any $F \in \mathcal{F}(P,Q)$, there is a permutation $\tau \in S_n$ such that p_j and $q_{\tau(j)}$ are in the same component of F; call this permutation τ_F .

Theorem 3.1 (Grove-determinant Formula). Let G be a finite B-graph labelled by elements of a field \mathbb{F} . Let P and Q be disjoint subsets of B with |P| = |Q| = n. Then

$$\det K_{P\cup I,Q\cup I} = (-1)^n \sum_{F \in \mathcal{F}(P,Q)} \operatorname{sgn} \tau_F \prod_{e \in E'(F)} a_e.$$

Proof. Let m = |I|. Let p_1, \ldots, p_{n+m} be the vertices of $P \cup I$ and q_1, \ldots, q_{n+m} be the vertices of $Q \cup I$, so that $P = \{p_1, \ldots, p_n\}$ and $Q = \{q_1, \ldots, q_n\}$ and for $j > n, p_j = q_j \in I$. Suppose $\sigma \in S_{n+m}$; if $p_j = q_{\sigma(j)}$, then p_j must be interior. Let m_{σ} be the number of indices with $p_j = q_{\sigma(j)}$. By definition, det $K_{P \cup I, Q \cup I}$ is

$$\sum_{\sigma \in S_{n+m}} \operatorname{sgn} \sigma \prod_{j=1}^{n+m} \kappa_{p_j, q_{\sigma(j)}}$$

$$= \sum_{\sigma \in S_{n+m}} \operatorname{sgn} \sigma \left(\prod_{\substack{p_j \neq q_{\sigma(j)} \\ \tau(e) = q_{\sigma(j)}}} \sum_{\substack{e:\iota(e) = p_j \\ \tau(e) = q_{\sigma(j)}}} (-a_e) \right) \left(\prod_{\substack{p_j = q_{\sigma(j)} \\ \tau(e) = q_{\sigma(j)}}} \sum_{e:\iota(e) = p_j} a_e \right)$$

$$= \sum_{\sigma \in S_{n+m}} (-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma \left(\prod_{\substack{p_j \neq q_{\sigma(j)} \\ \tau(e) = q_{\sigma(j)}}} \sum_{\substack{e:\iota(e) = p_j \\ \tau(e) = q_{\sigma(j)}}} a_e \right) \left(\prod_{\substack{p_j = q_{\sigma(j)} \\ \tau(e) = q_{\sigma(j)}}} \sum_{e:\iota(e) = p_j} a_e \right)$$

Our goal is to expand each of the sums inside the product. Fix σ ; choosing one term from each of the inner sums amounts to choosing for each j an edge e_j such that (1) $\iota(e_j) = p_j$ and (2) if $p_j \neq q_{\sigma(j)}$, then $\tau(e) = q_{\sigma(j)}$. Let \mathcal{Y} be the collection of all sets $Y = \{e_1, \ldots, e_{n+m}\}$ such that $\iota(e_j) = p_j$. We say $\sigma \in S_{n+m}$ and $Y \in \mathcal{Y}$ are compatible if (1) and (2) are satisfied for every $e_j \in Y$. Then

$$\det K_{P\cup I,Q\cup I} = \sum_{\sigma\in S_{n+m}} (-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma \sum_{\substack{\operatorname{compatible}\\Y\in\mathcal{Y}}} \prod_{e\in Y} a_e$$
$$= \sum_{Y\in\mathcal{Y}} \sum_{\substack{\operatorname{compatible}\\\sigma\in S_{n+m}}} (-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma \prod_{e\in Y} a_e$$

Suppose that Y contains a sequence of edges e_{j_1}, \ldots, e_{j_k} with $\tau(j_\ell) = \iota(j_{\ell+1})$ for $\ell = 1, \ldots, k-1$ and $\tau(e_{j_k}) = \iota(e_{j_1})$. (Either such a sequence forms a cycle or k = 2 and it is a pair e, \overline{e} .) If σ is compatible with Y, there are two possibilities: Either (1) $\sigma(j_\ell) = j_\ell$ for all ℓ or (2) $j_1 \mapsto j_2 \mapsto \ldots \mapsto j_k \mapsto j_1$ is a cycle of σ . In fact, there is a one-to-one correspondence between compatible permutations satisfying (1) and those satisfying (2), and we can partition the compatible permutations into pairs $\{\sigma, \xi\sigma\}$, where $\xi \in S_{n+m}$ is the cycle $j_1 \mapsto j_2 \mapsto \ldots \mapsto j_k \mapsto j_1$, such that σ satisfies (1) and $\xi\sigma$ satisfies (2). Then $m_{\xi\sigma} = m_{\sigma} - k$ and $\operatorname{sgn} \xi = (-1)^{k+1}$, so

$$(-1)^{n+m-m_{\xi\sigma}}\operatorname{sgn}(\xi\sigma) = (-1)^{n+m-m_{\sigma}-k}(-1)^{k+1}\operatorname{sgn}\sigma = -(-1)^{n+m-m_{\sigma}}\operatorname{sgn}\sigma.$$

Thus,

$$\sum_{\substack{\text{ompatible}\\\sigma\in S_{n+m}}} (-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma = 0$$

because the terms for σ and $\xi \sigma$ cancel.

с

Therefore, it suffices to consider elements $Y \in \mathcal{Y}$ which do not contain cycles or pairs $\{e, \overline{e}\}$. For any such Y, there is a unique spanning forest F with $E(F) = Y \cup \overline{Y}$. I claim that

- 1. If Y is compatible with σ , then the corresponding F is in $\mathcal{F}(P,Q)$,
- 2. There is a one-to-one correspondence between compatible (Y, σ) pairs and forests F, and
- 3. For each (Y, σ) , we have $(-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma = (-1)^n \operatorname{sgn} \tau_F$.

To prove (1), it suffices to show that every component of F includes exactly one vertex from $B \setminus P$, that is, one vertex from Q or one from $B \setminus \{P \cup Q\}$. For each p_j , there is a unique outgoing $e_j \in Y$ with $\iota(e_j) = p_j$. We start at p_j and construct a path following the oriented edges of Y. As long as the last vertex is in $P \cup I$, we can continue the path. Since Y has no cycles or conjugate pairs, we cannot repeat vertices, so eventually we will reach a vertex in $B \setminus P$, so every component has one vertex from $B \setminus P$. Suppose for the sake of contradiction that it had more than one. Then there would be $r, r' \in B \setminus P$ and a path from r to r' using oriented edges $\epsilon_1, \ldots, \epsilon_K \in Y \cup \overline{Y}$. We can assume without loss of generality that r and r' are the only vertices in $B \setminus P$ in the path. If $e \in Y$, then $\iota(e) \in P \cup Q$. Thus, $\epsilon_1 \notin Y$, $\epsilon_K \in Y$. Let k be the first index such that $\epsilon_k \in Y$. Then $\epsilon_{k-1} \in \overline{Y}$, so ϵ_k and $\overline{\epsilon_{k-1}}$ are two edges in Y with the same initial vertex, which contradicts our definition of \mathcal{Y} .

(2) Choose F, and we will show there is a unique (Y, σ) which corresponds to F. We obtain Y from E(F) by choosing one orientation for each edge. For an $e \in E(F)$, there is an embedded path from $\iota(e)$ to some $r \in B \setminus P$; this embedded path must be unique because F is a forest. There is also an embedded path from $\tau(e)$ to r, and one of the two paths must use e or \overline{e} . We choose the orientation which matches the orientation of the path. These orientations are uniquely determined: If we assume $e \in Y$ for some Y but that the orientation of e does not match the orientation of the path, then we reach a contradiction by the same argument as above.

To construct σ , we decompose τ_F into disjoint cycles η_1, \ldots, η_K . For each η_k , we define a cycle $\sigma_k \in S_{n+m}$ as follows: Let η_k be given by $i_1 \mapsto i_2 \mapsto i_R \mapsto i_1$ (the dependence on k has been suppressed in the notation). There is a unique embedded path in F from p_{i_r} to $q_{i_{r+1}}$ and the other vertices in the path are interior, so the vertices in all the paths have the form

We define the cycle ξ_k by

$$i_{1} \mapsto j_{1,1} \mapsto j_{1,1} \mapsto j_{1,2} \mapsto \dots \mapsto j_{1,n_{1}} \mapsto i_{2}$$
$$i_{2} \mapsto j_{2,1} \mapsto j_{2,1} \mapsto j_{2,2} \mapsto \dots \mapsto j_{2,n_{2}} \mapsto i_{3}$$
$$\dots$$
$$i_{R} \mapsto j_{R,1} \mapsto j_{R,1} \mapsto j_{R,2} \mapsto \dots \mapsto j_{R,n_{R}} \mapsto i_{1}.$$

Then let $\sigma = \xi_1 \xi_2 \dots \xi_n$. To show σ is uniquely determined, suppose σ is product of cycles ξ_1, \dots, x_k . Suppose ξ_k given by $j_1 \mapsto j_2 \mapsto \dots \mapsto j_L \mapsto j_1$. If each j_ℓ was greater than n (corresponding to an interior vertex), then we would have $p_{j_\ell} = q_{j_\ell} \in I$ and the edges $e_{j_1}, \dots, e_{j_L} \in Y$ would form a cycle or pair $\{e, \overline{e}\}$, contradicting our assumptions about Y. Thus, some of the indices in the cycle are $\leq n$; it follows that ξ_k must represent boundary-to-boundary paths just as in our original construction, and the paths are uniquely determined by F

(3) Consider a cycle η_k which maps $i_1 \mapsto i_2 \mapsto i_R \mapsto i_1$, and let $j_{i,1}, \ldots, j_{i,n_r}$ be as above. Let $z_k = \sum_{r=1}^R n_r$, which is the number of interior vertices in the paths corresponding to η_k . Then $\operatorname{sgn} \xi_k = (-1)^{z_k} \operatorname{sgn} \eta_k$. The total number of interior vertices in the paths is $\sum_{k=1}^K z_k$. The interior vertices not in the paths are exactly the vertices p_j for which $\sigma(j) = j$. Hence, $\sum_{k=1}^K z_k = m - m_{\sigma}$. Therefore,

$$\operatorname{sgn} \sigma = \operatorname{sgn}(\xi_1 \dots \xi_n) = (-1)^{\sum_k z_k} \operatorname{sgn}(\eta_1 \dots \eta_n) = (-1)^{m-m_\sigma} \operatorname{sgn} \tau_F.$$

Thus, $(-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma = (-1)^n \operatorname{sgn} \tau_F$. Therefore,

$$\sum_{Y \in \mathcal{Y}} \sum_{\substack{\sigma \in S_{n+m} \\ \sigma \in S_{n+m}}} (-1)^{n+m-m_{\sigma}} \operatorname{sgn} \sigma \prod_{e \in Y} a_e = (-1)^n \sum_{F \in \mathcal{F}(P,Q)} \operatorname{sgn} \tau_F \prod_{e \in E'(F)} a_e.$$

Corollary 3.2. Let $\mathcal{F} = \mathcal{F}(\emptyset, \emptyset)$. Then det $K_{I,I} = \sum_{F \in \mathcal{F}} \prod_{e \in E'(F)} a_e$.

Proof. The proof is the same except that n = 0 and there is no τ_F .

The matrix-tree theorem follows as a special case:

Corollary 3.3 (Matrix-Tree Theorem). Let G be a connected graph (without boundary). Let K be the Kirchhoff matrix of the electrical network where each edge has conductance $a_e = 1$. For $p, q \in V$, $(-1)^{p-q} \det K_{V \setminus \{p\}, V \setminus \{q\}}$ is the number of spanning trees of G.

Proof. If p = q, then make G into a graph with boundary by setting $B = \{p\}$. Reindex the vertices so that p occurs first; this does not change the determinant. Then by the previous theorem,

$$\det K_{V\setminus\{p\},V\setminus\{p\}} = \det K_{I,I} = \sum_{F\in\mathcal{F}} \operatorname{sgn} \tau_F.$$

Since p is the only boundary vertex, each forest is a spanning tree, so the result is the number of spanning trees. If $p \neq q$, set $B = \{p, q\}$. Reindex the vertices so that p and q occur first; this does not change the determinant, but it does change $(-1)^{p-q}$ to -1. Compute

$$\det K_{V\setminus\{p\},V\setminus\{q\}} = \det K_{I\cup\{q\},I\cup\{p\}} = -\sum_{F\in\mathcal{F}(\{q\},\{p\})} \operatorname{sgn} \tau_F.$$

Again, since p and q are the only boundary vertices, each spanning forest is a spanning tree, and τ_F is the identity.

The grove-determinant formula allows us to test when $K_{P\cup I,Q\cup I}$ is invertible for networks over various fields. In particular:

Proposition 3.4. Let G be a finite B-graph labelled with nonzero elements of some field F. Let $P, Q \subset B$ be disjoint with |P| = |Q| = n. Then

- a. If $\mathcal{F}(P,Q) = \emptyset$, then det $K_{P \cup I,Q \cup I} = 0$.
- b. If $\mathcal{F}(P,Q)$ has exactly one element, then det $K_{P\cup I,Q\cup I} \neq 0$.
- c. If $\mathcal{F}(P,Q)$ has more than one element, then we can label the edges of G with nonzero real numbers which will make det $K_{P\cup I,Q\cup I}$ positive, negative, or zero. The determinant is nonzero for some positive numbers.

