Pseudoline arrangements and the Inverse Boundary Value Problem in Nonlinear Electrical Networks

Will Johnson

January 9, 2012

Abstract

We consider the inverse boundary value problem in the case of discrete electrical networks containing nonlinear (non-ohmic) conductors. For a fixed nonlinear electrical network, we show that under reasonable assumptions, there are well-defined Dirichlet-to-Neumann and Neumann-to-Dirichlet maps relating boundary voltages and boundary currents. We also generalize work of Curtis, Morrow, and others, characterizing the circular planar graphs for which the inverse boundary value problem has a solution, in the sense that the Dirichlet-to-Neumann maps determine all interior conductors.

1 Introduction

Consider a network of resistors, containing several boundary nodes where current is allowed to flow in and out. If we fix the voltages at these boundary nodes, all interior voltages and currents will be determined, and in particular the amount of current flowing in and out at each boundary node is determined. Thus, boundary voltages determine boundary currents. The relationship between boundary voltages and boundary currents can be encapsulated in a matrix Λ, the *response matrix* or *Dirichlet-to-Neumann matrix*.

Viewing the circuit as a black box whose internals are hidden, the matrix Λ exactly describes the behavior of the circuit. It is never possible to determine the internal structure of a circuit from Λ, but if we are given the structure of the underlying graph, then in many cases we can determine all the resistances from Λ. This problem is a discretization of Calderon's inverse boundary value problem. TODO history. Curtis, Morrow, and students in the Mathematics REU at the University of Washington have essentially solved the discrete problem in the case where the network is *circular planar*, in the sense that it can be drawn as a planar graph with all boundary nodes along the outside [1]. In this case, the criterion for a graph to be recoverable can be obtained by considering an associated graph, the *medial graph*, and checking a simple condition.

In this paper, we generalize these results to allow for non-linear (non-ohmic) conductors. We first assume that at each edge in the graph, current is given as a continuous, monotone, zero-preserving function of voltage. In this case there is a well-defined map from boundary voltages to boundary currents. Dually, we next assume that voltage across each edge is given by a continuous, monotone, zero-preserving function of current, showing in this case that boundary currents almost uniquely determine boundary voltages. Finally, we consider the inverse problem of going from these relations to determine the current-voltage relation on each edge. For the inverse problem, we must make the stronger assumption that the relation between current and voltage at each edge is given by a zero-preserving bijection. We consider circular planar networks, and show that recovery is possible in exactly the same graphs as in the linear case. This is done using properties of critical medial graphs, which are equivalent to simple pseudoline arrangements.

Surprisingly, this recovery method works even when the conductance functions are not monotone or continuous. In particular, this shows that in the linear case, we can recover negative conductivities. Negative conductivities appear naturally in some generalizations of Y-Δ transformations, so our results have implications on the ordinary linear recoverability of some non-planar graphs.

2 The Dirichlet-to-Neumann Map

Definition 2.1. *A* graph with boundary Γ *is a graph whose set of vertices* V *have been partitioned into a set of boundary vertices* $\partial\Gamma$ *and interior vertices* Γ^{int}.

We allow our graphs to contain multiple edges and self-loops, though the latter are completely irrelevant in what follows. If e is a directed edge in our graph, we let $\overline{e}$ be the edge going in the opposite direction.

Definition 2.2. *If* Γ *is a graph with boundary, a* conductivity (conductance) network on Γ *is a function that assigns to each directed edge* e *a function* $\gamma_e : \mathbb{R} \to \mathbb{R}$ *which is*

- *zero-preserving:* $\gamma_e(0) = 0$
- *monotone:* $\gamma_e(x) \geq \gamma_e(y)$ *if* $x \geq y$
- *continuous*

such that $\gamma_e(-x) = -\gamma_{\overline{e}}(x)$.

We abuse notation and use Γ for both the network and the graph.

Definition 2.3. *A* voltage function *on a graph with boundary* Γ *is a function* ϕ *from the vertices of* Γ *to* $\mathbb{R}$. *A* current function *on* Γ *is a function* ι *from the directed edges of* Γ *to* $\mathbb{R}$, *such that*

$$\iota(e) = -\iota(\overline{e})$$

and for every vertex v *in the* interior *of* Γ, $\sum_e \iota(e) = 0$, *where the sum ranges over all edges leaving the vertex* v. *A voltage function* ϕ *and current function* ι *are* compatible (with respect to some conductance network on Γ) *if*

$$\iota(e) = \gamma_e(\phi(x) - \phi(y))$$

whenever e *is an edge leading from vertex* x *to vertex* y. *Given such a pair, the* boundary voltage function *is the restriction of* ϕ *to* $\partial\Gamma$, *while the* boundary current function *is the function that assigns to each vertex $v \in \partial\Gamma$, the sum $\sum_e \iota(e)$ where e ranges over the edges leaving from v.*

In other words, the interpretation of the function γ_e is the function that specifies current in terms of voltage along that edge. In the *linear* case, $\gamma_e(x) = c_e x$, where c_e is the conductivity (reciprocal of resistance) along edge e.

Given a conductivity network and a set of boundary voltages, the *Dirichlet boundary value problem* is to find a voltage function with a compatible current function extending the boundary values. In general, the solution to the Dirichlet problem may not be unique, but solutions always exist, and the current function is uniquely determined:

Theorem 2.4. *Let Γ be a conductivity network and let f be any function from $\partial\Gamma$ to $\mathbb{R}$. Then there is a voltage function ϕ and a compatible current function ι on Γ that has f as the boundary voltage function. Moreover, ι is uniquely determined.*

Note that ϕ need not be uniquely determined. This occurs for example if every function γ_e is the zero function and there are any interior vertices.

Proof. We first show existence. Let $\mathbb{R}^V$ be the vector space of all voltage functions. Let W be the affine subspace of $\mathbb{R}^V$ containing only functions which agree with f on $\partial\Gamma$. For each directed edge e, let

$$q_e(x) = \int_0^x \gamma_e(t)dt$$

This function is C^1, nonnegative, and convex because of the properties of γ_e. We also have $q_e(0) = 0$, and $q_e(-x) = q_{\overline{e}}(x)$. Moreover, if x and y are both nonnegative (or both nonpositive), and $|x| \leq |y|$, then

$$q_e(x) \leq q_e(y). \tag{1}$$

This is easy to verify.

For $\phi \in \mathbb{R}^V$, we define the *pseudopower* $Q(\phi)$ to be

$$\sum_{e \in E} q_e(\phi(\text{start}(e)) - \phi(\text{end}(e)))$$

where start(e) and end(e) are the start and end of e, and the sum is over all directed edges. Note that the terms for e and $\overline{e}$ agree, so we are essentially counting each edge twice.

In the linear case, this function is the usual power, the sum over voltage times current for each edge. In the more general case, however, it is not this value, so we call it the pseudopower.

As a sum of nonnegative convex C^1 functions, pseudopower is a nonnegative convex C^1 function on $\mathbb{R}^V$, and also on W. We aim to show that $Q(\cdot)$ attains a minimum on W. To see this, first let M be the maximum of $|f(v)|$ as v ranges over $\partial\Gamma$. Let K be the hypercube subset of W containing only voltage functions which take values in $[-M, M]$. Then $Q(\cdot)$ attains a minimum on K because K is compact. Let ϕ_0 be such a minimum on K. We claim that ϕ_0 is in fact a global minimum. To see this, let

ϕ be any other voltage function in W. Let $\mu : \mathbb{R} \to [-M, M]$ be the retraction of $\mathbb{R}$ onto $[-M, M]$ that sends (M, ∞) to M and $(-\infty, -M)$ to $-M$. Let $\phi' = \mu \circ \phi$. Then clearly ϕ' still agrees with f on the boundary, so $\phi' \in W$, and in fact $\phi' \in K$. Moreover, because μ is a monotone short map and because the q_e functions satisfy Equation (1), it follows easily that

$$Q(\phi) \geq Q(\phi') \geq Q(\phi_0).$$

Since $\phi \in W$ was arbitrary, ϕ_0 does indeed minimize Q on all of W.

For any directed edge e from vertex x to vertex y, let $\iota(e) = \gamma_e(\phi(x) - \phi(y))$. Then ι and ϕ will be compatible current and voltage functions as long as ι is a current function. The property that $\iota(e) = -\iota(\overline{e})$ follows easily from properties of γ_e, so it remains to show that for every interior vertex v,

$$\sum_{e \text{ with } v=\text{start}(e)} \iota(e) = 0.$$

This follows directly by taking the partial derivative of $Q(\phi)$ with respect to $\phi(v)$, and using the fact that ϕ_0 is a critical point. Specifically, the partial derivative

$$\frac{\partial Q(\phi)}{\phi(v)} = \sum_{e \text{ with } v=\text{start}(e)} q_e'(\phi(v) - \phi(\text{end}(e))) - \sum_{e \text{ with } v=\text{end}(e)} q_e'(\phi(\text{start}(e)) - \phi(v))$$

$$= \sum_{e \text{ with } v=\text{start}(e)} \gamma_e(\phi(v) - \phi(\text{end}(e))) - \sum_{e \text{ with } v=\text{end}(e)} \gamma_e(\phi(\text{start}(e)) - \phi(v)) =$$

$$2 \sum_{e \text{ with } v=\text{start}(e)} \gamma_e(\phi(v) - \phi(\text{end}(e)))$$

where the last line follows by the change of variables $e \leftrightarrow \overline{e}$ in the second sum, and the fact that

$$-\gamma_e(\phi(\text{start}(e)) - \phi(\text{end}(e))) = \gamma_{\overline{e}}(\phi(\text{start}(\overline{e})) - \phi(\text{end}(\overline{e}))).$$

This establishes existence. It remains to show that the current function ι is unique. Suppose we had two current functions ι_1 and ι_2 respectively compatible with two voltage functions ϕ_1 and ϕ_2, both of which agreed with f on $\partial\Gamma$. Define a new conductivity network $\hat{\Gamma}$ on the same graph by letting

$$\hat{\gamma_e}(x) = \gamma_e(x + \phi_2(\text{start}(e)) - \phi_2(\text{end})(e)) - \iota_2(e)$$

This is a valid conductance function because

$$\gamma_e(\phi_2(\text{start}(e)) - \phi_2(\text{end})(e)) = \iota_2(e)$$

implies that $\hat{\gamma_e}(0) = 0$, and the other properties are easy to verify. One easily verifies that $\hat{\phi} = \phi_1 - \phi_2$ and $\hat{\iota} = \iota_1 - \iota_2$ are now compatible voltage and current functions for $\hat{\Gamma}$. Moreover, the boundary values of $\hat{\phi}$ are all zero, because ϕ_1 and ϕ_2 agreed on the boundary.

Consider the directed subgraph of Γ in which we include a directed edge e if and only if $\hat{\imath}(e) > 0$. Since $\hat{\imath}$ equals

$$\hat{\gamma_e}(\hat{\phi}(\text{start}(e)) - \hat{\phi}(\text{end}(e)))$$

and $\hat{\gamma_e}$ is monotone and zero-preserving, every edge e in the subgraph must also satisfy

$$\hat{\phi}(\text{start}(e)) > \hat{\phi}(\text{end}(e))$$

In particular then, our subgraph must be acyclic. We claim that it has no interior 'sources' or 'sinks' in the following sense: there cannot be any interior node v which has at least one coming in and none coming out, or at least one edge going out and none coming in. These follow from the facts that the sum of $\hat{\imath}(e)$ for the edges incident to v is zero, and a sum of positive and nonnegative numbers cannot be zero.

For $x, y \in V$, let $x \to y$ indicate that there is a chain of edges in the subgraph leading from x to y, oriented in the correct way. This relation is obviously transitive, and irreflexive because $x \to y$ implies that $\hat{\phi(x)} > \hat{\phi(y)}$. Suppose for the sake of contradiction that there is at least one edge e in the subgraph. Let e go from x_0 to y_0. Then choose y with $y = y_0$ or $y \to y_0$, that is maximal in the sense that no y' exists with $y \to y'$. Similarly choose a minimal x' with $x' \to x_0$ or $x' = x_0$. Because there are no internal sources or sinks, x_0 and y_0 must be boundary nodes. We also have $x_0 \to y_0$. Thus $\hat{\phi}(x_0) > \hat{\phi}(y_0)$, contradicting the fact that $\hat{\phi}$ vanishes on the boundary.

So this subgraph is empty. This means that for every edge e, $\hat{\imath}(e) \leq 0$. But then for every edge e we also have $\hat{\imath}(\bar{e}) = -\hat{\imath}(e) \leq 0$. So $\hat{\imath}$ vanishes everywhere, and ι_1 and ι_2 agree. □

Since the current function determines the boundary current function, it follows that there is a well-defined map, the *Dirichlet-to-Neumann map*, from boundary voltage functions to boundary current functions.

3 The Neumann-to-Dirichlet Map

This section is completely dual to the previous one. Here we assume that voltage is given as a suitably nice function of current, and show essentially that boundary currents uniquely determine boundary voltages. There are some caveats, however: a given set of boundary currents must add up to zero on each connected component of a graph, or else there are no solutions at all. Additionally, there is latitude in choosing the voltages: adding a constant to each value of a voltage function preserves compatibility. In fact, we can add a separate constant on each connected component of Γ. The solution then is to normalize boundary currents and voltages by assuming both sum to zero on each connected component.

Definition 3.1. *If Γ is a graph with boundary, a* resistance network on Γ *is a function that assigns to each directed edge e a function $\rho_e : \mathbb{R} \to \mathbb{R}$ which is*

- *zero-preserving:* $\rho_e(0) = 0$
- *monotone:* $\rho_e(x) \geq \rho_e(y)$ *if* $x \geq y$

- *continuous*

such that $\rho_e(-x) = -\rho_{\overline{e}}(x)$.

Definition 3.2. *A voltage function* ϕ *and current function* ι *are* compatible (with respect to some resistance network on Γ) *if*

$$\rho_e(\iota(e)) = \phi(x) - \phi(y)$$

whenever e *is an edge leading from vertex* x *to vertex* y. *Given such a pair, the* boundary voltage function *is the restriction of* ϕ *to* $\partial\Gamma$, *while the* boundary current function *is the function that assigns to each vertex* $v \in \partial\Gamma$, *the sum* $\sum_e \iota(e)$ *where* e *ranges over the edges leaving from* v.

In other words, the interpretation of the function ρ_e is the function that specifies voltage in terms of current along that edge. In the *linear* case, $\rho_e(x) = r_e x$, where r_e is the resistance (reciprocal of conductance) along edge e.

Theorem 3.3. *Let* f *be a real-valued function on* $\partial\Gamma$, *summing to zero on each connected component of* Γ. *Then there is a voltage function* ϕ *and a compatible current function* ι *with* f *as the boundary current function. Moreover, given two such solutions* ϕ_1 *and* ϕ_2, $\phi_1 - \phi_2$ *is constant on each connected component of* Γ.

Proof. We first show existence. Let W be the affine space of current functions on Γ which have f as their boundary currents. This space is nonempty because of the requirement that f sum to zero on each connected component - we leave the proof of this as an exercise to the reader. For each edge e, let

$$q_e(x) = \int_0^x \rho_e(t)dt$$

As before, these functions are nonnegative, convex, and C^1 and satisfy Equation (1). For $\iota \in W$, let $Q(\iota)$ be the sum

$$Q(\iota) = \sum_e q_e(\iota(e))$$

Then as before, $Q(\cdot)$ is convex, nonnegative, and C^1.

Let $M = \sum_{v\in\partial\Gamma} |f(v)|$ and let K be the set of all current functions $\iota \in W$ for which there are no directed cycles of edges $e_1, e_2, \ldots, e_n$ with $\iota(e_i) > 0$ for all i. We claim that K is compact, because $|\iota(e)| \leq M$ for $\iota \in K$ and e any edge. This is true because the acyclicity of ι implies that we can partition the vertices of Γ into a set of vertices 'upstream' from e and a set of vertices 'downstream' from e, and the total current flowing through e must be at most the total current flowing into the upstream set, which is at most M. So K is bounded, and clearly closed. Thus Q attains a minimum on K, at some ι_0.