Proof. In case (a), the determinant-forest formula expresses the term is a sum over an empty index set, which is zero. In case (b), there is exactly one term in the sum, which is a product of nonzero numbers, hence nonzero. Now consider case (c). Let F_1 and F_2 be two distinct forests. All the forests must have the same number of edges, as is clear from the proof of the determinant-forest formula. Thus, there is some $e_0 \in E'(F_1) \setminus E'(F_2)$ and $e_1 \in E'(F_2) \setminus E'(F_1)$. Choose a sign sgn $e = \pm 1$ for each edge in E as follows: Set sgn e = 1 for

 $e \neq e_0, e_1$ and choose sgn e_0 such that sgn $e_0 \operatorname{sgn} \tau_{F_1} = 1$ and sgn $e_1 \operatorname{sgn} \tau_{F_2} = -1$. Choose $\epsilon < 1/|\mathcal{F}(P,Q)|$ and set

$$a_e = \begin{cases} \operatorname{sgn} e, & e \in E'(F_1) \\ \epsilon \operatorname{sgn} e, & \notin E'(F_1). \end{cases}$$

Then in the forest expansion of det $K_{P\cup I,Q\cup I}$, the term for F_1 dominates making the determinant positive. In the other hand, if

$$b_e = \begin{cases} \operatorname{sgn} e, & e \in E'(F_2) \\ \epsilon \operatorname{sgn} e, & \notin E'(F_2), \end{cases}$$

then the determinant is negative. Applying the intermediate value theorem to the connected region $\{c \in \mathbb{R}^{E'} : \operatorname{sgn} c_e = \operatorname{sgn} e\}$ shows that there are nonzero numbers c_e which will make the determinant zero.

The same argument shows that whatever sign we choose for the edges, we can make det $K_{P \cup I, Q \cup I}$ nonzero; in particular, this holds if we want the conductances to be positive.

3.3 Singular Networks over \mathbb{R}

A network for which the Dirichlet problem does not have a unique solution is called *Dirichlet-singular*; if the Neumann problem does not have a unique solution up to a constant, it is *Neumann-singular*. Using the spanning forest formula, we will show that the Dirichlet and Neumann problems have a unique solution for reasonable graphs when $a_e > 0$, but if we allow a_e to be positive or negative, one can generally find values of a_e which create a Dirichlet-singular or Neumann-singular network. We assume throughout that G is connected and has some boundary vertices.

As the reader can verify, this implies that there is at least one spanning forest in \mathcal{F} . Hence, if $a_e > 0$,

$$\det K_{I,I} = \sum_{F \in \mathcal{F}} \prod_{e \in E'(F)} a_e > 0.$$

Therefore, the Dirichlet problem has a unique solution. However, for most graphs there will be real values of $a_e \neq 0$ for which $K_{I,I}$ is not invertible. Indeed, if there exist some interior vertices, and if each interior vertex has degree ≥ 2 , then there are at least two forests in \mathcal{F} , hence the determinant is sometimes zero for nonzero real numbers by Proposition 3.4.

A more delicate question is, what are the possible values of dim ker $K_{I,I}$? This depends on the graph, but in some cases, it is easy to find a lower bound: Suppose G_1, \ldots, G_N form a suB-graph partition of G and $B(G_k) \subset B(G)$ for all k. Suppose there are Dirichlet-singular conductances for each G_k , and let the conductances on G be the same as the conductances on the G_k 's. Since ker $K_{I,I}$ is nontrivial for each G_k , there is a nonzero harmonic potential u_k on G_k , and we can extend it to G by setting it to zero on the other vertices. The potentials thus defined are linearly independent because u_k is nonzero on G_k , but u_j for $j \neq k$ is zero on G_k . Thus, dim ker $K_{I,I} \geq N$.

If $a_e > 0$, the Neumann problem has a unique solution. By similar reasoning as in Corollary 3.3, for any p, q,

$$(-1)^{p-q} \det K_{V \setminus \{p\}, V \setminus \{q\}} = \sum_{\substack{\text{spanning } e \in E'(T) \\ \text{trees } T}} \prod_{de} a_e.$$

Since G is connected, it has a spanning tree, so the right hand side is positive if $a_e > 0$. So K has rank |V| - 1 and the Neumann problem has a unique solution. This also shows that the determinant of any |V| - 1 by |V| - 1 submatrix of K is the same up to sign, so to see whether the Neumann problem has a unique solution, it suffices to check one of them.

If G is a tree (that is, it has no cycles), then there is only one spanning tree of G, which is all of G, so the Neumann problem has a unique solution. However, if G has a cycle, there is more than one spanning tree, so by Proposition 3.4, there exist signed a_e 's which produce a Neumann-singular network.

What are the possible values of dim ker K? It must be ≥ 1 . Now suppose G_1, \ldots, G_N form a subgraph partition of G, such that each G_k is connected and any cycle of G is contained in some G_k . Suppose there exist Neumann-singular conductances on each G_k , and use them to define conductances on G. Then for each G_k , there exists a non-constant harmonic potential u_k on G_k with net current zero on every vertex. We can extend u_k to G by defining it to be constant on each G_k ; this will be consistent because every cycle is contained in some G_k . Then the u_k 's are linearly independent, so dim ker $K \geq N + 1$.

For some networks, it is possible for a nonzero harmonic function to have potential and current zero on the boundary, even if there are no components without boundary vertices. Consider the "triangle-in-triangle" network with boundary vertices $\{1, \ldots, 6\}$ and interior vertices $\{7, 8, 9\}$ and edges with coefficients a_e shown in the figure. The Kirchhoff matrix is

(0	0	0	0	0	0	$^{-1}$	0	$1 \rangle$
0	0	0	0	0	0	1	-1	0
0	0	0	0	0	0	0	1	-1
0	0	0	0	0	0	-1	0	1
0	0	0	0	0	0	1	-1	0
0	0	0	0	0	0	0	1	-1
-1	1	0	-1	1	0	0	0	0
0	-1	1	0	$^{-1}$	1	0	0	0
1	0	$^{-1}$	1	0	-1	0	0	0 /

Let χ_p be the vector with 1 on vertex p and zero elsewhere. Then $\chi_7 + \chi_8 + \chi_9$ is a harmonic potential which is zero on the boundary and the corresponding current function has net current zero on the boundary.

Given this fact, it is conceivable that on a network, the Dirichlet problem does not have a unique solution, and yet the net currents on the boundary



Figure 1: Singular conductances on the triangle-in-triangle network. Boundary vertices are colored in.

vertices are uniquely determined since the only harmonic functions with zero potential also have zero current. In this case, we will say that the boundary behavior is not Dirichlet-singular even though the network is. The same consideration applies to the Neumann problem.

3.4 Ranks and Connections

An important consequence of the grove-determinant formula is the relationship between ranks and connections noted by [2] and others.

Let P and Q be sets of boundary vertices. A connection from P to Q is a collection of disjoint boundary-to-boundary paths through the graph such that each path has its intial vertex in P and its terminal vertex in Q. There may be a vertex $p \in P \cap Q$; in this case, any connection from P to Q must include the length-0 path from p to itself. Thus, there is a one-to-one correspondence between connections from P to Q and connections from $P \setminus Q$ to $Q \setminus P$. If there is a connection from P to Q, then P and Q must have the same cardinality.

Assume $K_{I,I}$ is invertible. Suppose P and Q are disjoint subsets of B with |P| = |Q|. Then the submatrix $\Lambda_{P,Q}$ is equal to the Schur complement $K_{P\cup I,Q\cup I}/K_{I,I}$ by elementary computation. The Hainsworth identity for Schur complements tells us that

$$\det \Lambda_{P,Q} = \det K_{P\cup I,Q\cup I} / \det K_{I,I},$$

and hence $\Lambda_{P,Q}$ is invertible if and only if $K_{P\cup I,Q\cup I}$ is invertible.

If there exists a connection from P to Q, then edges in the paths can be completed to a grove in $\mathcal{F}(P,Q)$, and conversely, any grove in $\mathcal{F}(P,Q)$ contains a connection. If there is no connection from P to Q, then Proposition 3.4 tells us that det $\Lambda_{P,Q} = 0$. If there is a connection, then we can choose positive numbers such that $K_{P\cup I,Q\cup I}$ is invertible (and we already know $K_{I,I}$ is invertible for positive conductances), and hence Λ is defined and det $\Lambda_{P,Q} \neq 0$.

Now suppose P and Q do not necessarily have the same cardinality, but are still disjoint sets of boundary vertices. Let m(P,Q) be the maximum size of a connection from some *subset* of P to some *subset* of Q. Then by considering all subsets of P and Q we see that rank $\Lambda_{P,Q} \leq m(P,Q)$ and equality is achieved for some positive conductances. Later, we will describe situations in which equality is guaranteed to hold for all conductances.

3.5 Properties of *L* for Networks over a Field

For linear conductances in a field \mathbb{F} , the space of harmonic potentials is the kernel of $K_{I,V}$, which has dimension at least |V| - |I| = |B|. If (u, c) is harmonic, then the boundary potentials and currents are given by $u|_B$ and $(Ku)|_B$. Let $\Phi : \ker K_{I,V} \to \mathbb{F}^B \times \mathbb{F}^B : u \mapsto (u|_B, (Ku)|_B)$. Then $L = \Phi(\ker K_{I,V})$. Hence, dim $L \leq \dim \ker K_{I,V}$. If there is a harmonic function with zero potential and current on the boundary, as in the last example, then ker Φ is nontrivial, so this inequality is strict.

In general, we would expect \mathcal{H} and L to have dimension |B|; this is the case if either the Dirichlet problem or the Neumann problem has a unique solution. Sometimes dim $\mathcal{H} > |B|$; however, in all cases,

Proposition 3.5. dim L = |B|.

Proof. The kernel of Φ consists of harmonic potentials which are zero on the boundary have zero current on the boundary, that is, ker Φ consists of elements of ker K whose boundary entries are zero. Hence, ker Φ is isomorphic to ker $K_{V,I}$. By the rank-nullity theorem and symmetry of K,

$$\operatorname{rank} \Phi + \dim \ker \Phi = \dim \ker K_{I,V}$$
$$= |V| - \operatorname{rank} K_{I,V}$$
$$= |V| - \operatorname{rank} K_{V,I}$$
$$= |V| - |I| + \dim \ker K_{V,I}$$
$$= |B| + \dim \ker \Phi.$$

Thus, dim $L = \operatorname{rank} \Phi = |B|$.

If the Dirichlet problem has a unique solution, then the Dirichlet-to-Neumann map $\Lambda = K_{B,B} - K_{B,I}K_{I,I}^{-1}K_{I,B}$ is symmetric. So if (ϕ_1, ψ_1) and (ϕ_2, ψ_2) are the boundary data of harmonic functions, then

$$\phi_1 \cdot \psi_2 = \phi_1^T \Lambda \phi_2 = \phi_2^T \Lambda \phi_1 = \phi_2 \cdot \psi_1.$$

Actually, this holds even for Dirichlet-singular networks:

Proposition 3.6. $\phi_1 \cdot \psi_2 = \phi_2 \cdot \psi_1$ for any $(\phi_1, \psi_1), (\phi_2, \psi_2) \in L$.

Proof. Suppose (ϕ_1, ψ_1) and (ϕ_2, ψ_2) are in L, and let u_1 and u_2 be the corresponding harmonic potentials. Let $w_1 = u_1|_I$ and $w_2 = u_2|_I$. Then $\psi_j = K_{B,B}\phi_j + K_{I,B}w_j$. Since $u_j \in \ker K_{I,V}$, we have $0 = K_{I,V}u_j = K_{I,B}\phi_j + K_{I,I}w_j$, which implies $K_{I,B}\phi_j = -K_{I,I}w_j$. Hence, applying the symmetry of K,

$$\begin{split} \phi_1 \cdot \psi_2 &= \phi_1^T (\psi_2 = \phi_1^T (K_{BB} \phi_2 + K_{BI} w_2) \\ &= \phi_1^T K_{B,B} \phi_2 + (K_{I,B} \phi_1)^T w_2 \\ &= \phi_1^T K_{B,B} \phi_2 - (K_{I,I} w_1)^T w_2 \\ &= \phi_1^T K_{B,B} \phi_2 - w_1^T K_{I,I} w_2 \\ &= \phi_2^T K_{B,B} \phi_1 - w_2^T K_{I,I} w_1 \\ &= \phi_2 \cdot \psi_1. \end{split}$$

3.6 Local Network Equivalences

As before, assume we are dealing with linear networks over a field; we leave it to the reader to work out when they generalize to rings. A *series* is the following configuration:



If $a + b \neq 0$, then it is electrically equivalent to



In other words, a series can be reduced to a single edge, and the resistances add: The original resistances were 1/a and 1/b, and the new resistance is 1/a + 1/b. This shows that the series is not recoverable; in fact, there is a one-parameter family of conductances on the series graph which produce the same boundary behavior.

If a + b = 0, then the series is Dirichlet-singular. The two boundary vertices must have the same potential. The potential of the interior vertex is independent of the boundary potentials, but depends on the current flowing from one boundary vertex to the other. In this case, changing the conductances to *ca* and *cb* for some $c \neq 0$ will produce an electrically equivalent network.

Any network which has a series as a subnetwork is not recoverable over the signed linear conductances. If $a+b \neq 0$, we can produce an electrically equivalent network by replacing the series subnetwork with a single-edge subnetwork, as follows from Corollary [refsubnetworksplicing]. This transformation is called a *series reduction* and we call it one type of *local electrical equivalence*. We also call the inverse operation is also a local electrical equivalence.