As before, we wish to show that ι_0 is a global minimum. To see this, take any $\iota \in W \setminus K$. Then ι contains a directed cycle of edges along which current is flowing in a positive direction. Letting m be the minimum of these positive current values, create another current function ι' by subtracting m from the current along each of these edges (and adding m to the currents along the edges going in the opposite directions). Then clearly ι' is still a current function, and it has the same boundary currents as ι. Moreover, since the current did not change in sign along any edge, and did not increase in magnitude along any edge, it follows by the analogue of Equation (1) that $Q(\iota') \leq Q(\iota)$. Now ι' may have more cycles, and so may not yet be in K. However, by repeating this procedure at most finitely many times, we get an $\iota'' \in K$ with $Q(\iota) \geq Q(\iota'') \geq Q(\iota_0)$. The process must terminate after finitely many steps because the number of edges along which no current is flowing increases at each step. Since ι was arbitrary, ι_0 is indeed a global minimum.

Take $\iota = \iota_0$. Then we need to produce a voltage function ϕ for which

$$\phi(x) - \phi(y) = \rho_e(\iota(e))$$

whenever e is an edge from x to y. This is possible exactly when

$$\sum_i \rho_{e_i}(\iota(e_i)) = 0 \tag{2}$$

for every directed cycle $e_1, e_2, \ldots, e_n$. To see that Equation (2) holds, take our cycle $e_1, e_2, \ldots, e_n$, and let ψ be the current function on Γ which takes the value 1 along each of the e_i, -1 along each $\overline{e_i}$, and 0 elsewhere. Then ψ is a current function with boundary current zero, so $\iota + t\psi \in W$ for every $t \in \mathbb{R}$. Then the fact that ι is a global minimum of $Q(\cdot)$ on W implies that

$$\frac{\partial Q(\iota + t\psi)}{\partial t}$$

vanishes at $t = 0$. But by an easy direct calculation, this gives Equation (2).

Next we show that ϕ is essentially uniquely determined. Let ϕ_1 and ϕ_2 be two solutions, respectively compatible with ι_1 and ι_2. As in the previous section, we can shift the resistance functions and produce a new network which has $\phi_1 - \phi_2$ as a voltage function compatible with the current function $\iota_1 - \iota_2$. So without loss of generality, f is identically zero, and we need to show that if ϕ and ι are compatible voltage and current functions with boundary current zero, then ϕ is constant on each connected component.

Similar to the previous theorem, we take a subgraph containing all edges e for which

$$\phi(\text{start}(e)) > \phi(\text{end}(e))$$

Since voltage is given as a monotone function of current, this implies that $\iota(e) > 0$ too. As before, there can then be no cycles and no internal sources and sinks. But in fact, since the boundary current is also zero, there can be no sources or sinks on the boundary either. But any finite directed acyclic graph without sources or sinks must be empty. Therefore, for every edge e we must have

$$\phi(\text{start}(e)) \leq \phi(\text{end}(e))$$

But then applying this to $\overline{e}$, we see that

$$\phi(\text{start}(e)) \geq \phi(\text{start}(e))$$

too. Thus ϕ must be constant on each connected component. □

Let A be the vector space of real-valued functions f on $\partial\Gamma$ with the property that for every connected component C of Γ,

$$\sum_{v \in C \cap \partial\Gamma} f(v) = 0$$

Then because we can add a separate constant of integration on each connected component to a voltage function, it follows that there is a well-defined Neumann-to-Dirichlet $A \to A$ from boundary voltage functions in A to boundary current functions in A. We could have similarly restricted the Dirichlet-to-Neumann map to a map from $A \to A$, without losing any information, since it is clear that changing all the boundary voltages on a connected component by a constant does not change boundary currents.

In the case of a conductance or resistance network where the conductance or resistance functions are *bijections*, it follows by combining this theorem and the previous one that we have a well-defined bijection between boundary voltage functions in A and boundary current functions in A. In other words, the Dirichlet-to-Neumann map and the Neumann-to-Dirichlet map are essentially inverses of each other. Also, boundary voltages uniquely determine interior voltages in this case, so the solution to the Dirichlet problem is unique.

4 The Inverse Problem

Given the Dirichlet-to-Neumann map or the Neumann-to-Dirichlet map, can we determine the resistance or conductance functions? Without putting additional restrictions on our conductance or resistance functions, the answer is no in almost all cases, because of the following problem. Consider the Y-shaped network of Figure 1. If we put conductance functions of $\arctan x$ on two of the three edges, then the total current through the third edge can never exceed π. Consequently, if we put a conductance function of $\gamma_3(x) = x$ on the third edge, this yields the same Dirichlet-to-Neumann map as the alternative conductance function

$$\tilde{\gamma}_3(x) = \begin{cases} x & \text{\textit{if} } -\pi \le x \le \pi \\ \pi & \text{\textit{if} } x \ge \pi \\ -\pi & \text{\textit{if} } x \le -\pi \end{cases}$$

Similar arguments work as long as any interior vertices exist. In the dual case where voltage is given as a function of current, a similar problem occurs in any graph that contains cycles.

To resolve this difficulty, we instead consider networks where the conductance functions or resistance functions are bijections:

Definition 4.1. *A* bijective network *is a graph* Γ *with boundary, and a function that assigns to each directed edge* e *in* Γ*, a* conductance function $\gamma_e : \mathbb{R} \to \mathbb{R}$*, which satisfies the following properties:*

- $\gamma_e(0) = 0$

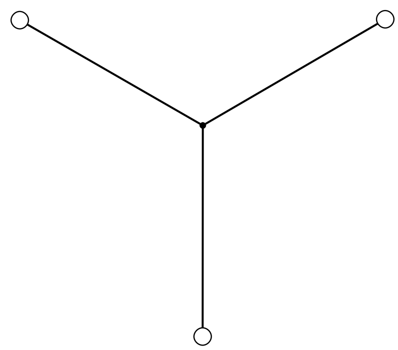

Figure 1: A graph for which the general recovery problem is impossible. The dark middle vertex is an interior vertex, while the light vertices are boundary vertices.

- *γ_e is a bijection for every e*
- *$\gamma_e(-x) = -\gamma_{\overline{e}}(x)$*

We say that the network if monotone *if every γ_e is order-preserving ($x \geq y \Rightarrow \gamma_e(x) \geq \gamma_e(y)$), and we say that the network is* linear *if every γ_e is linear (of the form $\gamma_e(x) = c_e x$ for some $c_e \in \mathbb{R}$).*

For a fixed graph Γ with boundary, a bijective conductance structure *on Γ is a choice of the functions γ_e.*

The definition of compatible current and voltage functions and boundary values is the same as for conductance networks.

Note that we do not require the γ_e to be continuous. In the case of monotone bijective networks, however, the γ_e are all continuous, so by combining Theorems 2.4 and Theorems 3.3 we have well-defined Dirichlet-to-Neumann and Neumann-to-Dirichlet maps. In the general case of bijective networks, current may not be uniquely determined by voltage, and vice versa, so we use a relation instead of a function:

Definition 4.2. *If Γ is a bijective network with n boundary nodes, then the* Dirichlet-Neumann relation *is the relation $\Lambda \subseteq \mathbb{R}^n \times \mathbb{R}^n$ consisting of all pairs (u, f), where u is the boundary voltage of a voltage function ϕ, f is the boundary current of a current function ι, and ϕ and ι are compatible.*

The inverse problem is then to determine all the conductance functions of a bijective network in terms of the Dirichlet-Neumann relation and the underlying graph with boundary.

Definition 4.3. *A graph Γ with boundary is* strongly recoverable *if no two distinct bijective network structures on Γ have the same Dirichlet-Neumann relation. Γ is* weakly recoverable *if no two distinct monotone linear bijective network structures on Γ have the same Dirichlet-Neumann relation.*

In other words, a graph is strongly recoverable if the inverse problem has a solution, in the sense that the conductance functions are uniquely determined by the Dirichlet-Neumann relation. A graph is weakly recoverable if this holds with the assumption

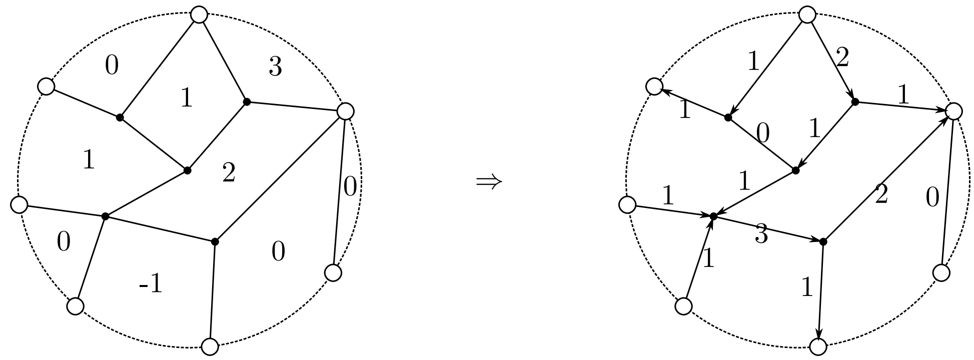

Figure 2: Covoltages on the left, and the corresponding currents on the right. Arrows point in the direction of current flow.

that the conductance functions are monotone linear. Weak recoverability is the type of recoverability considered in [1].

Definition 4.4. *Let* Γ *be a graph with boundary. We say that* Γ *is* circular planar *if it can be embedded in a disk* D *with all boundary vertices on* ∂D *and all interior vertices in the interior of* D*.*

Equivalently, Γ is circular planar if the graph obtained by adding an extra vertex v and connecting v to every vertex in $\partial\Gamma$ is planar.

Generalizing results of [1], we will determine exactly which circular planar graphs are recoverable, showing that for circular planar graphs, weak recoverability and strong recoverability are equivalent.

5 Covoltages

Working with voltages and currents is somewhat asymmetric because voltage functions live on vertices and are unconstrained, while current functions live on edges and are constrained by Kirchhoff's Current Law. To make matters more symmetric, it will help to work with *covoltages* rather than currents.

Definition 5.1. *If* Γ *is a circular planar graph with boundary (with a fixed embedding), a* covoltage function *on* Γ *is an arbitrary real-valued function on the faces of* Γ*. If* ψ *is a covoltage function, the* associated current function *is the current function which assigns a value of* $\psi(f_1) - \psi(f_2)$ *to a directed edge with face* f_1 *on its left and* f_2 *on its right.*

See Figure 2 for an example. We summarize the key facts about covoltage functions in the following lemma:

Lemma 5.2.

- *The associated current function of a covoltage function is a current function (it satisfies KCL).*

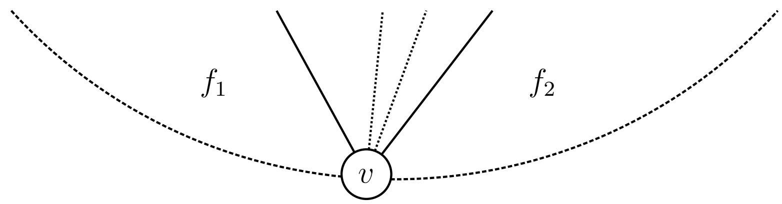

Figure 3: A boundary vertex v and the two faces f_1 and f_2 on either side.

- *If ψ is a covoltage function, with associated current function ι, then the boundary current of ι at a boundary vertex v is $\psi(f_1) - \psi(f_2)$, where f_1 and f_2 are the faces on the left and right sides of v (from v facing into the graph). See Figure 3*
- *A covoltage function is determined by its associated current function up to a constant of integration. In other words, if two covoltage functions have the same associated current function, then their difference is constant.*
- *Every current function is associated to a one covoltage function.*

Only the last of these is nontrivial. It follows from well-known facts about planar graph duality, and we leave its proof as an exercise to the reader.

From this lemma we see that covoltage functions and current functions are essentially equivalent. We will generally identify them, saying that a voltage function ϕ is compatible with a covoltage function ψ if ϕ is compatible with the associated current function of ψ.

More explicitly, a voltage function ϕ is compatible with a covoltage function ψ if at every edge e from vertex v_1 to v_2, with face f_1 on the left and f_2 on the right,

$$\phi(v_1) - \phi(v_2) = \gamma_e(\psi(f_1) - \psi(f_2)).$$

From the lemma, we also see that boundary covoltages determine boundary currents, and conversely boundary currents determine boundary covolages up to an additive constant. Consequently, the Neumann-Dirichlet relation contains the same information as the *voltage-covoltage relation* Ξ which can be defined by

$$\Xi = \{(\phi|_{\partial V}, \psi|_{\partial F}) : \ \phi \text{ is a voltage function and } \psi \text{ is a compatible covoltage function}\}$$

where ∂V and ∂F are the sets of vertices and faces of Γ along the boundary.

The relation between Ξ and Λ can be described explicitly as follows: let $v_1, f_1, v_2, f_2, \ldots, v_n, f_n$ be the boundary vertices and faces of Γ, in counterclockwise order around the boundary of Γ. Then we have

$$(g, h) \in \Xi \iff (g(v_1), g(v_2), \ldots, g(v_n), h(f_n) - h(f_1), h(f_1) - h(f_2), \ldots, h(f_{n-1}) - h(f_n)) \in \Lambda,$$

and conversely, $(x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n) \in \Lambda$ iff $y_1 + y_2 + \cdots + y_n = 0$ and $(g, h) \in \Xi$ where $g(v_i) = x_i$ and h is any function on the set of boundary faces with $h(f_n) - h(f_1) = y_1$, $h(f_1) - h(f_2) = y_2$, and so on. Such a function h exists because $y_1 + y_2 + \cdots + y_n = 0$, and the choice of h doesn't matter because Ξ has the property that $(g, h) \in \Xi$ iff $(g, h + \kappa) \in Xi$.

Because the voltage-covoltage relation contains the same information as the Neumann-Dirichlet relation, we can replace the Neumann-Dirichlet relation with the voltage-covoltage relation in the definition of weak and strong recoverability.

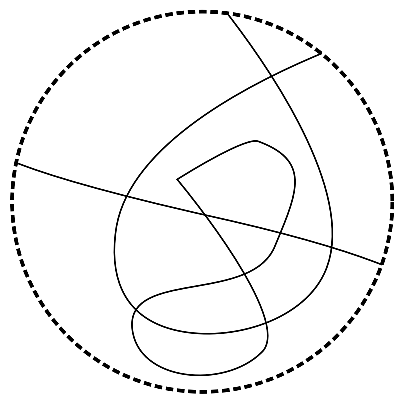

Figure 4: A sample medial graph. This one is rather degenerate; one of the geodesics closes back up on itself.

6 Medial Graphs

A graph's recoverability can be determined by considering an associated *medial graph*:

Definition 6.1. *A* medial graph *is a piecewise-smooth Jordan curve $C \subseteq \mathbb{R}^2$, called the* boundary curve *and a collection of piecewise-smooth curves (*geodesics*) in the inside of C, each of which is either closed, or has both endpoints on C, satisfying the following properties:*

- *Every intersection of a geodesic with C is transversal.*
- *Every intersection between two geodesics or between a geodesic and itself is transversal.*
- *There are no triple intersection points of geodesics or the boundary curve.*

The geodesics partition the inside of the boundary curve into connected components called cells*. Those along the boundary curve are called the* boundary cells*. The* interior vertices *of a medial graph are all the points of intersection between geodesics (including self-intersections). The* boundary vertices *are the intersections of the geodesics with the boundary curve.*

We draw the boundary curve with a dashed line. The medial graph of Figure 4 has three geodesics, eleven cells, eight interior vertices, and four boundary vertices.

To any circular planar graph Γ (with a fixed embedding in a disk D), we can construct a medial graph as follows: place a vertex along each edge of Γ, and connect two vertices if they are on consecutive edges around one of the faces of Γ. Also place two vertices along each arc of ∂D between consecutive boundary nodes of Γ, and connect these to the vertices on the closest edges. Figure 5 shows this whole process. The resulting vertices and edges can then be decomposed as a union of geodesics in a unique way. The result is a medial graph with at least $|\partial\Gamma|$ geodesics, and boundary curve ∂D.

Definition 6.2. *A* coloring *of a medial graph is a function from the cells to the set $\{black, white\}$ such that directly adjacent cells have opposite colors. Here we say that two cells are directly adjacent if they share an edge in common.*

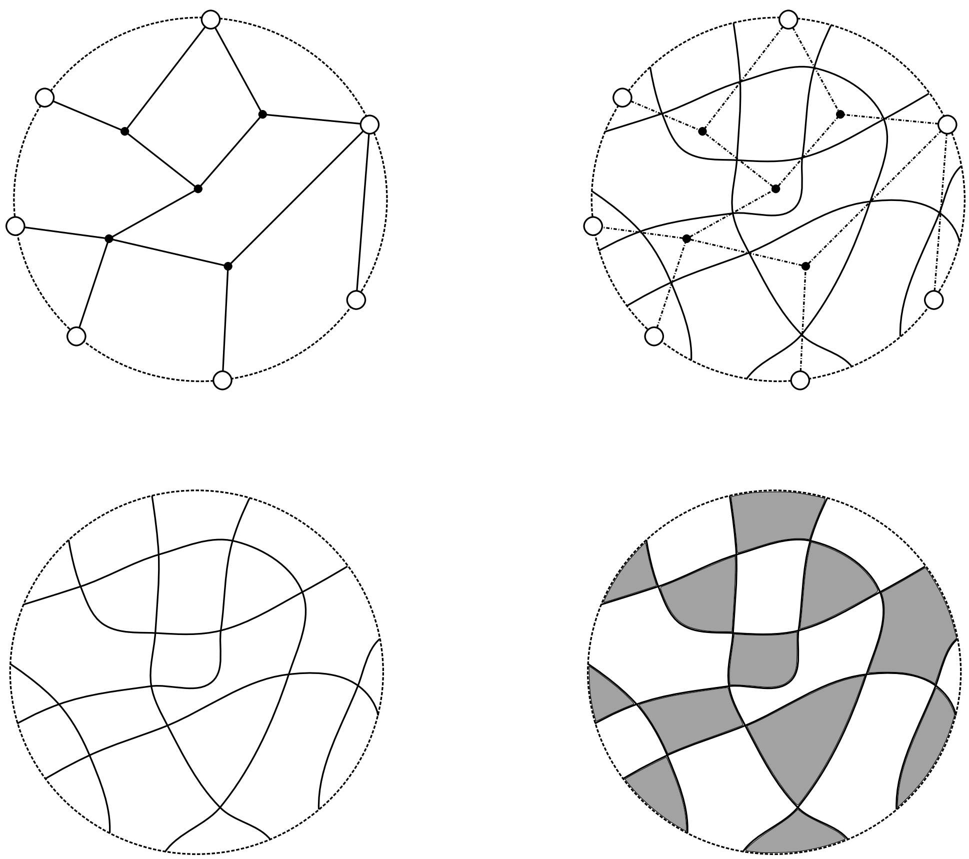

Figure 5: The process of converting a circular planar graph (top left) into a medial graph (bottom left). The natural coloring of this medial graph is shown on the bottom right.

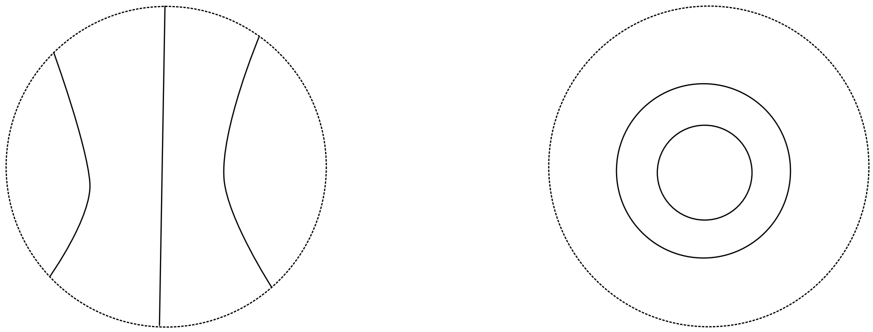

Figure 6: Medial graphs that cannot come from circular planar graphs (regardless of which way they are colored).

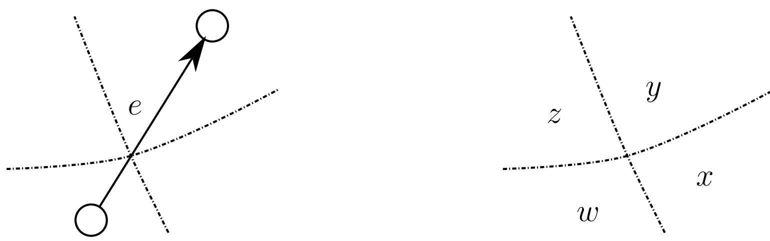

Figure 7: The four medial-graph cells surrounding an edge of the original graph.

Every medial graph has exactly two colorings. A medial graph that comes from a circular planar graph Γ has a canonical coloring, with the cells that come from vertices colored black, and the cells that come from faces colored white. The process of converting a circular planar graph into a colored medial graph is almost reversible. In particular, a circular planar graph Γ is determined (up to isotopy) by its associated colored medial graph. However, some degenerate medial graphs, like those in Figure 6, do not come from circular planar graphs.

Colored medial graphs therefore generalize circular planar graphs, and we can generalize the inverse boundary value problems to colored medial graphs in the following way. If M is the medial graph of Γ, then the black cells of M correspond to the vertices of Γ, while the white cells of M correspond to faces of Γ. So a voltage function is equivalent to a function on the black cells of M, while a covoltage function is equivalent to a function on the white cells of M. The compatibility condition then says that if w, x, y, z are four cells that meet at an interior vertex corresponding to a directed edge e, then

$$\psi(z) - \psi(x) = \gamma_e(\phi(w) - \phi(y))$$

if w, x, y, z are in counterclockwise order around e, and w is the cell at the start of e, as in Figure 7. From the medial graph point of view, we might as well combine ψ and ϕ into one function.

Remark 6.3. *To avoid a definitional nightmare, we will henceforth pretend that conductance functions γ_e are always odd. Equivalently, we pretend that $\gamma_e = \gamma_{\overline{e}}$. None of the following results will use this assumption, but they are much harder to state* *correctly without it! For example, in the next definition, we would need to define something like an "oriented interior vertex" which is tricky to define correctly in the case where geodesics loop back on themselves.*

Definition 6.4. *Let M be a colored medial graph. A* conductance structure *on M assigns to every interior vertex v an odd bijection $\gamma_v : \mathbb{R} \to \mathbb{R}$. A conductance structre is* monotone *if every γ_v is order-preserving, and* linear *if every γ_v is linear.*

When M comes from a circular planar graph Γ, a conductance structure on M is equivalent to a bijective conductance structure on Γ, and the conditions of monotonicity and linearity are equivalent for M and for Γ.

Definition 6.5. *Let M be a colored medial graph and γ be a conductance structure on M. A* labelling *of (M, γ) is a function ϕ from the cells of M to $\mathbb{R}$ such that for every interior vertex v in M,*

$$\phi(z) - \phi(x) = \gamma_v(\phi(w) - \phi(y)) \tag{3}$$

where w, x, y, z are the cells in counterclockwise order around v, with w and y colored black. The boundary data *of a labelling ϕ is the restriction $\phi|_\partial$ of ϕ to the boundary cells. The* boundary relation *Ξ_γ of γ is the set*

$$\Xi_\gamma = \{\phi|_\partial \ : \ \phi \textit{ is a labelling of } (M, \gamma)\}$$

of all possible boundary data of labellings of (M, γ). A colored medial graph M is strongly recoverable *if the function $\gamma \mapsto \Xi_\gamma$ is injective, and M is* weakly recoverable *if it is injective when restricted to monotone linear γ.*

Note that strong recoverability is a stronger condition then weak recoverability. In the case where M comes from a circular planar graph, it is not difficult to see that Ξ_γ contains the same information as the voltage-covoltage relation of (Γ, γ), and in particular, Γ is weakly (strongly) recoverable if and only if M is.

It turns out that M is weakly (strongly) recoverable if and only if M' is, where M' is the colored medial graph with the opposite colors. This is slightly easier to see for the case of strong recoverability. In particular, recoverability of a medial graph doesn't depend on the choice of the coloring. We won't use this fact in what follows, so we leave the (easy) proof to the reader.

Definition 6.6. *A medial graph is* semicritical *if the following hold:*

- *No geodesic intersects itself*
- *If g_1 and g_2 are geodesics, then g_1 and g_2 intersect at most once.*

If in addition, every geodesic starts and ends on the boundary curve (rather than being closed), we say that the medial graph is critical.

Critical medial graphs are essentially equivalent to *simple pseudoline arrangements.*

The main result we are working towards is the following:

Theorem 6.7. *Let M be a medial graph. Then the following are equivalent:*

- *M is strongly recoverable.*
- *M is weakly recoverable.*
- *M is semicritical.*

In particular, this gives a criterion for determining whether a circular planar graph is weakly or strongly recoverable.

The equivalence of weak recoverability and semicriticality for medial graphs coming from circular planar graphs was shown in [1] by arguments involving linear algebra and determinantal identities. It therefore comes as a surprise that this remains true in the highly nonlinear cases considered here. It is also interesting that for circular planar graphs, weak and strong recoverability are equivalent. It is unclear whether this remains true for non-planar graphs.

7 Lenses, Boundary Triangles, and Motions

In this section, we review facts about medial graphs from [1], and show why a weakly recoverable medial graph is semicritical. Almost all of the proofs of this section are based on those in [1]. We use unusual terminology, however.

Definition 7.1. *A* boundary digon *or* boundary triangle *is a simply-connected boundary cell with two or three sides, respectively. If g is a geodesic which does not intersect the boundary curve or any other geodesics and which is a simple closed curve, then we call the set of cells inside g a* circle. *If there is just one cell, we call it an* empty circle*.*

Theorem 7.2. *Let M be a critical medial graph. If M has more than one cell, then M has at least one boundary digon, or at least three boundary triangles.*

Proof.

Definition 7.3. *If M is a critical medial graph and g is a geodesic, then since M is critical, g divides M into two connected components, by the Jordan curve theorem. We call such a connected component a* side*. If g and h are two geodesics, then $g \cup h$ divides M into four pieces, which we call* wedges*. The boundary of any wedge consists of a segment of the boundary curve, which we call the* base*, and two segments of g and h, which we call the* legs*. We call the intersection of g and h the* apex *of a wedge. We say that a wedge W is a* semiboundary wedge *if at least one of the two legs contains no interior vertices other than the apex. A* boundary zone *is one of the following:*

- *A semiboundary wedge*
- *The side of a geodesic which intersects no other geodesics.*

We first prove the following:

Lemma 7.4. *Every boundary zone contains a boundary triangle or a boundary digon.*

Proof. Let Z be a boundary zone. Without loss of generality, Z contains no smaller boundary zones. I claim that no geodesic crosses the boundary of Z. This is trivial if Z is the side of a geodesic which intersects no other geodesics. In the other case, Z is a semiboundary wedge. Say Z has apex a and legs ab and ac. Without loss of generality, no geodesics cross ab, and it remains to show that no geodesics cross ac. Suppose a geodesic g crosses ac at an interior vertex d. Without loss of generality, d is as close to c as possible, and in particular, there are no interior vertices between c and d. After g enters Z at d, it must do one of three things:

- Exit Z by passing through ab.
- Exit Z by passing through ac.
- End on the boundary curve.

The first of these is impossible because there are no interior vertices on ab. The second is impossible because it would mean that g crosses ac in two places. The third is impossible because it would mean that there was a smaller semiboundary wedge inside Z, as in Figure 8, contradicting the choice of Z.

So no geodesic enters or exits Z. I claim that there are in fact no geodesics in Z. Let g be any geodesic in Z. If g intersects no other geodesics, then one side of g is a smaller boundary zone inside Z, a contradiction. Otherwise, there is at least one interior vertex along g. Let e be one which is as close to the boundary curve as possible, as in Figure 9. Then e is the apex of at least two semiboundary wedges contained inside Z, a contradiction.

Therefore, Z contains no geodesics in its interior. It follows that Z must be a single cell. Therefore Z is either a boundary triangle or a boundary digon. □

We next, show the existence of at least one boundary triangle or boundary digon. By the lemma, it suffices to find at least one boundary zone. Since M has more than one cell, it has more than zero geodesics. Let g be a geodesic. If g intersects no other geodesics, then both sides of g are boundary zones, so we have at least one boundary digon or boundary triangle. Otherwise, let v be a boundary vertex of g which is as close as possible to one of the endpoints of g. If h crosses g at v, then half the wedges formed by g and h are semiboundary wedges, so at least one boundary digon or boundary triangle exists.

If a boundary digon exists, then we are done. Otherwise, we have a boundary triangle. This boundary triangle is in fact a wedge formed by two geodesics g and h. Two of the other wedges formed by g and h are semiboundary wedges, so this gives two more boundary triangles or boundary digons. □

Definition 7.5. *A medial graph is* circleless *if it has no empty circles.*

Theorem 7.6. *A semicritical circleless graph M is critical.*

Proof. Given the definition of critical and semicritical, it suffices to show that if M is circleless, then it has no circles. Let C be a circle which minimizes the number of cells inside C. Since all the intersections of geodesics are transversal, if g is any geodesic, then g intersects (the boundary of) C an even number of times. This can only happen in a semicritical graph if g intersects (the boundary of) C exactly zero times. It follows that every geodesic in M is either inside C or outside C. But every geodesic inside C must itself be a circle, contradicting the choice of C. Therefore C is an empty circle, and M is not circleless. □

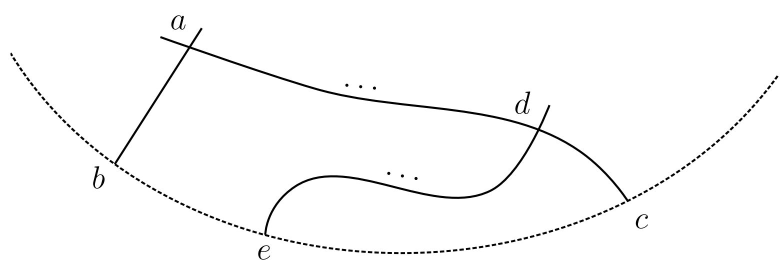

Figure 8: If there are any interior vertices along leg ac, then abc is not a minimal semiboundary wedge.

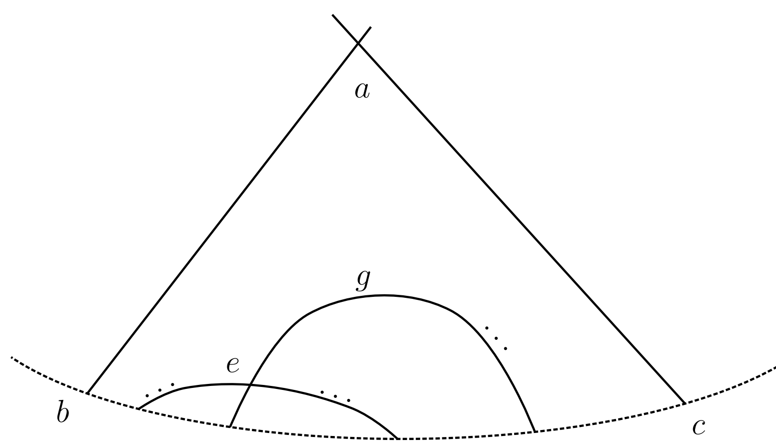

Figure 9: If any interior vertex e exists along g, then e is the apex of at least two semiboundary wedges inside abc.