Suppose a + b = 0 and p and q are the endpoints of the series, and r is the middle vertex. If the series is a subnetwork of a larger network in which p is an interior vertex, then we can produce an electrically equivalent network by "collapsing" the series-identifying p and q and removing r and the edges in the series. This is because any harmonic function must have the same potential on p and q, and the amount of current flowing from p to q is independent of the potentials. This is another type of local electrical equivalence.

A parallel circuit is the following configuration:



If $a + b \neq 0$, then this is equivalent to a single edge with conductance a + b. If a + b = 0, then it is equivalent to a network with no edges. Substituting a parallel edge for a single edge or no edge is another local electrical equivalence.

A Y (left) and a Δ (right) are the following types of networks:



For any Y with $a + b + c \neq 0$, there is a unique equivalent Δ with

$$A = \frac{bc}{a+b+c}, \quad B = \frac{ac}{a+b+c}, \quad C = \frac{ab}{a+b+c}.$$

This can be proved by computing the response matrix Λ for each network. If a + b + c = 0, then in the Y the Dirichlet problem does not always have a solution; however, this is impossible in a Δ , so there is no equivalent Δ . For any Δ with $1/A + 1/B + 1/C \neq 0$, there is a unique equivalent Y with

$$a = \frac{AB + BC + CA}{A}, \quad b = \frac{AB + BC + CA}{B}, \quad c = \frac{AB + BC + CA}{C}.$$

However, if 1/A + 1/B + 1/C = 0, then the Δ is Neumann-singular because it is a tree, so there is no equivalent Y. A Y- Δ transformation is the transformation that replaces a Y subnetwork with an equivalent Δ subnetwork or vice versa.

Over \mathbb{R} , Y- Δ transformations preserve recoverability over the positive linear conductances. For suppose G' is obtained from G by a Y- Δ transformation and G' is recoverable over the positive linear conductances. For any positive linear conductances on G, we can find equivalent conductances on G'. These conductances are uniquely determined by L over the positive linear conductances. In particular, the conductances on the Y or Δ in G' are determined, but then we can find the conductances on the corresponding Δ or Y in G, so G is also recoverable.

We say two graphs are Y- Δ equivalent if there is a sequence of Y- Δ transformations which will change one into the other. This is an equivalence relation. If G is Y- Δ equivalent to G' and G' has a series or parallel configuration, then G' is not recoverable, and hence G is not recoverable over the positive linear conductances. This is one of the best methods for showing a graph is not recoverable over \mathbb{R} , and it is applied in [2] to circular planar networks.

The final type of local electrical equivalence is the \bigstar - \mathcal{K} transformation described in [7] and [4]. An *n*-star is a graph with *n* boundary vertices and one interior vertex, and one edge from the interior vertex to each boundary vertex. The complete graph \mathcal{K}_n is a graph with *n* boundary vertices and one edge between each pair of distinct boundary vertices. For example, here are networks on 4-star and \mathcal{K}_4 graphs:



Index the vertices of the *n*-star and \mathcal{K}_n by $1, \ldots, n$. Let a_j be the conductance of the star edge incident to j and $b_{i,j}$ the conductance of the edge in the \mathcal{K}_n between vertices i and j. Let $\sigma = a_1 + \cdots + a_n$. For any star with $\sigma \neq 0$, there is an equivalent \mathcal{K}_n with conductances $b_{i,j} = a_i a_j / \sigma$. If $\sigma = 0$, then the star has Dirichlet-singular boundary behavior and hence is not equivalent to a \mathcal{K}_n . If $n \geq 4$, most \mathcal{K}_n 's are not equivalent to a star, unlike the n = 3 case of Y- Δ transformations:

Lemma 3.7. Let $n \ge 4$. A network on a \mathcal{K}_n is equivalent to a star if and only if

- It satisfies the quadrilateral rule: $b_{i,j}b_{k,\ell} = b_{i,k}b_{j,\ell}$ for distinct i, j, k, ℓ .
- It is not Neumann-singular.

Proof. If the network is equivalent to a star, then for distinct i, j, k, ℓ ,

$$b_{i,j}b_{k,\ell} = \frac{a_i a_j a_k a_\ell}{\sigma^2} = b_{i,k}b_{j,\ell}$$

A star is a tree and is therefore not Neumann-singular. Thus, the boundary behavior of the \mathcal{K}_n cannot be Neumann-singular, and since all its vertices are boundary vertices this is equivalent to the network itself not being Neumann-singular.

Suppose conversely that a \mathcal{K}_n network satisfies these two conditions. Fix *i* and choose distinct $k, \ell \neq i$, and let

$$a_i = \sum_{j \neq i} b_{i,j} + \frac{b_{i,k}b_{i,\ell}}{b_{k,\ell}}$$

The quadrilateral rule guarantees that the right hand side is independent of k and ℓ . But for a fixed k and ℓ , this is the current on vertex i of the potential $\chi_i - (b_{i,\ell}/b_{k,\ell})\chi_k$ on the \mathcal{K}_n network. This function has net current 0 on vertex ℓ , but since $b_{i,\ell}/b_{k,\ell}$ is independent of the choice of ℓ , it has current 0 on all vertices other than k and i. Since the potential is not constant and the network

is not Neumann-singular, there must be nonzero net current on i and k, so a_i must be nonzero.

Now we must verify that $\sigma = \sum a_i \neq 0$ and that $a_i a_j / \sigma = b_{ij}$. By extending F to a larger field if necessary, we can assume that there exists c_i with

$$c_i^2 = b_{i,k} b_{i,\ell} / b_{k,\ell}$$
 for distinct $k, \ell \neq i$,

and again this is independent of k, ℓ . Then

$$c_i^2 c_j^2 = \frac{b_{i,k} b_{i,j}}{b_{j,k}} \frac{b_{j,k} b_{i,j}}{b_{i,k}} = b_{i,j}^2$$

so that $c_i c_j = \pm b_{i,j}$. By choosing c_1 first and then modifying c_j for $j \neq 1$ if necessary, we can guarantee $c_1 c_j = b_{1,j}$ for $j \neq 1$. Then for $i \neq 1$ we have

$$c_i c_j = b_{1,i} b_{1,j} / c_1^2 = b_{i,j}$$

as well. Then

$$a_i = \sum_{j \neq i} b_{i,j} + \frac{b_{i,k}b_{i,\ell}}{b_{k,\ell}} = \sum_{j \neq i} c_i c_j + c_i^2 = c_i \sum_{j=1}^n c_j.$$

Since $a_i \neq 0$, the sum is nonzero; hence,

$$\sigma = \sum_{i=1}^{n} c_i \sum_{j=1}^{n} c_j = \left(\sum_{i=1}^{n} c_i\right)^2 \neq 0.$$

The \mathcal{K}_n is equivalent to the star because

$$\frac{a_i a_j}{\sigma} = \frac{(c_i \sum_{k=1}^n c_k) (c_j \sum_{k=1}^n c_k)}{(\sum_{k=1}^n c_k)^2} = c_i c_j = b_{i,j}.$$

For any finite graph G, there is a sequence of \bigstar - \mathcal{K} moves and parallel circuit reductions that will transform it into a graph with no interior vertices. Let Γ be a signed linear network on G, and suppose that at each step, the star is nonsingular, so an equivalent \mathcal{K} can be found. After the final step, the response matrix is exactly the Kirchhoff matrix because there are no interior vertices. So the \bigstar - \mathcal{K} transformation provides a way to compute the response matrix from the Kirchhoff matrix in small steps, and in some cases, this is a useful technique for determining recoverability over positive (real) linear conductances.

4 Layering I: Reduction Operations

4.1 Motivation: Layer-stripping for the Inverse Problem

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

A boundary spike on a B-graph 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 B-graph 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 B-graph G, we can define a new B-graph 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 [2] and Will Johnson [5] 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 layerstripping exhausts all the edges?

In this chapter, we will describe layer-stripping formally and address the second question above.

4.2 Definitions and Basic Properties

A reduction operation is a transformation of a B-graph G into a suB-graph 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 B-graph 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 Gor 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 suB-graphs $G = G_0 \supset G_1 \supset G_2 \supset \ldots$ such that $\bigcap_{n=0}^{\infty} G_n = \emptyset$. If G_{n+1} is obtained from G_n by a reduction operation, then we the filtration is called a *layerable filtration* and the B-graph is said to be *layerable*. A partial filtration is a sequence of suB-graphs $G = G_0 \supset G_1 \supset \ldots$, and it is a partial layerable filtration if each suB-graph is obtained from the previous one by a reduction operation.

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

Now we describe the electrical properties of reduction operations. Suppose that G' is obtained from G by a reduction operation. We will show that the boundary data L' is uniquely determined by L and the conductance functions of the edges removed in the reduction.

Lemma 4.1. Suppose that G' is obtained from G by contracting some nondegenerate 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 spike (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.

Lemma 4.2. 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.

Lemma 4.3. 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. \Box

5 Layering II: IO-Graphs and Factorization

5.1 The Category of Input-Output Graphs

In [1], John Baez and Brandan Fong describe "gluing networks together" in terms of a composition in a category of cospans. Here we present a simplified version of their definitions and some geometric-electrical applications. Roughly speaking, we will label some of the boundary vertices of our B-graph as "input" and some as "output" (allowing a vertex to be both input and output), and think of the B-graph (or a network) as a transformation from the input vertices to the output vertices. We compose such transformations by identifying the output vertices of the first with the input vertices of the second. Formally, a graph with input and output or IO-graph is a graph G together with two sets P and Q and injections $i: P \to V(G)$ and $j: Q \to V(G)$. In this case, we say the triple (G, i, j) is an IO-graph from P to Q. If (G, i, j) and (G, i', j') are two IO-graphs from P to Q, then we say they are isomorphic if there is a graph isomorphism $f: G \to G'$ such that the following commutes:



We define the category of IO-graphs as follows:

- The objects are sets.
- A morphism $\mathcal{G}: P \to Q$ is an isomorphism class of IO-graphs from P to Q.
- Composition is defined as follows: Suppose $\mathcal{G}_1 : P \to Q$ and $\mathcal{G}_2 : Q \to R$, and choose representatives (G_1, i_1, j_1) and (G_2, i_2, j_2) of the isomorphism classes. Let G be the graph obtained from the disjoint union of G_1 and G_2 by identifying $j_1(q)$ with $i_2(q)$ for each $q \in Q$. We define $i : P \to V(G)$ by composing i_1 with the obvious map $V(G_1) \to V(G)$ and $j : R \to V(G)$ is defined similarly. Then $\mathcal{G}_2 \circ \mathcal{G}_1 : P \to R$ is the isomorphism class represented by (G, i, j).
- The identity morphism $P \to P$ is a represented by a graph with no edges and V(G) = P, and *i* and *j* are the identity $P \to P$.

The reader may verify that this is well-defined. The reason to use isomorphism classes is that the disjoint union of graphs is only well-defined up to isomorphism.

Any IO-graph can be made into a B-graph by defining $B = i(P) \cup j(Q)$. Conversely, if we have a B-graph G and write B as a union of two sets P and Q, then G represents an IO-graph morphism from P to Q.

For any set S, one can define the category of input-output S-labelled graphs, and in particular, there is a category of IO-networks. To describe what the composition of IO-networks does to the boundary behavior of the networks, we first put the boundary behavior into a new form. Recall our networks take values in the abelian group M. Suppose that G is an IO graph from P to Qlabelled with potential-current relationships Θ_e . Consider G as a B-graph with $B = i(P) \cup j(Q)$, and let Γ be the corresponding network. We then define a relation

$$\Xi: (M^P \times M^P) \to (M^Q \times M^Q)$$

as follows: Let π_P be the projection $M^B \to M^P$ and $\pi_Q : M^B \to M^Q$. Let ι_P and ι_Q be the canonical inclusions $M^P \to M^Q$ and $M^Q \to M^B$. If x =

 $(x_1, x_2) \in M^P \times M^P$ and $y = (y_1, y_2) \in M^Q \times M^Q$, then we say $(x, y) \in \Xi$ if and only if there exists $(\phi, \psi) \in L$ such that

$$x_1 = \pi_P \phi, \qquad y_1 = \pi_Q \phi, \qquad \iota_P(x_2) - \iota_Q(y_2) = \psi.$$

To state it more concretely, $(x, y) \in \Xi$ if there exists a harmonic function on Γ with boundary potentials consistent with x_1 and y_1 and boundary net currents consistent with x_2 and y_2 . Here x_2 represents current flowing into the network at the input vertices, and y_2 represents current flowing out of the network at the output vertices. If a vertex is both input and output, then current can flow in at the input side and out at the output side.

If $\mathcal{G} : P \to Q$ is an IO-network morphism,¹ then $\Xi(\mathcal{G})$ is independent of the choice of representation for the isomorphism class. I claim we can define a functor $X : \mathbf{IO} - \mathbf{graph} \to \mathbf{Rel}$ by setting $X(P) = M^P \times M^P$ and $X(\mathcal{G}) =$ $\Xi(\mathcal{G})$. To see that X preserves composition, suppose $\mathcal{G}_1 : P \to Q$ and $\mathcal{G}_2 :$ $Q \to R$. Let $\mathcal{G} = \mathcal{G}_2 \circ \mathcal{G}_1$ and let $\Gamma_1, \Gamma_2, \Gamma$ be specific networks representing the isomorphism classes. We can assume without loss of generality Γ_1 and Γ_2 are subnetworks of Γ .

Composition in the category of relations gives us that $(x, z) \in \Xi(\mathcal{G}_2) \circ \Xi(\mathcal{G}_1)$ if and only if there exists some y with $(x, y) \in \Xi(\mathcal{G}_1)$ and $(y, z) \in \Xi(\mathcal{G}_2)$. In this case, (x, y) and (y, z) represent the boundary data of harmonic functions (u_1, c_1) and (u_2, c_2) on Γ_1 and Γ_2 . I claim these functions paste together to form a harmonic function Γ . Note that $V(\Gamma_1) \cap V(\Gamma_2) = i_2(Q) = j_1(Q)$, and the potentials of u_1 and u_2 on these vertices are both given by y_1 , so that u_1 and u_2 agree on $V(\Gamma_1) \cap V(\Gamma_2)$ and hence define a function $u : V(\Gamma) \to M$. Now $E(\Gamma_1) \cap E(\Gamma_2) = \emptyset$, so c_1 and c_2 define a function $c : E(\Gamma) \to M$. The net current at each interior vertex of Γ_1 or Γ_2 is already zero. On the vertices $i_2(Q) = j_1(Q)$, we know the output current on Γ_1 and the input current on Γ_2 are given by y_2 ; because of the difference in signs, these two net currents cancel. In particular, if one of these vertices is interior it has net current zero, and if it is in $i_1(P)$ and/or $j_2(R)$, then the net current is given by $x_2, -z_2$, or x_2 minus z_2 as appropriate. Thus, we can see $\Xi(\mathcal{G}_2) \circ \Xi(\mathcal{G}_1) \subset \Xi(\mathcal{G}_2 \circ \mathcal{G}_1)$, and the opposite inclusion follows by similar reasoning.

Thus, X respects composition, and the reader can check that it preserves the identity morphism and hence is a functor.

We make the convention that if $P = \varnothing$ then $M^P \times M^P$ is a one-element set.

5.2 Elementary IO-graphs

The category of IO graphs enables us to express complicated networks as compositions of simpler ones. Our building blocks are networks on the following four types of *elementary IO-graphs*:

1. A graph in which every component consists of either (a) an isolated vertex which is both an input and an output or (b) one edge and two vertices,

¹Unfortunately, I have used the same notation for isomorphism classes of networks as for graphs, but I don't suppose any confusion will result.

where one of the vertices is an input and the other is an output. See G_1, G_3, G_9 in Figure 2.

- 2. A graph in which all the vertices are both inputs and outputs. See G_2, G_4, G_8, G_{10} in Figure 2.
- 3. A graph with no edges in which every vertex is an input. We call the vertices which are not outputs *input stubs*. See G_5 in Figure 2.
- 4. A graph with no edges in which every vertex is an output. We call the vertices which are not inputs *output stubs*. See G_7 in Figure 2.

An elementary factorization of an IO-graph morphism $\mathcal{G}: P \to Q$ is a factorization $\mathcal{G} = \mathcal{G}_n \circ \circ \mathcal{G}_1$ such that each \mathcal{G}_j is represented by an elementary IO-graph and all the type 3 elementary IO-graphs come before (to the right of) the type 4 elementary IO-graphs. If \mathcal{G} is the identity morphism $P \to P = Q$ then we make the convention that it has an elementary factorization of length zero. If we are given a graph G representing $\mathcal{G}: P \to Q$ and an elementary factorization of \mathcal{G} , then we can assume without loss of generality that \mathcal{G}_j is represented by a suB-graph G_i of G, and we will often make this simplification.

Remark. The usefulness of the stipulation that the type 3 IO-graphs come before the type 4 ones will become clear in the next section. For the moment, the reader can verify that, if we allowed type 3 IO-graphs to come after type 4 IO-graph, then *any* morphism $\mathcal{G} : P \to Q$ represented by a finite graph would admit a factorization. It is unreasonable to expect such cheap factorizations to provide useful information.

Remark. In the case of finite graphs, any type 1 IO-graph can be factorized into type 1 IO-graphs with only one edge. It is more convenient for writing out specific factorizations if we allow several edges; but in proving general theorems we will often assume only one edge. The same considerations apply to the other types of elementary IO-graphs.

We define an elementary IO-network to be a network on an elementary IOgraph. The behavior of Ξ on such IO-networks is easy to describe in the cases we are interested in:

1. Suppose we have a type 1 network Γ with only one edge with input vertex $p = \iota(e)$ and output $q = \tau(e)$. Suppose Θ_e is given by a resistance function $\rho_e : M \to M$. Then $(x, y) = ((x_1, x_2), (y_1, y_2))$ is in Ξ if and only if

$$\begin{aligned} &(y_1)_q = (x_1)_p + \rho_e((x_2)_p) \\ &(y_2)_q = (x_2)_p \\ &(x_1)_r = (y_1)_r \text{ and } (x_2)_r = (y_2)_r \text{ for } r \neq p, q. \end{aligned}$$

As a result, Ξ defines a bijective function $M^P \times M^P \to M^Q \times M^Q$.



Figure 2: An elementary factorization. The inputs are shown in red and the outputs in blue.

2. Suppose we have a type 2 network Γ with only one edge with endpoints $p = \iota(e)$ and $q = \tau(e)$. Suppose Θ_e is given by a conductance function $\gamma_e : M \to M$. Then $(x, y) \in \Xi(\Gamma)$ if and only if

$$x_1 = y_1$$

$$(y_2)_p = (x_2)_p - \gamma_e((x_1)_p - (x_1)_q)$$

$$(y_2)_q = (x_2)_q + \gamma_e((x_1)_p - (x_1)_q)$$

$$(y_2)_r = (x_2)_r \text{ for } r \neq p, q.$$

As a result, Ξ defines a bijective function $M^P \times M^P \to M^Q \times M^Q$.

- 3. Suppose we have a type 3 network Γ with only one input stub p. Then $(x, y) \in \Xi$ if and only if $(x_2)_p = 0$ and for all $r \neq p$, $(x_1)_r = (y_1)_r$ and $(x_2)_r = (y_2)_r$.
- 4. The case of a type 4 network is symmetrical.

5.3 Factorization, Layerability, and Regularity

Any B-graph represents an IO-graph morphism $\mathcal{G}: B \to \emptyset$. Then

- 1. If \mathcal{G}' is a type 1 IO-graph morphism, then $\mathcal{G} \circ \mathcal{G}'$ is obtained from \mathcal{G} by adjoining a boundary spike. (Or to be precise, this holds for some pair of B-graphs representing $\mathcal{G} \circ \mathcal{G}'$ and \mathcal{G} .)
- 2. If \mathcal{G}' is a type 2 IO-graph morphism, then $\mathcal{G} \circ \mathcal{G}'$ is obtained from \mathcal{G} by adjoining a boundary edge.
- 3. Precomposing a type 3 IO-graph morphism corresponds to adding an isolated boundary vertex.
- 4. Precomposing a type 4 IO-graph morphism corresponds to changing a boundary vertex to interior.

The IO-graphs can thus be viewed as a geometric and categorical realization of reduction operations and other graph transformations. I invite the reader to reinterpret the proofs of §4.2 using elementary IO-networks.

If we consider $\mathcal{G} : \emptyset \to B$ instead and *post* compose the elementary B-graphs, then the roles of type 3 and type 4 networks are reversed. These considerations lead to ...

Lemma 5.1. Let G be a finite B-graph. The following are equivalent:

- a. G is layerable.
- b. The IO-graph morphism $B \rightarrow \emptyset$ represented by G admits an elementary factorization into networks of types 1, 2, and 3.
- c. For some $P, Q \subset B$ with $P \cup Q = B$, the morphism $P \to Q$ represented by G admits an elementary factorization.

Proof. (a) \implies (b). A sequence of reduction operations on G can be interpreted as a factorization into elementary IO-graphs. Details left to the reader.

(b) \implies (c) is trivial.

(c) \implies (a). Let *G* represent an IO-graph morphism \mathcal{G} with factorization $\mathcal{G} = \mathcal{G}_n \circ \cdots \circ \mathcal{G}_1$. We can choose some *k* such that $j \leq k$ for any type 3 network \mathcal{G}_j and j > k for any type 4 network \mathcal{G}_j . Then we define a layerable filtration of *G* using the suB-graphs that represent

$$\begin{array}{c}
\mathcal{G}_{n} \circ \cdots \circ \mathcal{G}_{1} \\
\mathcal{G}_{n-1} \circ \cdots \circ \mathcal{G}_{1} \\
\cdots \\
\mathcal{G}_{k+1} \circ \cdots \circ \mathcal{G}_{1} \\
\mathcal{G}_{k} \circ \cdots \circ \mathcal{G}_{1} \\
\mathcal{G}_{k} \circ \cdots \circ \mathcal{G}_{2} \\
\cdots \\
\mathcal{G}_{k} \circ \mathcal{G}_{k-1} \\
\mathcal{G}_{k} \\
\varnothing \qquad \qquad \Box
\end{array}$$

For layerable networks, this idea provides an easy way to parametrize the space of harmonic functions and the boundary behavior. If G is layerable, it is not hard to show that we can express $\mathcal{G}: B \to \emptyset$ in the form $\mathcal{G}_n \circ \cdots \circ \mathcal{G}_1 \circ \mathcal{G}_0$ where \mathcal{G}_j is a type 1 or type 2 network for $j \geq 1$ and \mathcal{G}_0 is a type 3 network with no outputs. Suppose Γ is a network on G with bijective conductance functions $\gamma_e: M \to M$. The boundary behavior of Γ_0 is $L_0 = \{(\phi, 0)\} \subset M^{(V(G_0))} \times M^{(V(G_0))}$ since any potentials are possible but the net current at each vertex must be zero. Then the boundary behavior of $\mathcal{G}_j \circ \cdots \circ \mathcal{G}_0$ is given by

$$L_j = \Xi(\Gamma_j) \circ \cdots \circ \Xi(\Gamma_1)(L_0).$$

Hence, we have a bijective parametrization of L_j by $M^{V(G_0)}$, and in particular $L = L_n$ has such a parametrization. In the process, we have also parametrized the space of harmonic functions \mathcal{H}_{Γ} . This, together with our explicit formula for $\Xi(\mathcal{G}_j)$ yields the following corollary:

Lemma 5.2. Suppose Γ is a layerable with bijective conductance functions γ_e . Then $\Phi : \mathcal{H}_{\Gamma} \to L_{\Gamma}$ is a bijection. Also,

- a. If M is a field F and γ_e is linear, then Φ is a linear isomorphism, and the space of harmonic functions has the "expected" dimension |B|.
- b. If M is a topological abelian group and γ_e is a homeomorphism $M \to M$, then Φ is a homeomorphism.
- c. If $M = \mathbb{R}$ or \mathbb{C} and γ_e is a diffeomorphism, then Φ is a diffeomorphism and L is a smooth manifold of dimension |B|.

We have already seen that (a) does not hold in general. It turns out that (c) can also fail without the assumption of layerability. Consider the following graph:



Define resistance functions $\mathbb{R} \to \mathbb{R}$ as follows: Let $\rho_{e_1}(t) = \rho_{e_3}(t) = t + \frac{1}{2} \sin t$ (the orientation of the edge does not matter since the function is odd), and let $\rho_{e_2}(t) = \rho_{e_3}(t) = -t$. These are bijective C^{∞} resistance functions with a C^{∞} inverse. The series with resistance functions ρ_{e_1} and ρ_{e_2} is equivalent to a single-edge with resistance $\rho_{e_1} + \rho_{e_2}$. Thus, the network is equivalent to a parallel connection



in which each edge has resistance function $\rho(t) = \frac{1}{2} \sin t$. Let e_1 and e_2 be the oriented edges shown in the picture. Thus, (u, c) is harmonic if and only if

$$u_1 - u_2 = \frac{1}{2}\sin c_{e_1} = \frac{1}{2}\sin c_{e_2}$$

Now $\sin c_{e_1} = \sin c_{e_2}$ is equivalent to $c_{e_2} = c_{e_1} + 2\pi n$ or $c_{e_2} = \pi - c_{e_1} + 2\pi n$. If $c_{e_1} = c_{e_2} + 2\pi n$, then the net current $\psi_1 = c_{e_1} + c_{e_2} = 2c_{e_1} + 2\pi n$ and $\psi_2 = -\psi_1$ and $u_1 - u_2$ must be $\frac{1}{2} \sin \psi_1/2$. If $c_{e_2} = \pi - c_{e_1} + 2\pi n$, then $\psi_1 = (2n+1)\pi$ and $\psi_2 = -\psi_1$ and $u_1 - u_2$ could be any number in [-1, 1]. Thus,

$$L = \{(\phi, \psi) : \phi_1 - \phi_2 = \frac{1}{2} \sin \psi_1 / 2, \ \psi_1 = -\psi_2 \}$$
$$\cup \{(\phi, \psi) : \phi_1 - \phi_2 \in [-1, 1], \ \psi_1 = (2n+1)\pi, \ \psi_2 = -\psi_1 \}.$$

This is not a smooth manifold because there is neighborhood of the points with $\phi_1 - \phi_2 = \pm 1$ and $\psi_1 = (2n+1)\pi$ which is not homeomorphic to \mathbb{R}^2 .