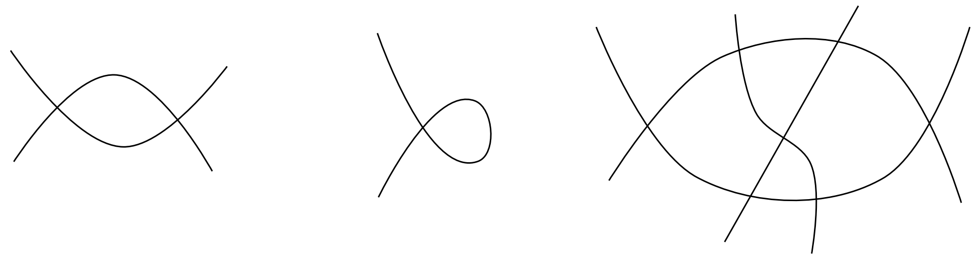

Figure 10: An empty 2-lens, an empty 1-lens, and a nonempty sharp 2-lens. TODO: better examples.

Definition 7.7. *If g is a geodesic, we can think of g as either a continuous function on $[0,1]$ or a continuous function on $\mathbb{R}$ with period 1. If $[a,b]$ is a subinterval of the domain of g, and g is injective on both half-open intervals $(a,b]$ and $[a,b)$, we call the curve $g([a,b])$ a* geodesic segment *with endpoints $g(a)$ and $g(b)$.*

Definition 7.8. *Let v be an interior vertex. A* 1-lens with pole at v *is the interior of a Jordan curve that is a geodesic segment with endpoints v and v. Let v and w be interior vertices. A* 2-lens with poles at v and w *is the interior of a Jordan curve that is a union of two geodesic segments, each having endpoints at v and w. For $n = 1, 2$, we say that an n-lens is* sharp *if for every pole v, only oneof the four cells that meet at v is inside the n-lens. We say that an n-lens is* empty *if it contains only one cell.*

Because every medial graph can be colored, no cell can be on both sides of a geodesic. From this it is easy to see that empty lenses must be sharp lenses. See Figure 10 for examples of lenses.

Theorem 7.9. *If M is a medial graph which is not semicritical, then M must contain a sharp 1-lens or sharp 2-lens.*

Proof. Since M is not semicritical, we either have a geodesic which self-intersects, or two geodesics which intersect too many times. First suppose that some geodesic g intersects itself. Then we have $g(s) = g(t)$ for some $s \neq t$, and in the case where g is a closed curve and we think of g as a function on $\mathbb{R}/\mathbb{Z}$, we want $s - b \notin \mathbb{Z}$. Choose a pair $\{s,t\}$ for which all this holds, minimizing $|s-t|$. Without loss of generality $s < t$. I claim that $g([s,t])$ is a geodesic segment. In the periodic case, clearly $|s-t| < 1$, or else we could replace s with $s+1$ and make $|s-t|$ smaller. If g fails to be injective on $(s,t]$, then there exist $s < r_1 < r_2 \leq t$ with $g(r_1) = g(r_2)$ and $|r_1 - r_2| < |s-t|$ (and $r_1 - r_2 \notin \mathbb{Z}$ in the periodic case), contradicting the choice of s, t. A similar argument shows that g is injective on $[s,t)$. So $g([s,t])$ is a geodesic segment with endpoints at $g(s)$ and $g(t)$. In particular, it is a 1-lens, with pole at $v = g(s) = g(t)$.

If this 1-lens fails to be sharp, then some geodesic h at v enters into the 1-lens. It eventually exits the 1-lens. If it does so without crossing itself, then we have a smaller 2-lens or 1-lens, as in Figure 11. Otherwise, we can repeat all of the above argument on h (or rather, on h restricted to the part between v and where h first exits the 1-lens), and get a smaller 1-lens, and continue inductively, until we get a sharp 1-lens.

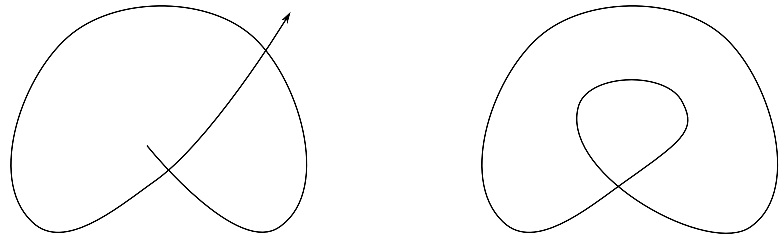

Figure 11: A non-sharp 1-lens must contain a smaller 2-lens or 1-lens. A third possibility, not pictured here, is that the geodesic crosses itself before entering and exiting the original 1-lens. Then one repeats the argument of Theorem 7.9 inductively.

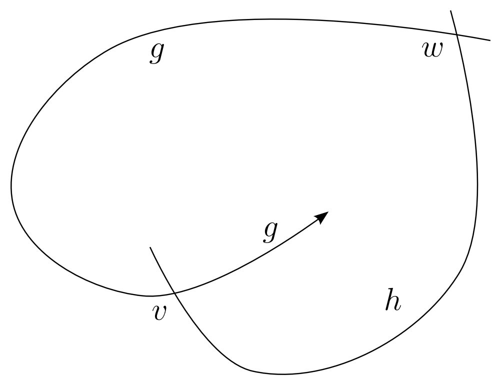

Figure 12: If sharpness fails at pole v, then geodesic g has no means of exiting the 2-lens. (The pole at w needn't be sharp in this arrangement.)

Next, suppose that no geodesic intersects itself. Then M can only fail to be semi-critical if two geodesics g and h intersect more than once. So there exist s, t in the domain of h with $h(s)$ and $h(t)$ in the range of g, and $s \neq t$. As before, we can take $s - t \notin \mathbb{Z}$ in the periodic case. Choose such an $\{s, t\}$ with $|s - t|$ minimized. Then as before, $|s - t| < 1$ in the periodic case. By choice of $\{s, t\}$, $h(r)$ cannot lie on g for any r between s and t. It follows immediately that the geodesic segments between $h(s)$ and $h(t)$ along h and g form a 2-lens, with poles at $v = h(s)$ and $w = h(t)$. To see that this 2-lens is sharp, suppose that more than one of the cells around v lies inside the 2-lens. Since $h \neq g$ and the intersection at v is transversal and the boundary of the 2-lens is a Jordan curve, exactly three of the cells lie inside, and we have the configuration of Figure 12. Then g continues on from v into the interior of the 2-lens. It has no way of escaping, however: it cannot intersect itself (by assumption), and it cannot intersect $h([s, t])$ (by choice of $\{s, t\}$). This is a contradiction. □

Definition 7.10. *A* Y-Δ motion *is the transformation shown in Figure 13.*

If our medial graphs come from circular planar networks, this is nothing but a standard Y-Δ transformation as in Figure 14.

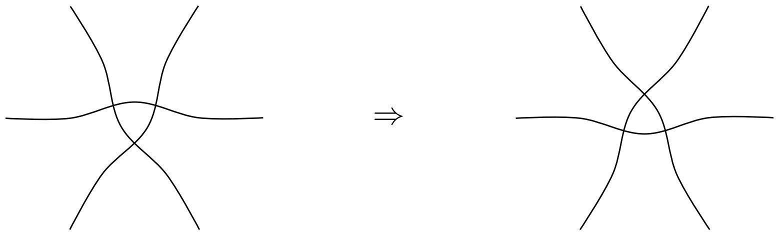

Figure 13: A Y-Δ motion in a medial graph.

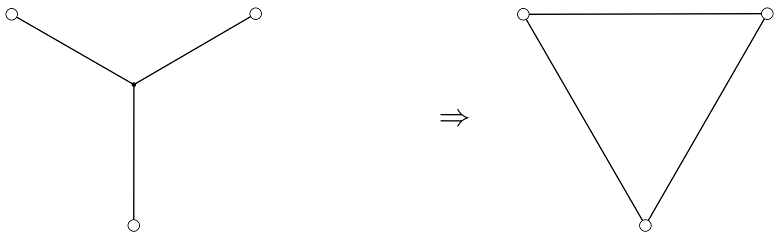

Figure 14: A Y-Δ transformation.

Theorem 7.11. *If M is a medial graph which is not semicritical, then we can apply some combination of the following operations to reach a medial graph which contains an empty 1-lens or 2-lens:*

- *A Y-Δ motion*
- *Removal of an empty circle.*

Proof. By the previous theorem, there is at least one sharp lens. We prove by induction on n that if M is a medial graph containing a sharp lens which contains n cells, then M can be converted into a medial graph which contains an empty lens. The base case where $n = 1$ is trivial, because then M already has an empty lens. Suppose $n > 1$, and we have a sharp lens L with n cells. If there are any smaller sharp lenses inside L, we are done by induction. Suppose that L contains no smaller sharp lenses. Since the boundary of L is a Jordan curve, we can make a new medial graph M' from the inside of L, with the boundary of L as its boundary curve. I claim that M' is semicritical. Otherwise, it contains a sharp lens, by the previous theorem, contradicting the fact that L contains no smaller sharp lenses. If M' contains an empty circle, then we can remove it, and continue by induction. So we can suppose that M' is circleless. Then by Theorem 7.6, M' is critical. Therefore, it contains at least three boundary triangles, or at least one boundary digon. If there are at least three boundary triangles, then at least one of them does not touch a pole of L. Therefore it corresponds to an actual triangle in M, and we can carry out a Y-Δ motion, pushing the intersection outside of L, while preserving the lens L. This decreases the number of cells in L, so we can continue by induction. Finally, suppose that there is a boundary digon. If it does not touch one of the poles of L, it corresponds to an empty lens, and we are done. Otherwise, it corresponds to a triangle in M, and we can push it outside of L by a Y-Δ motion, and then continue by induction. □

In the monotone linear case, a conductance structure on a medial graph is merely a function that assigns a conductance $c_v > 0$ to each interior vertex v of M. The conductance functions are of the form $\gamma_v(x) = c_v x$.

Lemma 7.12. *The following transformations do not effect the boundary relation of a medial graph. We assume that the vertices involved in the transformations have monotone linear conductance functions, but the rest of the medial graph need not.*

- *Removing a boundary digon.*
- *Removing an empty circle.*
- *A motion as in Figure 15. Here, the new conductances should be given by*

$$c_{12} = \frac{c_1 c_2}{c_1 + c_2 + c_3}$$
$$c_{13} = \frac{c_1 c_3}{c_1 + c_2 + c_3}$$
$$c_{23} = \frac{c_2 c_3}{c_1 + c_2 + c_3}$$

  in one direction, and

$$c_1 = \frac{c_1 c_2 + c_1 c_3 + c_2 c_3}{c_{23}}$$
$$c_2 = \frac{c_1 c_2 + c_1 c_3 + c_2 c_3}{c_{13}}$$
$$c_3 = \frac{c_1 c_2 + c_1 c_3 + c_2 c_3}{c_{12}}$$

  in the other direction.

- *Removing an immediate self-loop.*
- *Replacing an emtpy lens with a single intersection, as in Figure 16. In the top row of Figure 16, the conductance of the new vertex should be*

$$\frac{c_1 c_2}{c_1 + c_2}$$

  while in the bottom row it should be

$$c_1 + c_2$$

Proof. These all follow by direct calculation. The third rule is just the standard Y-Δ transformation, while the fifth consists of the standard rules for resistors in series or parallel. □

The $Y - \Delta$ transformation is reversible, while the process of replacing a series or parallel configuration with a single resistor throw away information. This immediately yields the following conclusion:

Corollary 7.13.

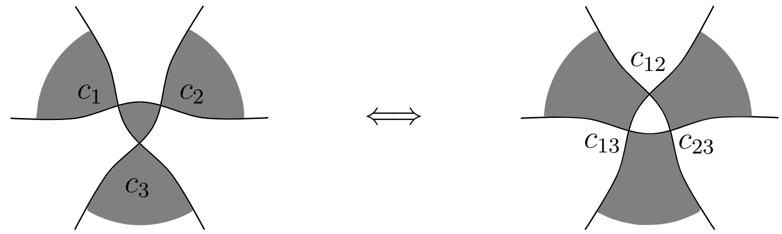

Figure 15: A Y-Δ motion in a colored medial graph.

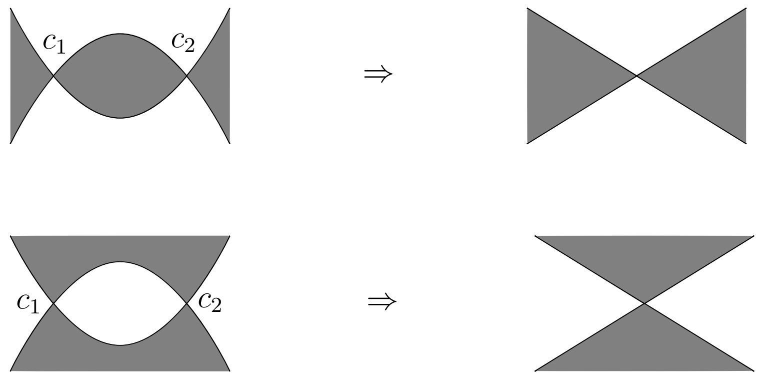

Figure 16: Replacing an empty lens with a single intersection.

- *Adding or removing a boundary digon does not affect weak or strong recoverability.*
- *Adding or removing an empty circle does not affect weak or strong recoverability.*
- *A Y-Δ motion does not affect weak recoverability.*
- *Any medial graph containing an immediate self-loop or an empty lens is not weakly recoverable.*

This yiels half of Theorem 6.7.

Corollary 7.14. *If M is a medial graph and M is not semicritical, then M is not weakly recoverable.*

Proof. Combine Corollary 7.13 and Theorem 7.11. □

For the other direction, we will use the following property of boundary triangles.

Theorem 7.15. *Let M be a medial graph with a boundary triangle abc, and let M' be the medial graph obtained by* uncrossing *the boundary triangle, as in Figure 17. Let a be the apex of abc. If we fix the conductance function γ_a, then the boundary relation of M' is determined by the boundary relation of M.*

Proof. Let w, x, y, z be the four cells around a, as in Figure 17. By (3), if ϕ is any labelling of M', then ϕ extends to a labelling of M uniquely by setting

$$\phi(w) = \phi(y) \pm \gamma_a(\phi(x) - \phi(y))$$

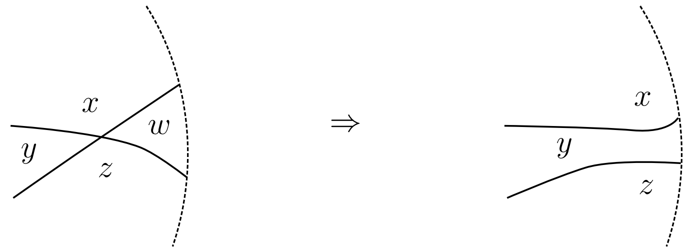

Figure 17: "Uncrossing" a boundary triangle.

where the $\pm$ depends on the coloring. It follows that if Ξ and Ξ' are the boundary relations of M and M', then

$$(\ldots, p, q, r, \ldots) \in \Xi' \iff (\ldots, p, q \pm \gamma_a(p - r), r, \ldots) \in \Xi,$$

so Ξ and Ξ' determine each other. □

In light of this, we can reduce Theorem 6.7 to one condition:

Theorem 7.16. *Suppose the following statement is true: if M is any critical colored medial graph, and abc is a boundary triangle of M, with apex a, then γ_a is determined by the boundary relation of M. Then Theorem 6.7 is true.*

Proof. We have already seen that weak recoverability implies semicriticality. Clearly, strong recoverability implies weak recoverability. It remains to show that semicriticality implies strong recoverability. Let M be a semicritical medial graph. Suppose M is not strongly recoverable, and take M to have as few cells as possible. By Theorem 7.13, we can remove empty circles from M without affecting strong recoverability. Therefore, M is circless, and by Theorem 7.6, it is critical. Similarly, M cannot have a boundary digon, by Theorem 7.13. If M has no geodesics, then M is strongly recoverable because the inverse boundary problem requires the recovery of no information. Otherwise, M must have a boundary triangle by Theorem 7.2. Let M' be the medial graph obtained by uncrossing the boundary triangle. By assumption, we can find the conductance function at the apex of the boundary triangle, and by the previous theorem, this in turn yields the boundary relation of M'. Uncrossing a boundary triangle does not break criticality, so M' is critical, and by choice of M it is therefore strongly recoverable. So given the boundary relation of M', we can find all the remaining conductance functions of M. Therefore M is strongly recoverable, a contradiction. □

It remains to show how to recover boundary the conductance functions of boundary triangles. We will do this by setting up special boundary value problems which allow us to read off the conductance function at a boundary triangle.[1] The main difficulty will be showing that our boundary value problem is well-defined, having a unique solution.