5.4 Factorization, Rank, and Connections

In this section, all graphs are assumed finite. Our goal is to find situations where we can guarantee rank $\Lambda_{P,Q} = m(P,Q)$ is valid for all conductances. First, we find a substitute for rank $\Lambda_{P,Q}$ that makes sense even when Λ is not defined and when P and Q are not necessarily disjoint.

For a linear IO-networks over a field representing a morphism $P \to Q, \Xi$ is a linear relationship, that is, a linear subspace of $(\mathbb{F}^P \times \mathbb{F}^P) \times (\mathbb{F}^Q \times \mathbb{F}^Q)$. In general, for finite-dimensional vector spaces W_1 and W_2 and a linear relationship $R: W_1 \to W_2$, we define rank R to be the maximal rank of a linear map T from some subspace of W_1 to a subspace of W_2 such that $(w, Tw) \in R$ for all w. Let π_1 and π_2 be the projections $R \to W_1$ and $R \to W_2$. Then R defines a linear isomorphism

$$\pi_1(R)/\pi_1 \circ \pi_2^{-1}(0) \to \pi_2(R)/\pi_2 \circ \pi_1^{-1}(0).$$

From this (with some linear algebra) we can see

rank
$$R = \dim \pi_1(R)/\pi_1 \circ \pi_2^{-1}(0) = \dim \pi_2(R)/\pi_2 \circ \pi_1^{-1}(0)$$

= $\dim \pi_1(R) - \dim \pi_2^{-1}(0) = \dim \pi_2(R) - \dim \pi_1^{-1}(0)$

(since dim $\pi_2 \circ \pi_1^{-1}(0) = \dim \pi_1^{-1}(0)$ and the same holds with π_1 and π_2 switched).

Proposition 5.3. Suppose the Dirichlet problem has a unique solution on a linear network Γ . Suppose P and Q are a partition of B, and let Ξ be the relation defined by Γ as an IO-network morphism $P \to Q$. Then rank $\Xi = 2 \operatorname{rank} \Lambda_{P,Q}$.

Proof. Let π_1, π_2 be as above for Ξ instead of R. Note that $\pi_2^{-1}(0)$ is isomorphic to the space of harmonic functions which have zero potential and net current on Q. Since the Dirichlet problem has a unique solution, these functions are parametrized by their potentials on P. Thus, the space is isomorphic to ker $\Lambda_{Q,P}$, so that dim $\pi_2^{-1}(0) = \dim \ker \Lambda_{Q,P}$. To compute $\pi_1(\Xi)$, we apply Λ to \mathbb{F}^B , then record the potential and current data on P. This means the matrix

$$\begin{pmatrix} I_P & 0 \\ \Lambda_{P,P} & \Lambda_{P,Q} \end{pmatrix}$$

maps $\mathbb{F}^B = \mathbb{F}^P \times \mathbb{F}^Q$ onto $\pi_1(\Xi) \subset \mathbb{F}^P \times \mathbb{F}^P$. Row reducing the left half will make the matrix block diagonal, and then we can see its rank is $|P| + \operatorname{rank} \Lambda_{P,Q}$. Hence,

$$\operatorname{rank} \Xi = \dim \pi_1(\Xi) - \dim \pi_2^{-1}(0)$$
$$= |P| + \operatorname{rank} \Lambda_{P,Q} - \dim \ker \Lambda_{Q,P}$$
$$= \operatorname{rank} \Lambda_{P,Q} + \operatorname{rank} \Lambda_{Q,P} = 2 \operatorname{rank} \Lambda_{P,Q}.$$

Exercise. Suppose P and Q are not disjoint but $B = P \cup Q$. Let $P' = P \setminus Q$ and $Q' = Q \setminus P$. Find rank Ξ in terms of rank $\Lambda_{P',Q'}$.

Suppose an IO-graph morphism $\mathcal{G}: P \to Q$ admits an elementary factorization $\mathcal{G}_n \circ \cdots \circ \mathcal{G}_1$ with $\mathcal{G}_j: P_{j-1} \to P_j$, then we define the *rank* of the factorization to be $\min_{0 \le j \le n} |P_j|$.

Proposition 5.4. Let G be a finite B-graph representing a morphism $\mathcal{G} : P \to Q$; let Γ be a linear network on G with nonzero conductances in \mathbb{F} . If \mathcal{G} has an elementary factorization of rank m, then rank $\Xi = 2n$.

Proof. We can assume without loss of generality each elementary IO-network has one edge or one stub. Let \mathcal{G}_j and P_j be as above. Let Ξ_j be the relation for \mathcal{G}_j , so that $\Xi = \Xi_n \circ \cdots \circ \Xi_1$. Choose k such that $\mathcal{G}_1, \ldots, \mathcal{G}_k$ are type 1, 2, or 3, and $\mathcal{G}_{k+1}, \ldots, \mathcal{G}_n$ are type 1, 2, or 4. We will parametrize the space of harmonic functions and Ξ by starting in the middle of the factorization and working our way outward.

Choose any $x_k \in \mathbb{F}^{P_k} \times \mathbb{F}^{P_k}$. If Ξ_k is type 1 or 2, there is a unique x_{k-1} such that $(x_{k-1}, x_k) \in \Xi_k$. If Ξ_k is type 3, then the data on the inputs which are not stubs is uniquely determined by x_k , and the current on the input stub must be zero, but the potential on the input stub is a free variable, so we have a one-dimensional affine space of functions x_{k-1} compatible with x_k . After determining the possible values of x_{k-1} , we repeat this process for $x_{k-2}, x_{k-3}, \ldots, x_0$. For each x_n , we can choose a compatible x_0, \ldots, x_{k-1} , and in fact the space of possible choices has dimension the number of type 3 networks (which we will call N_i). However, x_0 uniquely determines x_1 which uniquely determines x_2 , and so forth.

By a symmetrical argument, we can choose x_{k+1}, \ldots, x_n compatible with x_k and the space of choices has dimension the number of type 4 networks (call it N_o), and on the other hand, x_n uniquely determines x_k . The effect is that we have parametrized Ξ using x_n and the potentials on the input and output stubs.

In particular, $\pi_1(\Xi)$ is parametrized by x_k and the potentials on the input stubs and so has dimension $2|P_k| + N_i$, and $\dim \pi_2(\Xi) = 2|P_k| + N_o$. If $x_n = 0$, that forces x_k to be zero, so the possible choices of x_0 have dimension the number of input stubs; that is, $\dim \pi_2^{-1}(0) = N_i$, and similarly, $\dim \pi_1^{-1}(0) = N_o$. Therefore, rank $\Xi = 2|P_k| = 2m$ as desired.

Proposition 5.5. Let G be a finite B-graph representing a morphism $\mathcal{G} : P \to Q$. If \mathcal{G} has an elementary factorization of rank m, then m(P,Q) = m. Hence, rank $\Xi = 2m$.

Proof. Let \mathcal{G}_j , P_j , and k be as above. To create a connection of size m, we start from the middle of the factorization. We want each path to contain exactly one element of P_k . If \mathcal{G}_k is type 2 or 3, then our paths are length zero, and if it is type 1, we use the edges in the network for our paths and hence have a connection from P_k to P_{k-1} . We continue to extend the paths inductively. Once we have a connection from P_k to some subset R_j of P_j through $\mathcal{G}_k \circ \cdots \circ \mathcal{G}_j$, we extend the paths into \mathcal{G}_{j-1} -if it is type 2 or 3, there is nothing to do, and if it is type 1 we use the edges that have endpoints in R_j and thus obtain a connection to some subset of P_{j-1} . Hence, we have a connection from P_k to some subset of P_0 . In the same way, we can extend our paths from P_k through $\mathcal{G}_{k+1}, \ldots, \mathcal{G}_n$. Therefore, we have a connection of size m from P to Q, with the paths formed by edges from type 1 networks.

On the other hand, it is easy to verify (by induction on the number of elementary IO-graphs) that any path from a vertex in P to a vertex in Q must contain a vertex of every P_j . In particular, every path in a connection from a subset of P to a subset of Q must use a vertex from P_k , so there can be at most $|P_k| = m$ paths.

The maximum connection can thus be detected from L if we already know a factorization exists. Actually, this holds for nonlinear networks as well:

Proposition 5.6. Let Γ be a network on a bgraph G such that Θ_e is given by a conductance function $\gamma_e : \mathbb{R} \to \mathbb{R}$ which is a homeomorphism. Let $B = P \cup Q$ and suppose the IO-graph morphism $P \to Q$ represented by G has an elementary factorization of rank m with N_i input stubs and N_o output stubs. Then there are homeomorphisms

$$\pi_1(\Xi) \cong \mathbb{R}^{2m+N_i},$$

$$\pi_2(\Xi) \cong \mathbb{R}^{2m+N_o},$$

$$\pi_2^{-1}(x) \cong \mathbb{R}^{N_i} \text{ for any } x \in \pi_2(\Xi)$$

$$\pi_1^{-1}(x) \cong \mathbb{R}^{N_o} \text{ for any } x \in \pi_1(\Xi).$$

The proof is to parametrize each of the spaces in the same way as in Proposition 5.4. The upshot is this: From the invariance of domain theorem in topology, we know that each of the four spaces has a well-defined dimension as a topological manifold. If we define "rank Ξ " to be dim $\pi_1(\Xi) - \dim \pi_2^{-1}(x)$, then we again have rank $\Xi = 2m(P,Q)$, so that m(P,Q) is visible in the topology of Ξ (and hence L), even with nonlinear conductances that are merely bijective and continuous.

6 Layering III: Scaffolds

6.1 Motivation

While elementary factorizations allow us to describe a variety of electrical properties, they have several disadvantages:

- It can be cumbersome to list the vertices, edges, inputs, and outputs of each factor. This would also make it difficult to write out complete proofs constructing them.
- If $f: G \to G'$ is a locally injective B-graph morphism, then it should be possible to pull back a factorization of G' to produce a factorization of G. While this works for covering maps, it fails for inclusions. Indeed, if we view G' as an IO-graph morphism $P \to Q$ and factorize it, we cannot even say what the domain and codomain of the IO-graph morphism represented by G should be, since $f^{-1}(P)$ and $f^{-1}(Q)$ may be empty. Simply put, the system of factorizations does not define any nice functors on the category of B-graphs.
- Elementary factorizations do not generalize well to infinite graphs. We could perhaps make sense of composing an infinite sequence of IO-graphs, but we still have another issue: It was important that all the type 3 networks come before the type 4 networks, so if we have infinitely many

type 3 networks and type 4 networks, a sequence would be insufficient to describe our order of operations.

This chapter will develop more flexible machinery that addresses these issues. We can then describe harmonic continuation on infinite graphs, describe the recovery of boundary spikes and boundary edges in somewhat general situations, and formulate purely geometric conditions to guarantee recoverability over BZCF. This machinery lays the groundwork for the methods of constructing factorizations in later chapters.

6.2 Definitions

A scaffold S on G consists of

- A (strict) partial order \prec on E',
- A partition of E' into two sets Vert S and Hor S, whose elements are called respectively *vertical* and *horizontal edges*.
- Two functions $t, b : \operatorname{Vert} S \to V$ which assign a "top" endpoint t(e) and a "bottom" endpoint b(e) to each $e \in \operatorname{Vert} S$, 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 $e \in \operatorname{Vert} S$ 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 $e_1 \leq e_2$, then either $p_1 \in b(\operatorname{Vert} S)$ or $p_2 \in t(\operatorname{Vert} S)$.

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

Scaffolds behave nicely with respect to locally injective B-graph morphisms. Let Scaf G be the set of scaffolds on a B-graph G. Suppose $f : G_1 \to G_2$ is a locally injective B-graph morphism. Suppose $S \in \text{Scaf} G_2$. Then define f^*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 Vert $f^*\mathcal{S} = f^{-1}(\operatorname{Vert} \mathcal{S})$ and Hor $f^*\mathcal{S} = f^{-1}(\operatorname{Hor} \mathcal{S})$.
- Since the map is locally injective, we can define b, t for $f^*(S)$ such that f(b(e)) = b(f(e)) and f(t(e)) = t(f(e)).

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

Proposition 6.1. $G \mapsto \text{Scaf} G$ defines a contravariant functor from the category of B-graphs and locally injective B-graph morphisms to Set. It also satisfies $f^{-1}(\text{Mid} S) \subset \text{Mid} f^*S$.

Proof. Straightforward and left to the reader.

6.3 Scaffolds, Elementary Factorization, and Layerability

In the case of finite graphs, scaffolds describe the same structure as elementary factorization. We sketch the correspondence and leave some details to the reader. Suppose $P, Q \subset B$ and G has an elementary factorization as an IOgraph morphism $P \to Q$, given by $\mathcal{G} = \mathcal{G}_n \circ \cdots \circ \mathcal{G}_1$. Then we can define a scaffold \mathcal{S} by

- An edge is vertical if it occurs in one of the type 1 networks; the bottom endpoint is the input vertex and the top endpoint is the output vertex.
- We set $e \prec e'$ if e is in \mathcal{G}_j and e' is in \mathcal{G}_k for some j < k.

It is clear this is a partial order and satisfies axioms (1), (2), and (3) of the scaffold definition. For (4), note that if $p \notin b(\operatorname{Vert} S)$ is an interior vertex, then it must be an input stub in some elementary IO-graph, and if $q \in I(G) \setminus t(\operatorname{Vert} S)$, then it is an output stub. Hence, if e is incident to p and e' incident to q, then $e \prec e'$.

Conversely, suppose S is a scaffold on a finite graph G. We can easily complete the partial order on E' to a total order without changing the top, bottom, and middle of S, and write the edges in order as e_1, \ldots, e_n . it is possible that a boundary vertex is incident to two vertical edges (which does not happen when the scaffold comes from a factorization). In that case, we modify the scaffold as follows: Choose k such that Bot $S \subset \{e_j : j \leq k\}$ and Top $S \subset \{e_j : j > k\}$. Suppose that $e_i \prec e_j$ are vertical edges incident to a boundary vertex p. Then either $i \leq k$ or j > k. In the first case, we can change e_i to a horizontal edge, and we will still have a layering; if the other endpoint q of e_i was interior, then we have created a new interior vertex that is not the bottom endpoint of a vertical edge; but all the edges incident to q are $\leq e_i$, so that the Top and Bottom are still disjoint. In the case j > k, we change e_j to a horizontal edge.

By repeating this for each vertex, we obtain a scaffold S' in which every boundary vertex is incident to at most one vertical edge, and the edges are totally ordered. Take $P = B \setminus t(\operatorname{Vert} S')$ and $Q = B \setminus b(\operatorname{Vert} S')$. From the scaffold conditions, we deduce that e_1 is either a boundary spike with one endpoint in P or a boundary edge with both endpoints in P, so that we have a factorization $\mathcal{G} = \mathcal{S}_1 \circ \mathcal{G}_1$ where \mathcal{G}_1 is a type 1 or type 2 IO-graph. We repeat this process for the edges e_1, \ldots, e_k to get a factorization

$$\mathcal{G} = \mathcal{S}_k \circ \mathcal{G}_k \circ \cdots \circ \mathcal{G}_1.$$

Next, in a symmetrical way, we start at e_n and work our way downwards to factorize

$$\mathcal{S}_k = \mathcal{G}_n \circ \cdots \circ \mathcal{G}_{k+1} \circ \mathcal{T}.$$

Then \mathcal{T} cannot have any interior vertices, and so $\mathcal{T} = \mathcal{T}_1 \circ \mathcal{T}_2$ where \mathcal{T}_1 is an output-stub IO-graph and \mathcal{T}_1 is an input-stub IO-graph, and this complete the factorization.

Even in the infinite case, layerability is related to scaffolds through the following lemma:

Lemma 6.2. Let G be a B-graph 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 S on G with $\text{Top } S = \emptyset$.
- d. For any $e \in E'(G)$, there is a scaffold S on G with $e \notin \text{Top } S$.
- 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 $G = G_0 \supset G_1 \supset \ldots$ 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 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 S' with the same vertical edges and t and b functions as in 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 S, which will automatically be minimal with respect to S'.

I claim that for any $e \in E'(G)$, there are only finitely many edges $e \leq e_0$ in 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, \ldots, e_n which incident to and less than e_0 , and $\{e \leq e_0\} = \bigcup_{j=1}^n \{e \leq e_j\} \cup \{e_0\}$ 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 , $\{e \leq e_j\}$ is finite, which implies $\{e \leq e_0\}$ is finite, which is a contradiction.

Now choose e. Let $e_1, \ldots, e_k = e$ be the edges $\leq e$ in 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 suB-graph 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, \ldots . For each e_n , choose a k_n and a sequence of suB-graphs $G = G_{n,1} \supset \cdots \supset G_{n,k_n}$ as in (e). Then consider the following filtration:

$$G = G_{1,1}, \quad G_{1,2}, \quad \dots \quad G_{1,k_1},$$

$$G_{1,k_1} \cap G_{2,1}, \quad G_{1,k_1} \cap G_{2,2}, \quad \dots \quad G_{1,k_1} \cap G_{2,k_2},$$

$$G_{1,k_1} \cap G_{2,k_2} \cap G_{3,1}, \quad \dots \quad G_{1,k_1} \cap G_{2,k_2} \cap G_{3,k_3},$$

$$\dots \dots$$

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 isolated boundary vertices.

6.4 Scaffolds, Harmonic Extensions, and Recovery

We need two more pieces of terminology: Let $T \subset E'$. The suB-graph 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-B-graph $G' \subset G$ 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 S be a scaffold on G. We say that $G' \subset G$ is a *lower sub-B-graph* if $e \prec e' \in E(G')$ implies $e \in E(G')$. We say that $G' \subset G$ is an *upper sub-B-graph* if $e \succ e' \in E(G')$ implies $e \in E(G')$.

Let S be a scaffold on G, and that G' is an induced lower sub-B-graph. Then by definition of scaffold, $E'(G) \setminus E'(G')$ has some minimal element e_0 . Let G''be the sub-B-graph induced by $E(G') \cup \{e_0\}$. Then

- Suppose e_0 is vertical. By the conditions on the partial order, we know that $t(e_0)$ is not in G'. If $e_0 \notin Bot S$, that implies $t(e_0)$ is the bottom vertex of some vertical edge $e_1 \succ e_0$. Since G' is a lower suB-graph, this e_1 is not in G', hence it is not in G''. In particular, not all the edges incident to $t(e_0)$ are in G'', which implies $t(e_0)$ is a boundary vertex of G''. So e_0 is a boundary spike of G''. Thus, for BZCF networks, we can extend any harmonic function on G' to a harmonic function on G''.
- Suppose e_0 is horizontal and $e_0 \notin \text{Bot } S$. Then both endpoints are the bottom endpoint of some vertical edge, which cannot be in G'', which implies e_0 is a boundary edge of G''. Hence, any harmonic function on G'' extends to a harmonic function on G'.

This proves the following lemma:

Lemma 6.3. Suppose S is a scaffold on G and G' is an induced lower suBgraph. Let G'' be the suB-graph induced by $E'(G') \cup \{e_0\}$. Let e_0 be a minimal edge not in G' and suppose that $e_0 \notin Bot S$. Let Γ be a BZCF network on G. Then any harmonic function on Γ' has some extension to Γ'' .

The next lemma follows using transfinite induction:

Lemma 6.4. Suppose S is a scaffold on G and G' is an induced lower subgraph with Bot $S \subset E'(G')$. Let Γ be a BZCF network on G. Then any harmonic function on Γ' extends to a harmonic function on Γ .

Proof. Let u' be a harmonic potential on Γ' . Consider the set \mathcal{Z} of pairs (Σ, v) , where Σ is an induced lower subnetwork of Γ , $\Gamma' \subset \Sigma$, and v is a harmonic potential on Σ which equals u' on Γ' . 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^*) .

If Σ^* is not all of Γ , then there is a minimal edge not in Σ^* . Then by the previous lemma, we can extend v^* to a larger induced lower subnetwork, which contradicts maximality of (Σ^*, v^*) . So we are done.

Now instead of considering the existence of extensions, we consider uniqueness:

Lemma 6.5. Suppose S is a scaffold on G, and Γ is a BZCF network. Let Γ' be an induced lower subnetwork of Γ and suppose that Top $S \cap E(\Gamma') = \emptyset$. Then any harmonic function on Γ' is uniquely determined by

- The potential on the vertices of $B(\Gamma) \cap V(\Gamma')$.
- The net current on the vertices of $B(\Gamma) \cap b(\operatorname{Vert} S \cap E(\Gamma'))$.

Proof. Suppose that u' and v' are two harmonic potentials on Γ' with the same potential on $B(\Gamma) \cap V(\Gamma')$ and net current on $B(\Gamma) \cap b(\operatorname{Vert} S \cap E(\Gamma'))$. Suppose for contradiction that u' and v' do not agree on all of Γ' . Let T be the set of edges e in Γ' such that u' and v' disagree on one or both endpoints of e. Then T has a minimal element e_0 . Then

- Suppose e_0 is vertical. If $b(e_0)$ is a boundary vertex of Γ , then by assumption u' and v' have the same potential and net current at $b(e_0)$. Also, by minimality of e_0 , u' and v' agree on all edges less than e_0 , and in particular all other edges incident to $b(e_0)$. Thus, u' and v' must have the same current on e_0 , and hence the same voltage at $t(e_0)$, which contradicts our choice of e_0 .
- In the case where e_0 is vertical and $b(e_0)$ is interior in Γ , we know from $e_0 \notin \text{Top } \mathcal{S}$ that $b(e_0)$ has some edges incident to it at $b(e_0)$, and hence u' and v' have the same potential and net current on $b(e_0)$. The same argument yields a contradiction to the minimality of e_0 .
- Suppose e_0 is horizontal. Since $e_0 \notin \text{Top } S$, we conclude that each endpoint is either a boundary vertex of Γ or incident to edges less than e_0 . Thus, u' and v' have potentials which agree on both endpoints of e_0 , which contradicts our choice of e_0 .

The contradiction shows that T is empty and u' = v'.

Lemma 6.6. Let G be a B-graph. Suppose that either

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

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

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

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

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

I claim that the potential u_2 on Γ_2 which is t at p and 0 everywhere else is harmonic. Because $e_0 \in \text{Mid } S$, 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 . Hence, p is not adjacent to any interior vertices of Γ_2 , so u_2 is harmonic.

By Lemma 6.4, u_2 extends to a harmonic function u on Γ satisfying:

- a. u(p) = t.
- b. u = 0 on $B(\Gamma) \cap V(\Gamma_0)$.
- c. *u* has net current zero on $B(\Gamma) \cap b(\operatorname{Vert} S \cap E(\Gamma_0))$.

By Lemma 6.5, any harmonic function satisfying properties (a) - (c) must be identically zero on Γ_0 , and in particular have potential zero at q. 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 (a) - (c), 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_2 , note $p, q \in B(\Gamma_2)$ and e_0 is the only edge incident to p in Γ_2 . Define u_1 to be zero on Γ_2 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.

6.5 Solvable and Totally Layerable B-graphs

Let G_0, G_1, \ldots, G_n be a layerable filtration of a B-graph 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 B-graph 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 6.7. Any solvable B-graph 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 6.6, the conductance functions of the edges removed from G_n are uniquely determined by L_n over BZCF. By Lemma 4.3, 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. \Box

Proposition 6.8. Let $f: G \to G'$ be a locally injective B-graph 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 6.1.

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

Proposition 6.9.

- 1. If $f: G \to G'$ is a locally injective B-graph morphism and G' is totally layerable, then so is G.
- 2. Any totally layerable B-graph 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 6.1. (2) follows from Lemma 6.2. 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.

7 Graphs on Surfaces

7.1 Embeddings of Graphs, Strand Arrangements, and Medial Strands

The study of graphs on surfaces, especially the disk, has made heavy use of graph embeddings and the medial graph. Before we detail how to construct scaffolds and elementary factorizations using the medial graph, we need some technical definitions. Our goal here is to define "embedding" and "medial graph" in enough generality to cover a lot of degenerate cases that are not usually considered, so that we can safely handle very small subgraphs without having to modify our definitions ad hoc. We do not care whether the medial graph is well-defined, only that *some* medial graph is there for us to use.

For any graph G, there is a corresponding topological space, the quotient space obtained from $E \times [0, 1]$ by identifying (e, t) with $(\overline{e}, 1 - t)$ and identifying (e, 0) and (e', 0) if $\iota(e) = \iota(e')$. We will call this topological space G as well since no confusion will result. An *embedding* of a graph on a surface with boundary S is a function $f: G \to \overline{S}$ which is a homeomorphism onto its image, such that $f(x) \in \partial S$ if and only if x corresponds to a boundary vertex. Let's identify f(G) with G.

The embedding is *non-degenerate* if each component of $S \setminus G$ is homeomorphic to an open disk. Unfortunately, a non-degenerate embedding can easy become degenerate when we pass to a subgraph, or even delete a boundary edge.

It will be helpful to have a generalization of a "chord diagram" or "pseudoline arrangement," which we will call a "strand arrangement." A strand arrangement on a surface with boundary S is a collection of curves on S called strands such that

- Each strand s admits a parametrization f_s by [0, 1], S^1 , \mathbb{R} , or $[0, \infty)$ which is a closed map and is locally a homeomorphism.
- The endpoints of any strand parametrized by [0,1] or $[0,\infty)$ must be on ∂S . No other points of the strands are allowed to be on ∂S .
- For each $x \in \overline{S}$ there are at most two strand segments which intersect there. That is, $\bigcup_s f_s^{-1}(x)$ contains at most two points.
- Call a point where two strands intersect or one strand intersects itself a vertex. We assume the vertices form a discrete set, and none of these points are on ∂S .
- For any point $x \in \overline{S}$, there is a neighborhood that intersects at most two strands. This prevents infinitely many strands from accumulating near a point.

If \overline{S} is compact, then any strand arrangement will form a B-graph embedded on S where the vertices of the strands are interior vertices of the graph and the endpoints of the strands are boundary vertices (by some tedious topological argument). This fails in the non-compact case because some strands may run off to ∞ , but we do not give a hoot about it.

A *lens* in a strand arrangement is a loop formed by one or two arcs of strands. If an arrangement has no lenses, it is called *lensless*.

The components of S minus the union of the strands are called *cells*. Two cells \mathcal{A} and \mathcal{B} are *adjacent* if $\partial \mathcal{A} \cap \partial \mathcal{B}$ contains some strand segment. A *two-coloring* of the cells is an assignment of a "white" or "black" color to each cell such that no adjacent cells are the same color. Depending on the surface, not all strand arrangements may admit a two-coloring of the cells.