[1]Curtis and Morrow [1] recover the conductances of boundary triangles using linear algebra. Since we allow nonlinear conductances, this approach won't work here.

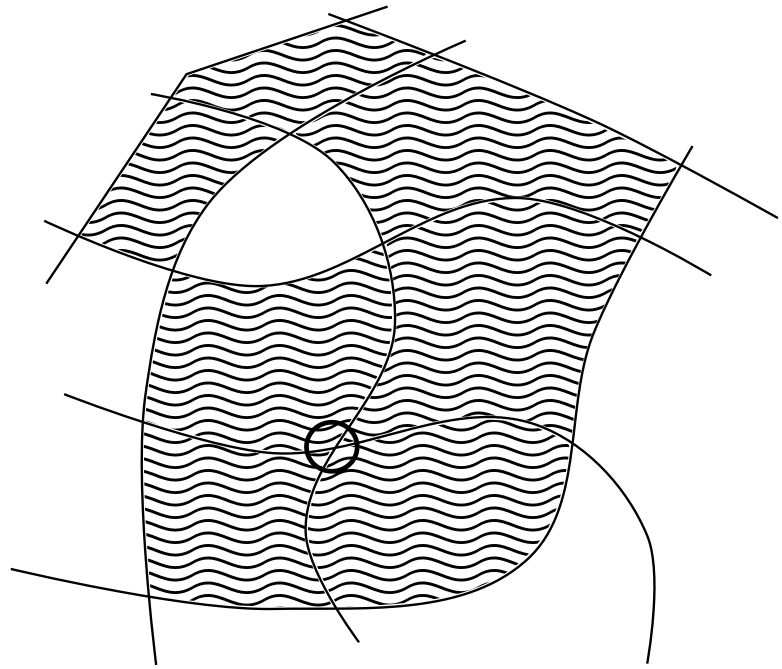

Figure 18: The wavy cellset has rank six, because it contains seven cells, and there is one interior vertex completely surrounded by S (the circled one).

8 The Propagation of Information

To accomplish this, we will need a way of modelling the propagation of information through a medial graph. We begin by defining partial labellings of medial graphs.

Definition 8.1. *If M is a fixed medial graph, a* cellset *is a set of cells in M. If M is the medial graph associated to a bijective network Γ, and S is a cellset, then a* labelling *of S is a function $\phi : S \to \mathbb{R}$ which satisfies (3) whenever $w, x, y, z \in S$.*

Note that the notion of a labelling depends on the conductance functions.

Next we turn to combinatorial definitions that depend solely on M, not even on its coloring. These will be our model for the propagation of information.

Definition 8.2. *Let S be a cellset. Then define* $\text{rank}(S)$ *to be $|S| - v(S)$ where $|S|$ is the number of cells in S, and $v(S)$ is the number of interior vertices for which all four surrounding cells are in S.*

The rank of S estimates the number of degrees of freedom of a labelling of S, since it counts the number of cells (the number of variables) minus the number of constraining equations. See Figure 18 for an example.

Definition 8.3. *If S is a cellset and x is a cell not in S, then we say that $S \cup \{x\}$ is a* simple extension *of S if there is an interior vertex v of M for which three of the cells around v are in S and x is the fourth. If there does not exist another interior vertex v' with the same property, then we say that $S \cup \{x\}$ is a* safe simple extension*.*

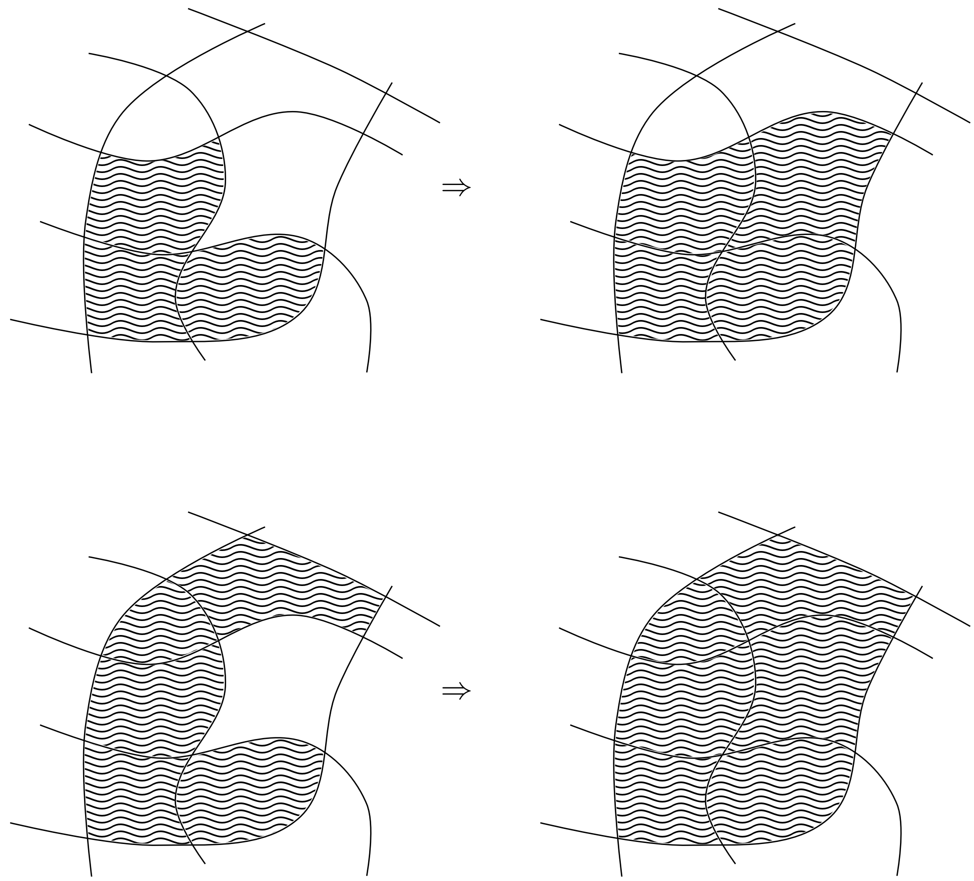

Figure 19: The top right is a safe simple extension of the top left, while the bottom right is a (non-safe) simple extension of the bottom left.

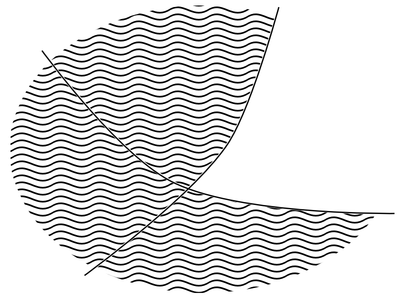

Figure 20: The forbidden configuration in a closed set. A cellset is closed exactly when this configuration never occurs.

Lemma 8.4. *Let S be a cellset and x be a cell not in S. Then* $\mathrm{rank}(S\cup\{x\}) \leq \mathrm{rank}(S)$ *if $S \cup \{x\}$ is a simple extension of S, and* $\mathrm{rank}(S \cup \{x\}) = \mathrm{rank}(S) + 1$ *otherwise. Moreover,* $\mathrm{rank}(S \cup \{x\}) = \mathrm{rank}(S)$ *if and only if $S \cup \{x\}$ is a safe simple extension of S.*

Proof. Obviously $|S \cup \{x\}| = |S| + 1$. So letting $v(T) = \mathrm{rank}(S) - |S|$ as above, we need to show that

- If $S \cup \{x\}$ is not a simple extension of S, then $v(S \cup \{x\}) = v(S)$
- If $S \cup \{x\}$ is a safe simple extension of S, then $v(S \cup \{x\}) = v(S) + 1$.
- If $S \cup \{x\}$ is a simple extension of S that is not safe, then $v(S \cup \{x\}) > v(S) + 1$.

But these are all obvious from the definitions of simple extensions and safe simple extensions. □

Definition 8.5. *A cellset S is* closed *if it has the property that for every interior vertex v of M, if three of the four cells around v are in S, then so is the fourth. The smallest closed set containing S is the* closure *of S, denoted $\overline{S}$.*

The closure $\overline{S}$ exists because the intersection of closed cellsets is closed. Note that S is closed if and only if it has no simple extensions.

Lemma 8.6. *A cellset S is closed if and only if it has no simple extensions. If S' is a simple extension of S, then $\overline{S'} = \overline{S}$. Also, there is a chain of cellsets*

$$S = S_1 \subset S_2 \subset \cdots \subset S_n = \overline{S}$$

where S_{i+1} is a simple extension of S_i for each i. Moreover $\mathrm{rank}(\overline{S}) \leq \mathrm{rank}(S)$, *with equality if and only if S_{i+1} is a safe simple extension of S_i for each i.*

Proof. The first claim is obvious. The second claim follows because $S \subset S'$ implies that $\overline{S} \subseteq \overline{S'}$, but conversely S' is clearly a subset of $\overline{S}$ and so $\overline{S'} \subseteq \overline{S}$ because $\overline{S'}$ is the smallest closed subset containing $\overline{S}$.

For the third claim, inductively define S_{i+1} to be any simple extension of S_i, stopping at S_n which has no simple extension. Then S_n is closed by the first claim, and $S_n = \overline{S_n} = \overline{S}$ by the second claim.

For the final claim, we know by the previous lemma that $\operatorname{rank}(S_{i+1}) \leq \operatorname{rank}(S_i)$ for every i. Then by transitivity $\operatorname{rank}(\overline{S}) \leq \operatorname{rank}(S)$, and if equality holds then $\operatorname{rank}(S_{i+1}) = \operatorname{rank}(S_i)$ for every i, so that each step in the chain is a safe simple extension. □

The significance of rank and closure is the following:

Theorem 8.7. *Let S be a cellset and ϕ be a labelling of S. If S' is a simple extension of S, then ϕ extends to a labelling of S' in at most one way, and if S' is a safe simple extension of S, then ϕ extends in exactly one way. Consequently, ϕ extends to a labelling of $\overline{S}$ in at most one way, and if* $\operatorname{rank}(\overline{S}) = \operatorname{rank}(S)$*, then it extends in exactly one way.*

Proof. First note that Equation (3) above has the property that any three of $\phi(w), \phi(x), \phi(y)$, and $\phi(z)$ uniquely determine the fourth, because we stipulated that γ_e be bijective.

If $S' = S \cup \{x\}$, then by definition of simple extension, there is some interior vertex v of the medial graph such that x is one of the four cells around v, and the other three are already in S. Then $\phi(x)$ is completely determined by (3) at the edge corresponding to v. Moreover, if $S \cup \{x\}$ is a safe simple extension, then v is the unique vertex which determines $\phi(x)$, so $\phi(x)$ is uniquely determined. Thus a labelling of S extends to a labelling of S' in at most one way, and at least one way if S' is a safe simple extension of S.

The statements about extending ϕ to $\overline{S}$ then follow easily from the previous lemma. □

Definition 8.8. *A cellset S is* safe *if* $\operatorname{rank}(\overline{S}) = \operatorname{rank}(S)$.

We have just shown that if S is safe, then a labelling of S extends to a labelling of $\overline{S}$ in a unique way.

9 Convexity in Medial Graphs

From now on we will assume that our medial graph M is critical, so that each geodesic starts and ends on the boundary, no geodesic intersects itself, and no two geodesics intersect more than once. The results of this section are completely independent of conductances, and are entirely combinatorial in nature. Since critical medial graphs are equivalent to simple pseudoline arrangements, our results can all be viewed as statements about pseudoline arrangements, but we will prefer to work with medial graphs (which contain a boundary curve) to simplify some proofs.

If g is a geodesic in a medial graph, the Jordan curve theorem implies that g divides the medial graph into two parts, which we call *pseudo halfplanes.* We can identify these with cellsets.

Definition 9.1. *A cellset S is* convex *if it is an intersection of zero or more pseudo halfplanes. The* convex closure *of a cellset S is the smallest convex set containing S, i.e., the intersection of all half-planes containing S.*

Since half-planes are closed, so are all convex sets. We will roughly show the converse.

Definition 9.2. *If a and b are two cells in a medial graph, we say that a and b are* adjacent *if they share an edge. A* path *between a and b is a sequence of cells $a = x_0, x_1, \ldots, x_n = b$ for $n \geq 0$, with x_i adjacent to x_{i+1} for each i. The number n is the* length *of the path. A* minimal path *between two cells is a path of minimal length. The* distance between a and b, *denoted* $\mathrm{d}(a, b)$, *is defined to be the length of a minimal path between a and b. A cellset S is* connected *if for every a and b in S, there is a path between a and b containing only cells in S. The* connected components *of S are the maximal connected subsets of S.*

Each cellset is the disjoint union of its connected components. The medial graph as a whole is always connected, so that $\mathrm{d}(a, b)$ is well-defined.

Lemma 9.3. *If S is a connected cellset, then so is $\overline{S}$.*

Proof. Because $\overline{S}$ can be obtained from S by a series of simple extensions (by Lemma 8.6), it suffices to show that a simple extension of a connected cellset is connected, which is obvious. □

Definition 9.4. *Let g be a geodesic, and a, b be cells. We say that g* separates *a and b if a and b are on opposite sides of g. We let $\mathrm{d}'(a, b)$ denote the number of geodesics which separate a and b. If S is a cellset, we say that g* separates cells of *S if there exist $a, b \in S$ on opposite sides of g.*

Clearly, $\mathrm{d}'(a, b) \leq \mathrm{d}(a, b)$.

Theorem 9.5. *Let S be a nonempty convex cellset. Then there is a critical medial graph M' whose cells are the same as the cells in S. In particular, S is connected and the boundary of the closure of the union of the cells in S is a Jordan curve. Additionally, the geodesics in M' are in one to one correspondence with the geodesics of M which separate cells of S.*

Proof. If H is a pseudo halfplane, then we can make a medial graph out of the cells in H, as shown in Figure 21. This new medial graph will still be critical. Moreover, if H' is any other pseudo halfplane, then $H' \cap H$ is a pseudo halfplane in the new medial graph, by criticality. We show by induction on r that if M is a critical medial graph and S is the intersection of r pseudo halfplanes in M, then S is the set of cells of a critical medial graph. The base case where $r = 0$ is trivial. Otherwise, write S as $H \cap S'$, where H is a half-plane, and $S' = \bigcap_{i=1}^{r-1} H_i$ for pseudo halfplanes H_i. Then we can consider H as a medial graph M', and $H \cap H_i$ is a pseudo halfplane of M' for each i. By induction, $S = H \cap S' = \bigcap_{i=1}^{r-1}(H \cap H_i)$ is the set of cells of a critical medial graph, and we are done. We leave the second claim as an exercise to the reader. □

TODO: the following corollary is also a lemma in Branko Grünbaum's *Arrangements and Spreads.*

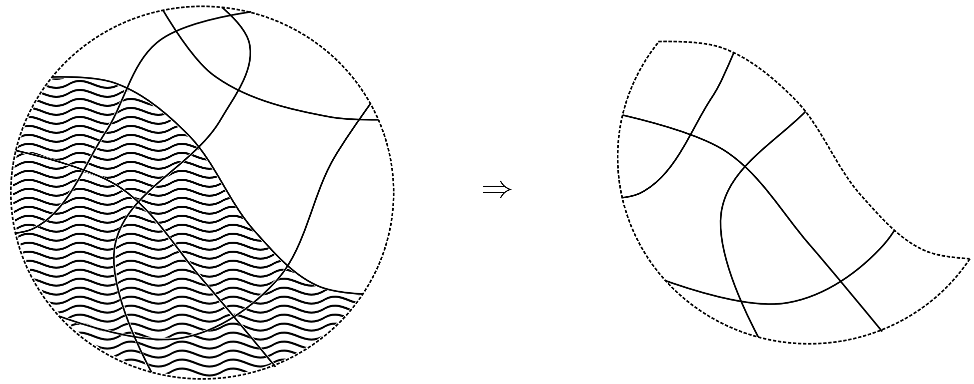

Figure 21: A pseudo halfplane (left), considered as a medial graph in its own right (right).