For a graph G embedded on S, a *compatible medial strand arrangement* is a strand arrangement with a two-coloring of the cells such that

- Each black cell is homeomorphic to the disk (though the closure might not be homeomorphic to the closed disk).
- There is a bijective correspondence between vertices of G and black cells such that each black cell contains the corresponding vertex.
- If \mathcal{A} is a black cell, then $\partial \mathcal{A}$ intersects ∂S if and only if the corresponding vertex of G is a boundary vertex.
- There is a bijective correspondence between the edges of G and vertices of the strand arrangement ("medial vertices") such that each edge of G contains the corresponding medial vertex and no other points of any strand.
- At each medial vertex, the two strands cross the edge e of G. That is, there is some neighborhood N of the vertex, homeomorphic to \mathbb{D} , such that the each parametrized strand moves from one component of $N \setminus e$ to the other.

Depending on how degenerate the embedding is, there may be many different compatible medial strand arrangements.

INCOMPLETE: Subgraph partitions and Elementary factorizations compatible with the embedding.

7.2 Producing Scaffolds from the Medial Graph

An *orientation* of a strand arrangement is a choice of orientation for each strand. It is *acyclic* if there is a no loop formed by oriented strand segments.

An orientation \mathcal{O} of the medial strands naturally produces a relation on E'and an assignment of vertical and horizontal edges (which may or may not form a scaffold). We can define a relation \prec on the medial vertices by setting $a \prec b$ if there is an increasing path from a to b along medial edges. Define \prec on the edges E'(G) by the relation on the corresponding medial vertices.

Suppose $e \in E'(G)$ corresponds to a medial vertex a. Define e to be vertical if and only if the ingoing medial edges at a are on the boundary of one black cell, and the outgoing edges are on the boundary of the other black cell. In this case, b(e) is the G vertex corresponding to the black cell bounded by the ingoing edges, and t(e) is the black cell bounded by the outgoing edges. Equivalently, e is vertical if the oriented strands cross e in opposite directions. Otherwise, e is horizontal.

For this to be a bona fide scaffold, we need to guarantee several things:

- ≺ defines a partial order; this is a equivalent to saying that the orientation of the medial strands is acyclic.
- For infinite graphs, every subset has a minimal element.
- The vertical and horizontal edges incident to a vertex have the correct behavior and the Top and Bottom are disjoint (scaffold conditions 2,3, 4).

The last condition will hold if we arrange that, for any interior medial black cell \mathcal{A} , $\partial \mathcal{A}$ can be partitioned into two arcs, the first arc when oriented according to \mathcal{O} moves counterclockwise around $\partial \mathcal{A}$, and the second arc moves clockwise (see Figure ...). For boundary medial black cells, we want the same behavior, except that either arc of $\partial \mathcal{A}$ is allowed to contain portions of the boundary of the surface. That is, all the strand segments in the first arc are oriented clockwise. This will guarantee that there are at most two vertical edges incident to any vertex of G, and the vertical edges are comparable with the other edges in the correct way. It also guarantees that each interior vertex is both the top and the bottom endpoint of some vertical edge, so that every edge will be in Mid \mathcal{S} and scaffold condition (4) is trivially satisfied.

To produce a scaffold, the orientation must be chosen judiciously. We will explain how to do this on the disk and half-plane, and leave other surfaces to future researchers.

7.3 B-graphs on the Disk

Lensless circular planar graphs have been well-studied by **??**. We will give an alternative proof of several of their results using scaffolds produced by the medial graph.

We use two types of orientations for the medial strands: First, fix $e^{i\theta} \in \partial \mathbb{D}$ and assume it is not the endpoint of any strrand. Define \mathcal{O}_{θ} as follows: If a strand *s* has endpoints $e^{i\alpha}$ and $e^{i\beta}$ with $\theta < \alpha < \beta < \theta + 2\pi$, then the positive direction moves from $e^{i\alpha}$ to $e^{i\beta}$.

The second orientation $\mathcal{O}_{\theta,\phi}$ is given by two points $e^{i\theta}, e^{i\phi} \in \partial \mathbb{D}$ which are not endpoints of any strand. We can assume $\theta < \phi \leq \theta + 2\pi$. Suppose s is a strand with endpoints $e^{i\alpha}$ and $e^{i\beta}$. After possibly switching α and β or changing their period, we can arrange that exactly one of the following holds:

- $\theta < \alpha < \beta < \phi$.
- $\theta < \alpha < \phi < \beta < \theta + 2\pi$.
- $\phi < \beta < \alpha < \theta + 2\pi$.

Then we say that the positive orientation of s moves from $e^{i\alpha}$ to $e^{i\beta}$.

Note that \mathcal{O}_{θ} is exactly $\mathcal{O}_{\theta,\theta+2\pi}$, so if we allow $e^{i\theta} = e^{i\phi}$, then we only have one type of orientation to worry about. Now to verify it is acyclic ...

Lemma 7.1. For any lensless strand arrangement in \mathbb{D} , $\mathcal{O}_{\theta,\phi}$ defines an acyclic orientation.

Proof. The proof is by induction on the number of strands. It clearly holds for one strand. Suppose it holds for n-1 strands and consider n strands s_1, \ldots, s_n with endpoints $e^{i\alpha_j}$ and $e^{i\beta_j}$ where α_j and β_j satisfy the relations above.

From the Jordan curve theorem, we know that $\mathbb{D} \setminus s_j$ has two components, one on the left of s_j and one on the right of s_j . Since the strand arrangement is

lensless, s_j can only cross s_k in one direction and the direction can be detected from the positions of the start and endpoints of s_j and s_k on $\partial \mathbb{D}$. I claim there is some s_k such that no s_j crosses from the left to the right of s_k . There are two cases:

- If there are any strands which satisfy $\phi < \beta_j < \alpha_j < \theta + 2\pi$, then let s_k be the one with the minimal value of α_k . Any s_j which crosses from left to right of s_k would have to have α_j between α_k and β_k . Then s_j must be one of the strands with $\phi < \beta_j < \alpha_j < \theta + 2\pi$ since otherwise α_j would be less than ϕ . But in that case, our choice of s_k implies $\alpha_k < \alpha_j$, so $\beta_k < \alpha_j < \alpha_k$ is impossible.
- Otherwise, we can assume all strands have $\alpha_j < \phi$ and $\alpha_j < \beta_j$. Choose s_k to have the minimum possible value of α_k . If s_j crossed from the left to the right of s_k , then we would have $\alpha_k < \beta_j < \beta_k < \alpha_k$ which contradicts our assumption that $\alpha_j < \beta_j$.

By reindexing assume k = n. From the induction hypothesis, s_1, \ldots, s_{n-1} do not form any oriented loops. Thus, if a loop exists it must contain some segment of s_n and clearly it cannot be all contained in s_n . When the loop exists s_n , it must move into the left component of $\mathbb{D} \setminus s_n$ because no strand crosses s_n from left to right. But then at some point the loop must return to or cross s_n from the left component of $\mathbb{D} \setminus s_n$ which implies there is some strand which crosses s_n from left to right, causing a contradiction. So there is no loop.

To describe the behavior of $\mathcal{O}_{\theta,\phi}$ on the boundary of a medial cell, we use

Lemma 7.2. Let \mathcal{A} be a cell of a lensless strand arrangement on \mathbb{D} . Let s_1, \ldots, s_n be the strands that intersect $\partial \mathcal{A}$, listed in CCW order around $\partial \mathcal{A}$ and oriented in the same direction as the CCW orientation of $\partial \mathcal{A}$ (with \mathcal{A} on the left of each s_j). Let x_j and y_j be respectively the start and end of s_j . Then x_1, \ldots, x_n occur in CCW order around $\partial \mathbb{D}$, and so do y_1, \ldots, y_n .

Remark. We do not assume in the hypothesis that s_1, \ldots, s_n are distinct, although that turns out to be true.

Proof. Suppose \mathcal{A} is an interior cell. Let z be the vertex of $\partial \mathcal{A}$ where s_1 and s_2 intersect. Let C be the counterclockwise arc of ∂D from x_1 to x_2 . Let h_1 and h_2 be the arcs of s_1 and s_2 from x_1 and x_2 to z, so that C, h_1 , and h_2 bound a geodesic triangle T.

Suppose for contradiction that there is some other $x_j \in C$. Let w be the first point where s_j hits ∂T . If $w \in h_2$, then s_j crosses s_2 there from left to right. It cannot intersect s_2 again since \mathcal{M} is lensless, but that implies it cannot intersect $\partial \mathcal{A}$ because \mathcal{A} is on the left side of s_2 . So suppose $w \in h_1$. Then at w, s_j crosses from the left to the right side of s_2 , and this occurs before the point z along s_2 , which implies $z \in \partial \mathcal{A}$ is on the right side of s_j . This also is impossible because \mathcal{A} is supposed to be on the left side of s_j .

This contradiction proves that there is no x_j between x_1 and x_2 , and the same argument applies to x_k and x_{k+1} for all k, hence x_1, \ldots, x_n occur in

counterclockwise order. By a symmetrical argument, y_1, \ldots, y_n occur in counterclockwise order. In the case of a boundary cell, similar reasoning applies except that arcs of $\partial \mathbb{D}$ may intervene between the strand segments; details left to the reader.

Now consider a medial cell \mathcal{A} , with s_j and x_j and y_j as above, but let s_j be oriented according to $\mathcal{O}_{\theta,\phi}$. Assume without loss of generality that x_1 is the first x_j on the counterclockwise side of $e^{i\theta}$. Note that as we proceed through the list s_1, \ldots, s_n ,

- 1. As long as $e^{i\theta}, x_j, e^{i\phi}$ and $e^{i\theta}, x_j, y_j$ occur in CCW order, then x_j is the start point of s_j in $\mathcal{O}_{\theta,\phi}$.
- 2. At some point x_j may pass $e^{i\phi}$ so that $e^{i\phi}, x_j, e^{i(\theta+2\pi)}$ are in CCW order.
- 3. At some point y_j may pass $e^{i\theta}$ so that $e^{i\theta}$, y_j , x_j occur in CCW order.
- 4. Either (2) or (3) may occur first, but after either one occurs, y_j will be the start point of s_j . And (2) or (3) will continue to hold for the rest of the strands.

Thus, $\partial \mathcal{A}$ can be divided into two arcs; before (2) or (3) occurs the orientation of s_j matches the CCW orientation of $\partial \mathcal{A}$, and after (2) or (3) occurs they are opposite.

This concludes the proof that $\mathcal{O}_{\theta,\phi}$ produces a scaffold; call it $\mathcal{S}_{\theta,\phi}$. We then have

Theorem 7.3 (cf. [2] Theorem ?? and [5] Theorem 6.7). A circular planar B-graph with a lensless medial graph is totally layerable, hence recoverable over BCZF for any M.

Proof. Let e be any edge and let a be the corresponding medial vertex, and s_1 and s_2 the geodesics that meet there. Note s_1 and s_2 divide D into four components, and e is contained in two opposite components. If $e^{i\theta}$ is on the boundary of one of the components that contains e, then e is horizontal in the scaffold S_{θ} , and if $e^{i\theta}$ is on the boundary of one of the boundary of one of the components. \Box

Any two "cut-points" $e^{i\theta}$ and $e^{i\phi}$ divide $\partial \mathbb{D}$ into two arcs; let C_1 be the CCW arc from $e^{i\theta}$ to $e^{i\phi}$ and let C_2 be the other arc. Let P and Q be the sets of vertices of G whose medial cells touch C_1 and C_2 respectively. Then P and Q are called a *circular pair*. $P \cap Q$ contains at most two vertices. The strands fall into three categories:

- A strand with both endpoints on C_1 is called *reentrant on* C_1 .
- A strand with both endpoints on C_2 is called *reentrant on* C_2 .
- A strand with one endpoint on C_1 and one on C_2 is called *transverse*.

Theorem 7.4 (cf. [2] Lemma ??). Let G be a lensless B-graph on \mathbb{D} . Assume the boundary of each medial cell intersects $\partial \mathbb{D}$ in at most one arc. Suppose P and Q are a circular pair induced by cut-points $e^{i\theta}$ and $e^{i\phi}$. Then

- a. The IO-graph morphism $\mathcal{G}: P \to Q$ represented by G admits an elementary factorization compatible with the medial strand arrangement.
- b. Hence, rank $\Xi(\mathcal{G}) = 2m(P,Q)$ for any network on G where a suitable notion of rank is defined.
- c. Also, $2 \cdot m(P,Q) = \#(\text{transverse strands}) + |P \cap Q|$.

Proof. We produce a factorization from the scaffold $S_{\theta,\phi}$ by a similar method to §6.3 (and S_{θ} would work as well). Using the Jordan curve theorem, the reentrant strands on C_1 do not intersect those on C_2 . Thus, the medial vertices on the C_1 -reentrant strands come before those on the C_2 -reentrant geodesics in our partial order, when they are comparable. Let W be the set of medial vertices a such that $a \not\succeq b$ for some b on a C_2 -reentrant geodesic. Then

- If there is a C_1 -reentrant strand with no medial vertices on it, then there is an isolated boundary vertex of G with its medial cell touching C_1 but not C_2 . (We assume that a medial cell only intersects C_1 in one arc, not two.)
- Otherwise, if W is nonempty, then choose a minimal element and let e be the corresponding edge in G. The strands which meet at e are tranverse or C_1 -reentrant, so e is either a boundary spike with an endpoint on C_1 or an boundary edge with both endpoints' medial cells touching C_1 .

In either case, we can write $\mathcal{G} = \mathcal{G}' \circ \mathcal{G}_1$ where \mathcal{G}_1 is an elementary IO-graph of type 1, 2, or 3 and the factorization can be represented by cutting \mathbb{D} into two components with a curve g_1 from $e^{i\theta}$ to $e^{i\phi}$.

The main component U_1 of $\mathbb{D} \setminus g_1$ is homeomorphic to \mathbb{D} and the scaffold satisfies all the same properties as before. (We may produce medial cells which intersect ∂U_1 in two arcs, but not we cannot produce any which intersect g_1 in two arcs.) We can repeat this process with U_1 instead of \mathbb{D} until there are no C_1 -reentrant strands without medial vertices, and W is empty. Continuing inductively, we obtain a factorization $\mathcal{G} = \mathcal{G}'' \circ \mathcal{G}_n \circ \cdots \circ \mathcal{G}_1$, where the \mathcal{G}_j 's are type 1, 2, or 3.

Next, we repeat this process starting at C_2 with a C_2 -reentrant strands with no medial vertices or a maximal medial vertex in our partial order. So we produce elementary IO-graph morphisms \mathcal{G}'_j of type 1, 2, or 4. In the end, we have $\mathcal{G} = \mathcal{G}'_1 \circ \cdots \circ \mathcal{G}'_m \circ \mathcal{G}^* \circ \mathcal{G}_n \circ \cdots \circ \mathcal{G}_1$. This \mathcal{G}^* has no reentrant strands and no edges in the graph; it represents the identity IO-graph morphism. Thus, the factorization is complete, proving (a), and (b) follows from earlier theorems.

In \mathcal{G}^* , all the medial cells touch both boundary arcs (so " $P^* = Q^*$ "). Since \mathcal{G}^* is in the middle of the factorization, the maximum connection between the two boundary arcs is the same for \mathcal{G}^* as for \mathcal{G} , that is the number of vertices of

 \mathcal{G}^* . Each cell in \mathcal{G}^* has either one or two strands on its boundary, and the cells with one strand correspond exactly to $P \cap Q$. From this (c) follows.

7.4 B-graphs on the Half-Plane

Consider a B-graph G embedded in the upper half-plane $\mathbb{H} \subset \mathbb{C}$ with a compatible medial graph such that

- Each medial strand begins and ends on \mathbb{R} rather than going off to ∞ .
- The medial graph is lensless.

These are the graphs [9] calls *supercritical*. Fix $t \in \mathbb{R}$ which is not the endpoint of any strand. Define an orientation \mathcal{O}_t as follows: If a strand s has endpoints x and $y \in \mathbb{R}$ with x < y, then we orient s from y to x if y < t and from x to y if y > t. I claim \mathcal{O}_t induces a scaffold on G.

To show \mathcal{O}_t is acyclic, choose a linear fractional transformation F mapping the half-plane onto the disk and suppose $t \mapsto e^{i\theta}$ and $\infty \mapsto e^{i\phi}$. Any oriented loop is formed by finitely many strands s_1, \ldots, s_n . By construction, $F(s_1), \ldots, F(s_n)$ are strands on the disk oriented according to $\mathcal{O}_{\theta,\phi}$, and hence they cannot form an oriented loop by Lemma 7.1. The same argument shows that the behavior of the boundary of each medial cells is the same as it was for the disk.

Thus, we only have to prove

Lemma 7.5. In the order induced by \mathcal{O}_t , every set of medial vertices has a minimal element.

Proof. Each strand s divides \mathbb{H} into two components, a bounded component which we will call the inside and an unbounded component which we will call the outside. As on the disk, we will divide the strands into three categories: reentrant on $(-\infty, t)$, reentrant on $(t, +\infty)$, and transverse.

I claim there exists a strand s_0 such that no other strand s crosses from the outside to the inside of s_0 . There are two cases:

- If there are any $(-\infty, t)$ -reentrant strands, let s_0 be the one with right endpoint y_0 closest to t (to the left of t). (There are only finitely many endpoints of strands in any bounded interval.) Any $(-\infty, t)$ -reentrant strand s that has its right endpoint left of x_0 cannot cross s_0 , but if its right endpoint is between x_0 and y_0 , then it starts on the inside of s_0 and hence cannot cross it from outside to inside. A transverse or (t, ∞) reentrant strand ends on the outside of s_0 and so cannot cross from outside to inside.
- Otherwise, we choose s_0 to have its right endpoint y_0 closest to t (to the right of t). Since all the other strands have their endpoints to the right of y_0 , they end outside it and hence cannot cross from outside to inside.

The same argument shows that for any subcollection of strands, there exists some s_0 such that no other strand in the collection crosses from outside to inside of s_0 .

To show that every set of medial vertices has a minimal element, it suffices to show that any decreasing path in the medial graph must terminate. The collection of strands which form the path has an element s_0 which no other strand in the collection crosses from outside to inside. At some point the decreasing path intersects s_0 . After that it can never move to the outside of s_0 . But there are only finitely many medial vertices inside s_0 or on s_0 and there are no oriented cycles, so the path must terminate.

Thus, by the same argument as for the disk, we have

Theorem 7.6. Any supercritical half-planar B-graph is totally layerable.

Remark. In [9], the networks have positive linear conductances and the network is recovered from the boundary data of minimum power solutions only, rather than all harmonic functions. The harmonic functions used for recovery of a boundary spike or boundary edge e_0 are thus chosen to be finitely supported. The method of scaffolds in §6.4 will produce a finitely supported function if we can arrange that only finitely many edges are $\succ e_0$.

The scaffold produced by \mathcal{O}_t does not do this. However, a slightly more complicated version does work. Choose t_0 and t_1 , and orient the strands as follows:

- If a strand has endpoints x and y with $t_0 < x < y < t_1$ or $x < t_0 < t_1 < y$, then the positive orientation moves from x to y.
- If a strand has one endpoint on $(-\infty, t_0) \cup (t_1, +\infty)$ and one on (t_0, t_1) , then the positive orientation moves from $(-\infty, t_0) \cup (t_1, +\infty)$ to (t_0, t_1) .
- If a strand has endpoints $x < y < t_0$ the positive orientation moves from y to x.
- If a strand has endpoints $t_1 < x < y$, the positive orientation moves from x to y.

The argument that this produces a scaffold is similar to what we have already done but with more casework, which we leave to the reader.

For any edge e_0 , let s_1 and s_2 be the strands which meet at e_0 , with endpoints $x_1 < x_2 < y_1 < y_2$. If we choose t_0 and t_1 such that $x_1 < t_0 < x_2$ and $y_2 < t_1$, then there will be only finitely many edges $\succ e_0$ in the scaffold. The same holds if we choose $t_0 < x_1$ and $y_1 < t_1 < y_2$. One of the two choices make e_0 vertical and the other makes it horizontal.

8 A Theorem on Unique Connections

In the case of factorizations into type 1 and type 2 networks, the relationship between connections, factorizations, and mixed problems is particularly strong, yielding equivalence between algebraic and geometric conditions, and the unexpected implication (e) \implies (d):

Theorem 8.1. Let G be a finite B-graph and assume each interior vertex has valence at least 2. Let $B = P \cup Q$ and $P' = P \setminus Q$ and $Q' = Q \setminus P$. The following are equivalent:

- a. There is a unique connection between P' and Q', and this connection uses all the interior vertices.
- b. There exists a scaffold S with $b(\operatorname{Vert} S) = V \setminus Q$ and $t(\operatorname{Vert} S) = V \setminus P$.
- c. The IO-graph morphism $P \rightarrow Q$ represented by G admits a factorization into type 1 and type 2 networks.
- d. For any M and any network Γ given by bijective conductance functions, potentials on P and net currents on P' determine a unique harmonic function on Γ .
- e. For any signed linear network Γ over \mathbb{R} , potentials on P and net currents on P' determine a unique harmonic function on the network.

Remark. Let (*) be the condition that potentials on P and net currents on P' determine a unique harmonic function. In (2) and (3) it is important that (*) holds for *all* conductances. Even if it holds for *most* signed linear conductances, the stubless layering may not exist.

Proof. (b) \implies (c) \implies (d) follows from the general theory developed so far, and (d) \implies (e) is trivial.

To prove (e) \implies (a), note that tor signed linear conductances $\{a_e\}$, (*) is equivalent to the submatrix $K_{P'\cup I,Q'\cup I}$ being invertible. If this holds for all signed linear conductances, then $\mathcal{F}(P,Q)$ has exactly one element by Proposition 3.4. Let F be this element. Each component contains either one vertex in $P \cap Q$, or it contains one vertex in P' and one in Q'. Each component is a tree, but I claim that each component is actually a path. Otherwise, there would be an interior vertex p with only one edge e in F incident to it. By assumption, there is another edge e' incident to p. The other endpoint of e' is in some component of F, so $F \setminus \{e\} \cup \{e'\}$ is another spanning forest. The components of F thus provide a connection from P to Q. The connection is unique because if there were another connection, then we could add edges to complete it to a different spanning forest.

(a) \implies (b). There is a unique connection between P' and Q' if and only if there is a unique connection between P and Q, as a simple consequence of our definition of connection. We define a scaffold S as follows:

1. The vertical edges are the edges in the paths of the connection, and the "increasing" orientation is the same as their orientation in the path.

2. We define \prec by setting $e \prec e'$ if $e \in \operatorname{Vert} S$ and e' are incident at t(e) and $e \succ e'$ if $e \in \operatorname{Vert} S$ and e' are incident at b(e), and then taking the transitive closure.

By construction this satisfies the conditions in (b), and every interior vertex is both the top and bottom endpoint of a vertical edge. The only thing left to check is that \prec defines a partial order. It suffices to show that there is no "precedence loop"

$$e_1 \prec e_2 \prec \cdots \prec e_K \prec e_1,$$

in which each pair of edges is comparable by the primitive relations given in (2); let's assume each e_1 is an oriented edge such that e_1, \ldots, e_K is a path. The basic idea is that if we had such a loop, then we could construct a different connection between P and Q as indicated in Figure 3.

To make this rigorous, consider the precedence loops with the minimal number of horizontal edges, and from those choose one with the minimal number of edges. Let $\alpha_1, \ldots, \alpha_n$ be the paths in the connection. Then observe:

- Any precedence loop must contain horizontal edges, since otherwise it would have to be contained in one of the α_m 's, which is impossible. We also cannot have two horizontal edges in a row since the primitive relations do not compare horizontal edges.
- In the loop which we chose, e_1, \ldots, e_K must be distinct, since otherwise we could find a loop with either fewer horizontal edges or the same number of horizontal edges and fewer vertical edges.
- Suppose there are some i < j < k with e_j horizontal and e_i and e_k vertical edges in the same path α_m , and that e_i comes before e_k in the path α_m . If we replace the segment $e_{i+1} \dots e_{k-1}$ of the loop with the segment of α_m from e_i to e_k , then we get a precedence loop with fewer horizontal edges. Thus, this cannot happen in our chosen loop. The same reasoning holds for any cyclic permutation of the indices $1, \dots, K$. Thus, the loop must intersect each path in an "interval"; that is, $I_m = \{k : e_k \in \alpha_m\}$ is of the form $\{1, \dots, k\}$ after some cyclic permutation of the indices.

Hence, our loop has the following form: It move upward along some path of the connection (which we will call α_1 after reindexing), then crosses by a horizontal edge to some other α_2 , and it continues in the same way until it crosses from some α_ℓ back to α_1 . The paths $\alpha_1, \ldots, \alpha_\ell$ are distinct. It follows that the vertices in our loop must be distinct and the loop looks essentially like the one portrayed in the Figure except that it might not visit every path. If the remaining paths are $\alpha_{\ell+1}, \ldots, \alpha_n$, then we construct our new connection as follows: $\alpha'_j = \alpha_j$ for $j = \ell + 1, \ldots, n$. For $j = 1, \ldots, \ell, \alpha'_j$ follows α_j until it meets an endpoint of a horizontal edge from the loop, then it crosses to α_{j-1} following this horizontal edge in the reverse orientation from the loop, and it continues along α_{j-1} until it reaches Q (indices written mod ℓ). So (b) is proved.



Figure 3: Proof of Theorem 8.1: (a) \implies (b)

References

- John C. Baez and Brandan Fong. "A Compositional Framework for Passive Linear Networks." http://arxiv.org/pdf/1504.05625v2.pdf
- [2] Edward B. Curtis and James A. Morrow. Inverse Problems for Electrical Networks. World Scientific. 2000.
- [3] Robin Forman. "Determinants of Laplacians on Graphs." *Topology* Vol. 32, No. 1, pp. 35-46, 1993.
- [4] David Ingerman. "Theory of Equivalent Networks and Some of its Applications." http://www.math.washington.edu/~reu/papers/1992/ingerman/ingerman.pdf
- [5] Will Johnson. "Circular Planar Resistor Networks with Nonlinear and Signed Conductors." http://arxiv.org/abs/1203.4045
- [6] Thomas Lam and Pavlo Pylyavkyy. "Electrical Networks and Lie Theory." http://arxiv.org/pdf/1103.3475v2.pdf
- [7] Jeffrey Russell. "★ and K Solve the Inverse Problem."
 http://www.math.washington.edu/~reu/papers/2003/russell/recovery.pdf
- [8] Konrad Schrøder. "Mixed-Sign Conductor Networks." http://www.math.washington.edu/~reu/papers/1993/schroder/sign.pdf
- [9] Ian Zemke. "Infinite Electrical Networks: Forward and Inverse Problems." https://www.math.washington.edu/~morrow/papers/ian-thesis.pdf