Corollary 9.6. *If a and b are two cells of a critical medial graph M, then* $\mathrm{d}'(a,b) = \mathrm{d}(a,b)$*. If c is a third cell, then c is on a minimal path between a and b if and only if c is in the convex closure of $\{a,b\}$.*

Proof. We proceed by induction on $\mathrm{d}'(a,b)$. If a and b are the same cell, then clearly $\mathrm{d}(a,b) = \mathrm{d}'(a,b) = 0$. Assume therefore that $a \neq b$. Let S be the convex closure of $\{a,b\}$. By the theorem, S is connected. Since S is connected, there is a path from a to b in S, and in particular there is some cell $a' \in S$ which is adjacent to a. Let g be the geodesic which separates a' from a. Then g also separates a from b. Otherwise, the pseudo halfplane H on the side of g containing a and b would not contain a', so $a' \notin S$. Since g separates a from b and a from a', it follows that $\mathrm{d}'(a',b) = \mathrm{d}'(a,b) - 1$. By induction, $\mathrm{d}'(a,b) = \mathrm{d}'(a',b) + 1 = \mathrm{d}(a',b) + 1$. Therefore,

$$\mathrm{d}'(a,b) \leq \mathrm{d}(a,b) \leq \mathrm{d}(a,a') + \mathrm{d}(a',b) = \mathrm{d}(a',b) + 1 = \mathrm{d}'(a,b),$$

so the desired equality holds.

For the second claim, note that c is on a minimal path between a and b if and only if $\mathrm{d}(a,b) = \mathrm{d}(a,c) + \mathrm{d}(c,b)$. Define $\sigma_g(x,y)$ for cells x, y and geodesic g to be 1 if g separates a from b, and 0 otherwise. Then the first claim of this corollary implies that $\mathrm{d}(x,y) = \sum_g \sigma_g(x,y)$. Now any geodesic which separates a from b must separate a from c or c from b, so

$$\sigma_g(a,b) \leq \sigma_g(a,c) + \sigma_g(c,b).$$

Summing over g, we see that $\mathrm{d}(a,b) \leq \mathrm{d}(a,c) + \mathrm{d}(c,b)$, with equality if and only if $\sigma_g(a,b) = \sigma_g(a,c) + \sigma_g(c,b)$ for every geodesic g. Thus c is on a minimal path from a to b if and only if for every geodesic g, either g does not separate any of a, b, c, or g separates a from b. Equivalently, c is on a minimal path from a to b if and only if c is in every half plane which contains both a and c. But this is just the convex closure of $\{a,b\}$. □

Corollary 9.7. *Every cell in a critical medial graph has a boundary that is a Jordan curve.*

Proof. Let a be a cell, and let S be the convex closure of $\{a\}$. If b is any adjacent cell to a, then some geodesic g separates a from b. Then there is a pseudo halfplane H

which contains b but not a, so that $b \notin S$. Thus S contains no cells adjacent to a. But by Theorem 9.5, S is connected, so $S = \{a\}$. It then follows that a itself must have the shape of a medial graph, and in particular its boundary must be a piecewise smooth Jordan curve. □

Definition 9.8. *Let $(x_1, x_2, \ldots, x_n)$ and $(y_1, y_2, \ldots, y_n)$ be two paths from a to b. We say that $(y_1, \ldots, y_n)$ is a* simple deformation *of $(x_1, \ldots, x_n)$ if there is some $1 < i < n$ such that $x_j = y_j$ for $j \neq i$, and x_{i-1}, x_{i+1}, x_i and y_i are the four cells surrounding some internal vertex. We say that $(y_1, y_2, \ldots, y_n)$ is a* deformation *of $(x_1, \ldots, x_n)$ if it is obtained by a series of zero or more simple deformations.*

Theorem 9.9. *If a, b are cells, then all minimal paths between a and b are deformations of each other.*

Proof. We proceed by induction on $\mathrm{d}(a, b)$. The base cases where $\mathrm{d}(a, b) = 0$ or 1 are obvious, so assume that $\mathrm{d}(a, b) > 1$.

By Corollary 9.6, the union of all the minimal paths between a and b are the convex closure of $\{a, b\}$, which by Theorem 9.5 is a medial graph itself. So we can assume without loss of generality that the convex closure of $\{a, b\}$ is everything, and so every geodesic separates a from b.

Let S be the set of all cells immediately adjacent to a. We claim that we can enumerate the elements of S as $s_1, s_2, \ldots, s_n$ in such a way that for each i, the three cells a, s_i, and s_{i+1} meet at a vertex (with a fourth cell). This setup is shown in Figure 22. This is obvious if a is not a boundary cell (which probably never happens). If a is a boundary cell, then this is still clear, unless a meets the boundary of the medial graph in two or more disjoint boundary arcs, as in Figure 23. But in this case, we can add a new geodesic connecting two of these boundary arcs, and throw away the side of the geodesic that is opposite g, making a smaller medial graph. Any geodesic discarded by this operation could not have separated a from b, and could not have been present. Consequently, the discarded part could not have contained any additional cells or geodesics, implying that the two boundary arcs we connected were in fact the same, a contradiction.

So we can number the neighbors of a as $s_1, s_2, \ldots, s_n$. Any minimal path from a to b must begin by moving from a to one of the s_i. By induction, we know that all minimal paths from s_i to b are equivalent through deformations. So it suffices to show for each i that at least one path from a to b through s_i is a deformation of a path from a to b through s_{i+1}.

To see this, let c be the fourth cell meeting a, s_i, and s_{i+1}, as in Figure 24. If we let g_i and g_{i+1} be the geodesics separating a from s_i and s_{i+1}, respectively, then both g_i and g_{i+1} separate a from b. Consequently, $\mathrm{d}(c, b) = \mathrm{d}(a, b) - 2$. Thus if we take a minimal path from c to b, and add on (a, s_i) or (a, s_{i+1}) to the beginning, we will get minimal paths from a to b. But these two alternatives are deformations of each other, so we are done. □

The following theorem gives several conditions equivalent to convexity.

Theorem 9.10. *Let S be a cellset. The following are equivalent:*

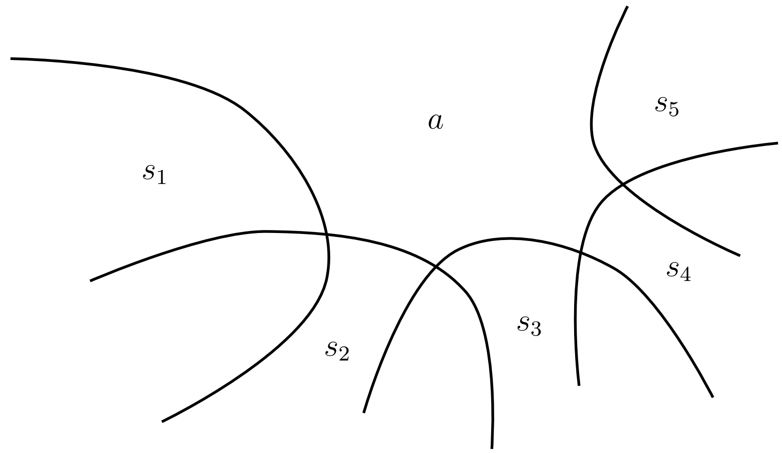

Figure 22: The desired setup in the proof of Theorem 9.9. Here, we can enumerate the neighbors of a in order. In particular, they occur contiguously.

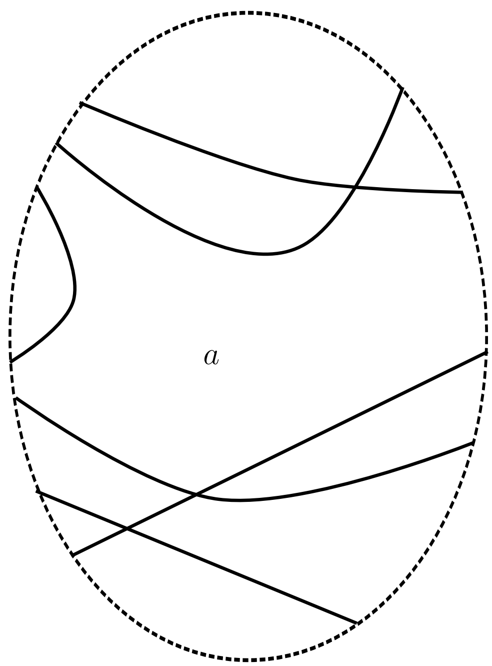

Figure 23: A cell whose neighbors are not contiguous. This only occurs if the cell touches the boundary more than once, as a does here.

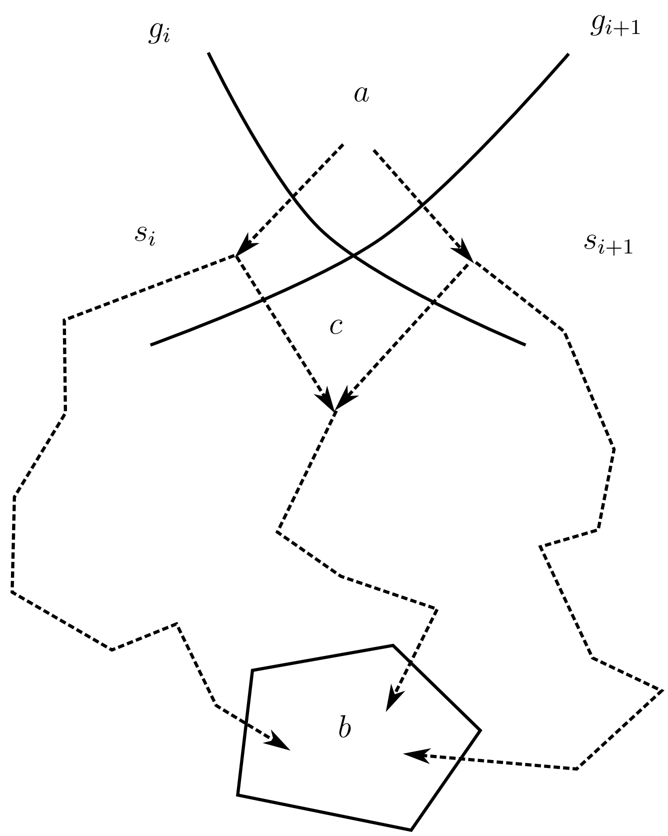

Figure 24: The setup of the rest of the proof of Theorem 9.9.

(a) *S is closed and connected.*

(b) *S is connected, and for every $a, b \in S$, if one minimal path between a and b is in S, then every minimal path between a and b is in S.*

(c) *If $a, b \in S$, then every minimal path from a to b is in S.*

(d) *If a, b are adjacent cells separated by geodesic g, and $a \in S$ but $b \notin S$, then every cell in S is on the same side of g as a.*

(e) *S is convex*

Proof.

(a)⇒(b) By Theorem 9.9, it suffices to show that if S is closed, and P is a path in S, then every deformation of P is in S. This is trivial, by definition of closed.

(b)⇒(c) Let S satisfy (b). We first show that S is closed. Let w, x, y, z be four cells in S that meet around an interior vertex, in counterclockwise order, Suppose that $x, y, z \in S$. If $x = z$, then the cell x contradicts Corollary 9.7. So $x \neq z$. And since medial graphs are always two-colorable, x and z are not adjacent. Consequently $x \to w \to z$ and $x \to y \to z$ are both minimal paths from x to z. Since the former is in S, the latter must also be in S, showing that S is closed.

Using this, we show that for every $a, b \in S$, there is a minimal path between a and b in S, which suffices to show (c) given (b). Since S is connected, between any two $a, b \in S$ there is at least one path. Call a path *S-minimal* if it has minimal length among all paths that lie in S. Then it remains to show that every S-minimal path is truly minimal, since an S-minimal path always exists between two cells in S.

We proceed by induction on the length n of our S-minimal path. The base cases where $n < 2$ are trivial. So suppose that $x_0 \to x_1 \to \cdots \to x_n$ is an S-minimal path for $n \geq 2$. Then the path $x_1 \to \cdots \to x_n$ is also S-minimal, so by induction it is truly minimal. In particular,

$$\mathrm{d}(x_1, x_n) = n - 1.$$

Now let g be the geodesic which separates x_0 from x_1. If x_n is on the same side of g as x_1, then by Corollary 9.6, $\mathrm{d}(x_0, x_n) = \mathrm{d}(x_1, x_n) + 1 = n$ and we are done. So we can assume that x_n is on the same side of g as x_0.

Similarly, the path $x_0 \to \cdots \to x_{n-1}$ is minimal. Since it begins with a step across the geodesic g, x_{n-1} must be on the opposite side of g from x_0. This makes x_{n-1} and x_n on opposite sides of g, so they must be separated by g. In other words, the path $x_0 \to \cdots \to x_n$ begins and ends by stepping across g. Now the path $x_1 \to \cdots \to x_{n-1}$ is S-minimal, so by induction it is truly minimal. But then property (b) implies that every minimal path between x_1 and x_{n-1} is in S. In particular, the path from x_1 to x_{n-1} that proceeds directly along the side of geodesic g, as in Figure 25, must be in S. This path is minimal because it steps over no geodesic twice, by criticality.

Consequently, every cell on the *far* side of g from x_1 to x_{n-1} is in S, as are the cells x_0 and x_n which are opposite x_1 and x_{n-1}, respectively. Thus we have the

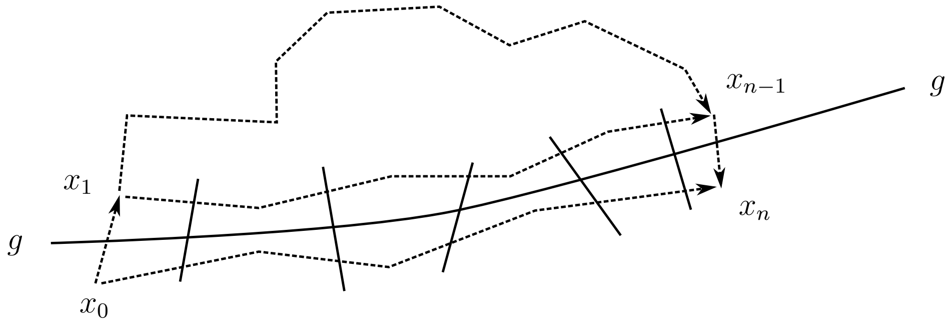

Figure 25: The path in question must begin and end by crossing g. By induction, the path that proceeds directly from x_1 to x_{n-1} along the side of g must be in S. Once all these cells are in S, the cells on the near side of g must also be in S, because S is closed. This shows that a true minimal path from x_0 to x_n is in S.

setup of Figure 25. But then the fact that S is closed easily implies that every cell on the *near* side of g from x_0 to x_n is in S. In particular, there is a truly minimal path from x_0 to x_n in S, so every S-minimal path from x_0 to x_n must be truly minimal.

(c)⇒(d) Let S satisfy (c), and suppose that a and b are adjacent cells, with $a \in S$ and $b \notin S$. Let H be the pseudo-halfplane containing a but not b. We wish to show that $S \subseteq H$. Suppose for the sake of contradiction that $c \in S$ and $c \notin H$. Then by property (c) and Corollary 9.6, the convex closure of $\{a, c\}$ is in S. We claim that $b \in \overline{\{a, c\}} \subseteq S$, a contradiction. Otherwise, there must be some pseudo-halfplane H' containing a and c but not b. But then the geodesic which defines H' must separate a from b, so H' must be H or its complement. Either way, H' cannot contain both a and c, since $a \in H$ and $c \notin H$.

(d)⇒(e) Let S satisfy (d). If S is empty, the result is trivial, so suppose that S is nonempty. Call g a *defining geodesic* of S if there are adjacent cells a, b separated by g, with $a \in S$ and $b \notin S$. In such a circumstance, call the half-plane on the side of g containing a a *defining halfplane*. Then condition (d) implies that S is contained in every defining halfplane. Let S' be the intersection of all defining halfplanes. We have $S \subseteq S'$, and S' is convex. It remains to show that $S = S'$. Otherwise there exists some cell in $S' \setminus S$. By Theorem 9.5, S' is connected, so there is some path from a cell in S to a cell in $S' \setminus S$. In particular, there exists some pair of adjacent cells $a, b \in S'$ with $a \in S$ and $b \notin S$. But then the halfplane containing a but not b is a defining halfplane of S, so $b \notin S'$, a contradiction.

(e)⇒(a) Let S be convex. Then Theorem 9.5 shows that S is connected. Convex sets are always closed, because half-planes are closed and intersections of closed sets are closed.

□

Corollary 9.11. *Let S be a cellset. Then S is closed if and only if each connected component of S is convex.*

Proof. It is easy to check that S is closed if and only if each connected component of S is closed. The result then follows by the equivalence (a) $\iff$ (e) of Theorem 9.10. $\square$

Next we return to the notion of rank:

Theorem 9.12. *If S is a convex cellset, then* $\text{rank}(S)$ *is one plus the number of geodesics g which separate cells of S.*

Proof. By Theorem 9.5, it suffices to show this in the case where S is the entire medial graph. Then we need to show that the rank of S is one plus the number of geodesics. Let c count the number of cells, v_i count the number of interior vertices, v_b count the number of boundary vertices, and e count the number of edges including the segments of the boundary curve. Interior vertices have degree 4 while boundary vertices have degree 3, so that

$$2e = 4v_i + 3v_b.$$

Since all cells in a critical medial graph are homeomorphic to disks (by Corollary 9.7), Euler's formula applies, yielding

$$c + v_i + v_b = e + 1 = 2v_i + \frac{3}{2}v_b + 1.$$

Thus

$$\text{rank}(S) = c - v_i = \frac{3}{2}v_b - v_b + 1 = \frac{v_b}{2} + 1.$$

But $v_b/2$ counts the number of geodesics, because each geodesic has two endpoints. $\square$

Recall that a cellset is "safe" if its closure has the same rank as it. The main statement that we will need to set up our boundary value problems is the following:

Lemma 9.13. *Let S be a safe cellset with connected closure. Let $b \notin \overline{S}$ be a cell adjacent to a cell in $\overline{S}$. Then $S \cup b$ is a safe cellset with connected closure.*

Proof. Note that $\overline{S \cup \{b\}} = \overline{\overline{S} \cup \{b\}}$ by general abstract properties of closure operations. Since b is adjacent to $\overline{S}$ and $\overline{S}$ is connected, Lemma 9.3 shows that $\overline{S \cup \{b\}} = \overline{\overline{S} \cup \{b\}}$ is connected.

Since both $\overline{S}$ and $\overline{S \cup \{b\}}$ are closed and connected, they are convex by Theorem 9.10. So Theorem 9.12 tells us that the ranks of $\overline{S}$ and $\overline{S \cup \{b\}}$ are determined by the number of geodesics which cut across them. Now any geodesic which cuts across $\overline{S}$ will also cut across $\overline{S \cup \{b\}}$, so

$$\text{rank}(\overline{S}) \leq \text{rank}(\overline{S \cup \{b\}}).$$

But in fact at least one geodesic which cuts across $\overline{S \cup \{b\}}$ does not cut across $\overline{S}$, specifically the geodesic g which separates b from its neighbor in $\overline{S}$. This geodesic cannot cut across $\overline{S}$ by condition (d) in Theorem 9.10. So we have

$$\text{rank}(\overline{S}) + 1 \leq \text{rank}(\overline{S \cup \{b\}}). \tag{4}$$

Now

$$\text{rank}(S) = \text{rank}(\overline{S})$$

because S is safe. And since $b \notin \overline{S}$, $S \cup \{b\}$ is not a simple extension of S, so that

$$\mathrm{rank}(S \cup \{b\}) = \mathrm{rank}(S) + 1$$

by Lemma 8.4. And by Lemma 8.6,

$$\mathrm{rank}(\overline{S \cup \{b\}}) \leq \mathrm{rank}(S \cup \{b\}) = \mathrm{rank}(S) + 1 = \mathrm{rank}(\overline{S}) + 1.$$

Combining this with (4), we see that equality must hold, so in particular

$$\mathrm{rank}(\overline{S \cup \{b\}}) = \mathrm{rank}(S \cup \{b\}),$$

so $S \cup \{b\}$ is safe. □

We will use this lemma to inductively build up safe sets of boundary cells.

10 Nonlinear recovery

To complete the proof of Theorem 6.7, we need to show how to recover the conductance functions of boundary triangles in critical medial graphs.

Lemma 10.1. *Let M be a critical medial graph, b be a boundary cell, and g be one of the geodesics by b, as in Figure 26. Then there are two sets of boundary cells T, S such that*

- *S is a safe cellset:* $\mathrm{rank}(S) = \mathrm{rank}(\overline{S})$.
- *$T \subseteq S$.*
- *$b \in S \setminus T$.*
- *$\overline{T}$ is the halfplane on the far side of g from b.*
- *$\overline{S}$ is the entire medial graph.*

Proof. Let H be the half-plane on the far side of g from b. Number the boundary cells in H $a_1, a_2, \ldots, a_n$ in consecutive order. Let $T_0 = \{a_1\}$, and then recursively define T_{i+1} to be $T_i \cup \{a_{j_i}\}$ where j_i is the minimal j such that $a_j \notin \overline{T_i}$. Stop once $\overline{T_i}$ contains every a_j. This gives us a chain of sets of boundary cells

$$T_1 \subset T_2 \subset \cdots \subset T_n.$$

By choice of a_{j_i}, we must have $a_{j_i - 1} \in \overline{T_i}$ for each i. Consequently, a_{j_i} is adjacent to a cell in $\overline{T_i}$. Then by Lemma 9.13, we see that if T_i is safe with connected closure, so is $T_{i+1} = T_i \cup \{a_{j_i}\}$. Since $T_1 = \{a\}$ clearly has $\overline{T_1} = T_1$, T_1 is safe with connected closure. Thus every T_i is safe with connected closure.

Let $T = T_n$. Then T is safe with connected closure, and $\overline{T}$ contains every boundary cell in H. Since $T \subseteq H$, $\overline{T} \subseteq \overline{H} = H$. But conversely, if any half-plane H' contains $\overline{T}$, then its defining geodesic g' must start and end on the far side of H', so by criticality it cannot cross g, and $H \subseteq H'$. Consequently the convex closure of $\overline{T}$ contains H. But by Theorem 9.10, $\overline{T}$ is already convex, so we see that $H \subseteq \overline{T}$. So $H = \overline{T}$.

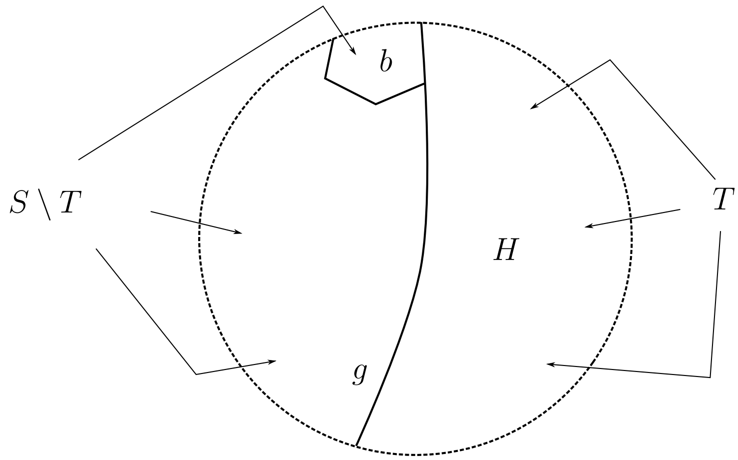

Figure 26: The setup of Lemma 10.1. Given boundary cell b and adjacent geodesic g, we can find sets S and T of boundary cells having the desired properties.

We then let $T_{n+1} = T_n \cup \{b\}$. Now $b \notin \overline{T} = H$, so by Lemma 9.13 again, T_{n+1} is still safe with connected closure. Continuing on inductively we can keep adding more boundary cells until we have a set S which is safe, has connected closure, contains $T \cup \{b\}$, and has every boundary cell in its closure. But $\overline{S}$ is convex by Theorem 9.10, and any convex set containing every boundary cell must be the whole medial graph, since no halfplane contains every boundary cell.

See Figure 27 for an illustration of all these steps in a small example. □

Using this we prove the other direction of Theorem 6.7

Proof (of Theorem 6.7). By Theorem 7.16, we only need to show that conductance functions of boundary triangles are recoverable. Let b be a boundary triangle, and let g be one of the two geodesics which define b. Find cellsets T and S as in Lemma 10.1. Fix boundary values of x at cell b, and 0 on $S \setminus \{b\}$. Since S is safe and has closure the entire medial graph, it follows by Theorem 8.7 that these boundary values uniquely determine a labelling ϕ of the entire medial graph. The restriction of ϕ to $\overline{T}$ is a labelling of $\overline{T}$, and so is the constant zero function. Since these extend the same labelling of T, Theorem 8.7 shows that ϕ must identically vanish on $\overline{T}$. By Lemma 10.1, $\overline{T}$ contains a neighbor a of b, and the cell c that diagonally touches a. Letting d be the fourth cell that meets at the apex of b, we have from Equation (3)

$$\phi(d) = \phi(d) - \phi(b) = \pm\gamma^{\pm}(\phi(a) - \phi(c)) = \pm\gamma^{\pm}(x)$$

where the signs depend on the orientation of e, the order of a, b, c, d around their meeting point, and whether or not e is a boundary spike. But we can determine $\phi(d)$ from the boundary relation Ξ, since d is a boundary cell. It follows that we can read off $\pm\gamma^{\pm}(x)$ directly from Ξ. □

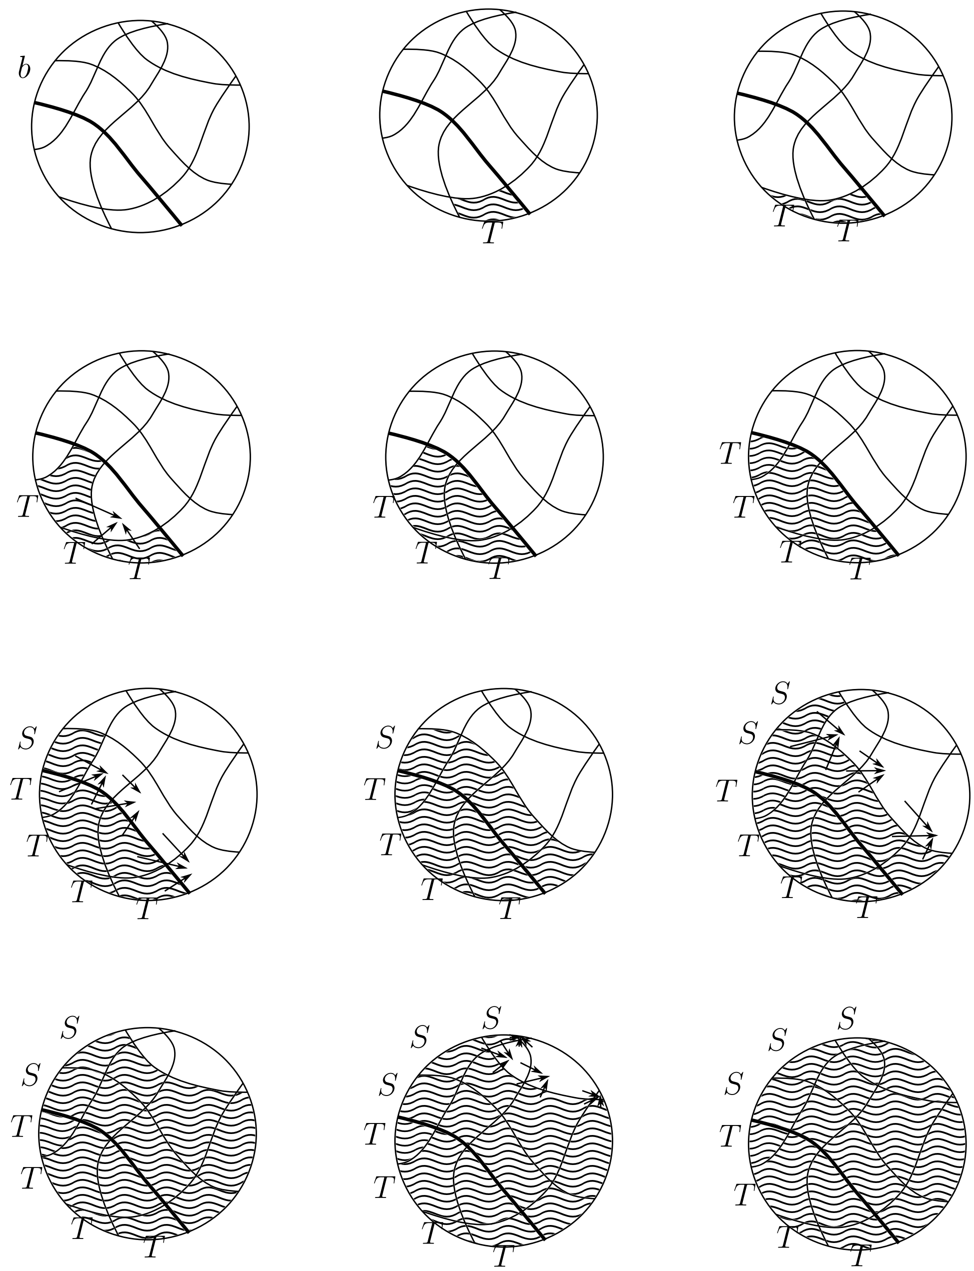

Figure 27: A small example of the proof of Lemma 10.1. By a series of safe simple extensions, we first fill up the desired pseudo halfplane (the one below the darkened geodesic), and then proceed to fill up the rest of the medial graph. Arrows show the propagation of information. Here, g is the darknened geodesic and the boundary cell b is the one above the left end of g. S consists of all four boundary cells below g, while T additionally contains b and the next two boundary cells after b in clockwise order.

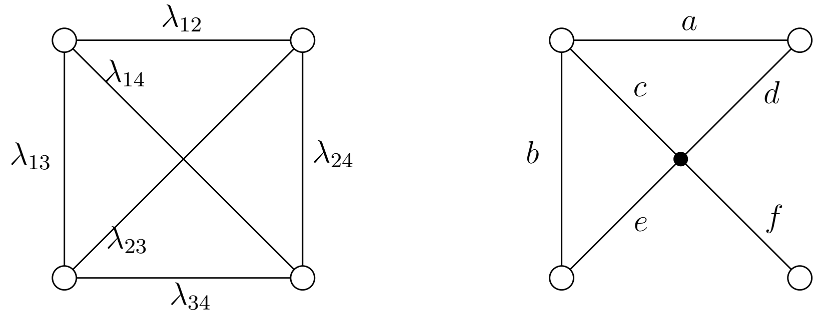

Figure 28: A complete graph on four vertices, and a critical circular planar graph. Both have six edges, and generically, they are interchangeable, after an appropriate birational change of variables.

11 An application of negative conductivities

Even though our original motivation for Theorem 6.7 was the recovery of networks that had nonlinear *monotone* conductance functions, monotonicity was never needed in the proof. In fact, we could have allowed our currents and voltages to take values in an arbitrary abelian group. This result is rather strong, and one wonders what sort of applications it might have.

One interesting implication is that we can allow linear conductance functions $\gamma(x) = cx$ arbitrary $c \in \mathbb{C} \setminus \{0\}$, and recovery will be possible. While these nonpositive conductivities might seem arbitrary, there are a couple of circumstances where they come up naturally. One comes from knot theory and will be discussed in a later paper.

The other comes from a generalization of the Y-Δ transformation of Section 7.[2] Consider the circular planar network shown on the right side of Figure 28, and suppose that all conductors are linear. By a direct calculation, one can show that this network is electrically equivalent to the network on the complete graph K_4 on the left side of Figure 28 with the conductivities specified as follows:

$$\lambda_{12} = a + \frac{cd}{\sigma}$$
$$\lambda_{13} = b + \frac{ce}{\sigma}$$
$$\lambda_{14} = \frac{cf}{\sigma}$$
$$\lambda_{23} = \frac{de}{\sigma}$$
$$\lambda_{24} = \frac{df}{\sigma}$$
$$\lambda_{34} = \frac{ef}{\sigma}$$

[2]The following idea first appeared in Michael Goff's REU paper *Recovering Networks with Signed Conductivities* in 2003. This paper claims to have a proof that signed conductivities can be recovered in critical circular planar networks, provided that the Dirichlet-Neumann relation is a function. The proof seems flawed, though.

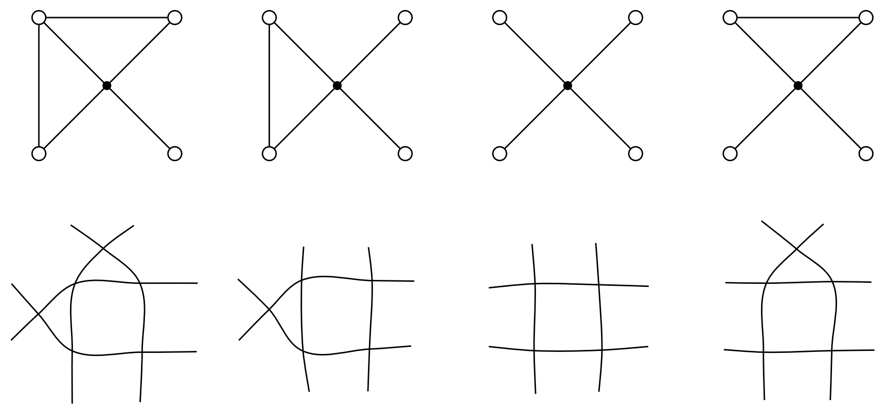

Figure 29: A K_4 with positive conductances can be turned into one of the four critical circular planar graphs of the top row, possibly introducing negative conductivities in the process. The four associated medial graphs are shown on the bottom row.

where $\sigma = c+d+e+f$, provided that $c+d+e+f$. Conversely, if the λ_{ij} are positive numbers, then the K_4 on the left side of Figure 28 is electrically equivalent to the graph on the right side of Figure 28 with conductances given as follows:

$$a = \lambda_{12} - \frac{\lambda_{14}\lambda_{23}}{\lambda_{34}}$$
$$b = \lambda_{13} - \frac{\lambda_{14}\lambda_{23}}{\lambda_{24}}$$
$$c = \frac{q\lambda_{14}}{\lambda_{24}\lambda_{34}}$$
$$d = \frac{q}{\lambda_{34}}$$
$$e = \frac{q}{\lambda_{24}}$$
$$f = \frac{q}{\lambda_{23}}$$

where

$$q = \lambda_{14}\lambda_{23} + \lambda_{23}\lambda_{24} + \lambda_{23}\lambda_{34} + \lambda_{24}\lambda_{34}$$

So we have a way of transforming any K_4 (with positive conductivities) into an electrically equivalent network on the graph on the right side of Figure 28, possibly introducing nonpositive conductivities in the process. Now from the equations for $a, \ldots, f$, we see that only a or b can be zero or negative. A conductivity of zero is equivalent to a deleted edge, so any K_4 (with positive conductivities) can be turned into one of the four graphs of Figure 29 (with conductivities in $\mathbb{R} \setminus \{0\}$). These graphs' associated medial graphs are also shown in Figure 29.

Like a Y-Δ transform, we can apply this transformation to a K_4 embedded locally within a larger graph. For example, the graph of Figure 30 (with positive conductivities) will be electrically equivalent to one of the graphs of Figure 31 (with nonzero conductivities). Now these resultant graphs are circular planar, and it is easy to check

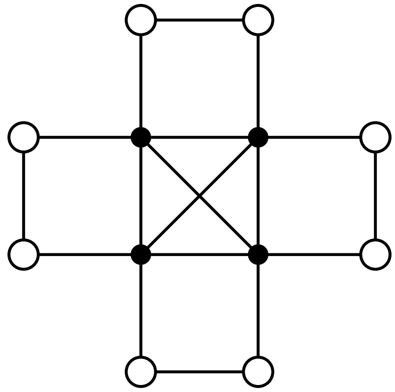

Figure 30: If all the conductivities in this graph are positive, it will be electrically equivalent to one of the four graphs in Figure 31, all of which are recoverable.

that they are critical. Consequently, the original graph of Figure 30 is (weakly) recoverable - as long as the response matrix provides enough information to distinguish the four graphs of Figure 31. This is implied by the following lemma, which generalizes results in [1].

Definition 11.1. *Let M_1 and M_2 be two critical medial graphs, embedded on the same disks with boundary vertices in the same locations. Then we say that M_1 and M_2 are* equivalent *if for every boundary vertex v, the geodesics g_{M_1} of M_1 and g_{M_2} of M_2 beginning at v end at the same boundary vertex.*

Lemma 11.2. *Let M_1 and M_2 be two colored critical medial graphs with conductivities, with the same boundary vertices. If M_1 and M_2 have the same boundary relation, then M_1 and M_2 are equivalent.*

Proof. Let w be one of the common boundary vertices of M_1 and M_2. Let g_i be the geodesic in M_i with one endpoint at w, and let v_i be the other endpoint of g_i. Suppose for the sake of contradiction that $v_1 \neq v_2$. Let H be the half-plane in M_1 on the side of g_1 containing v_2, and let b be the cell of M_1 that is by w but not in H, as in Figure 32. By Lemma 10.1, we can find sets of boundary cells $T \subseteq S$, with $b \in S \setminus T$, S and T safe, $\overline{T} = H$, and $\overline{S}$ equal to the whole medial graph M_1. So we can set boundary values of 0 on T, 1 at b, and whatever we like on $S \setminus \{b\}$, and there will be at least one labelling of M_1 with these boundary values. In fact, since $\overline{T} = H$, the labelling must take the value zero on all of H, so we can even assert boundary values of zero on every boundary cell in H.

In particular, we see that for M_1, there is at least one set of valid boundary values that takes the value 1 at b, and 0 at every boundary cell in H. We claim that this fails in M_2, a contradiction. To see this, let Q be the set of boundary cells in M_2 corresponding to the boundary cells of H. That is, Q consists of all the boundary cells in M_2 along the boundary arc from w to v_1 that passes through v_2. Let b' be the cell of M_2 corresponding to b. If M_1 and M_2 had the same boundary relation, then there should be a labelling of M_2 which takes the value 0 at every cell in Q, and 1 at b'. But this is impossible because $b' \in \overline{Q}$, so that boundary values of zero on Q should

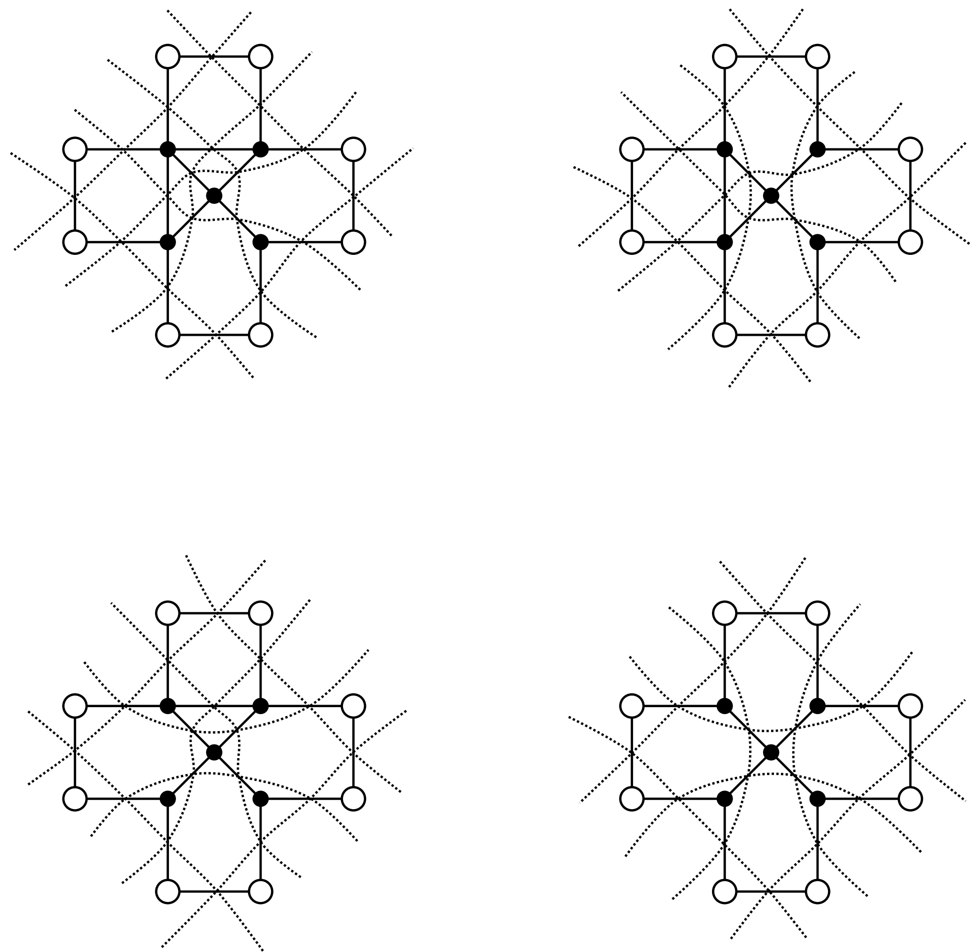

Figure 31: These four graphs are recoverable. Furthermore, they are electrically distinguishable from each other, by Lemma 11.2.

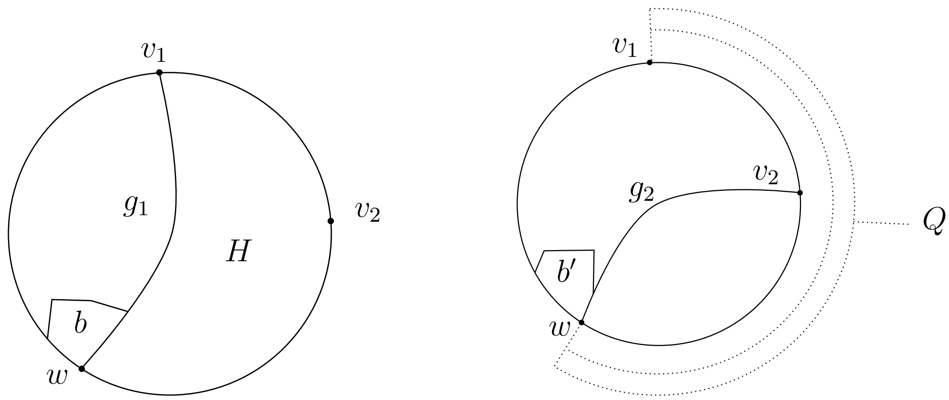

(a) M_1 (b) M_2

Figure 32: Two nonequivalent medial graphs.

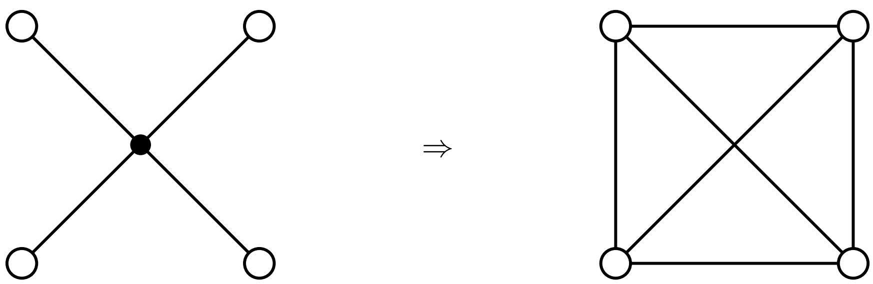

Figure 33: Turning a "star" into a complete graph.

force a value of zero at b'. To see that $b' \in \overline{Q}$, note first that Q is connected, so $\overline{Q}$ is connected and thus convex. But b' is directly across g_2 from a cell of Q, so if any geodesic separates b' from $\overline{Q}$, it would be g_2, which is impossible because cells on both sides of g_2 are in Q, specifically the two cells around v_2. So no geodesic separates b' from Q, and $b' \in \overline{Q}$.

So the assumption that $v_1 \neq v_2$ was false. Consequently g_1 and g_2 have the same start and endpoints. □

By Lemma 11.2 and an examination of the medial graphs of Figure 31, we see that all are electrically distinct, and consequently the original graph of Figure 30 is weakly recoverable.

In contrast, the graph of Figure 34a is not recoverable because in one of the cases, it is equivalent to Figure 34b, which is not recoverable. To verify this, note that we can turn a "four-star" into a K_4 as in Figure 33, without producing negative or zero conductivities. This map is invertible (because the "four-star" is recoverable), so if Figure 34a were recoverable, then Figure 34b would also be recoverable, but it is not since its medial graph clearly contains a lens.

However, we can say that Figure 34a is *generically recoverable*, because generically it is equivalent to Figure 34c which is weakly recoverable. In fact, the map from conductivities on Figure 34a to response matrix is birational.

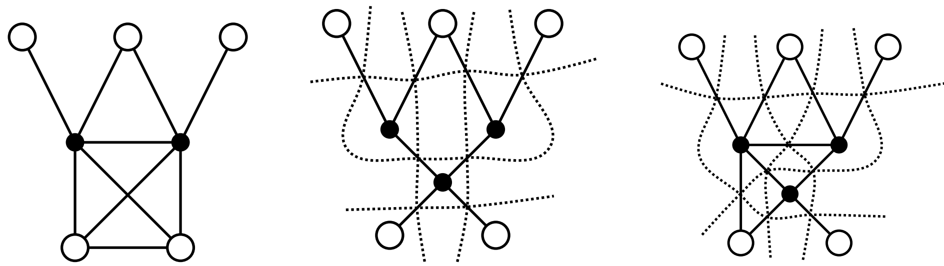

Figure 34: Only 34c is recoverable, but 34a is *generically* recoverable, in some sense.

12 Closing Remarks

We have shown that if Γ is a circular planar graph, then Γ is strongly recoverable iff its medial graph is critical. This extends results of Curtis and Morrow and students in the UW Mathematics REU [1], who showed the same for weak recoverability. An interesting corollary is that Γ is weakly recoverable if and only if it is strongly recoverable. An obvious conjecture is that this holds for general Γ, not necessarily circular planar. However, it may be necessary to put further restrictions on conductance functions (like monotonicity or smoothness).

The recovery algorithm presented above makes surprisingly few assumptions on conductivity functions. This suggests that there might be interesting applications of our method. Letting currents and voltages live in finite fields might have interesting combinatorial applications. In the linear case, we have shown that the rational map from conductivities to response matrices remains injective and somewhat well-defined when extended from $(\mathbb{R}_{>0})^n$ to $(\mathbb{C} \setminus \{0\})^n$. This is a purely algebraic statement, which might have some uses.

Ian Zemke has also generalized some of the results of this paper to infinite networks and discrete Schrödinger Networks.

References

[1] Edward B. Curtis and James A. Morrow. *Inverse Problems for Electrical Networks.* Number 13 in Series on Applied Mathematics. World Scientific, Singapore, 2000.