Model Theoretic Methods in Discrete Inverse Problems: o-minimality and n-to-1 Graphs in the Study of Electrical Networks

Reid Dale

November 9, 2014

Contents

Chapter 1

Introduction

Consider the following problem: you are given a black box containing fragile circuitry with some of its nodes and conductors sticking out, as well as a blueprint for the whole network. You are able to measure the voltages and the currents of the visible nodes and conductors (the *boundary* data), but opening the black box would cause the device to break. The *electrical inverse problem* is to deduce the conductances of each conductor on the inside of the black box (the *interior* of the network) from this boundary information so as to preserve the integrity of the box.

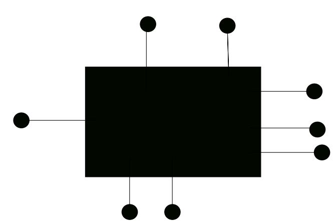

Black Box Network

This problem has been studied over the last 25 years in Jim Morrow's Summer REU program, resulting in a large number of results that draw from ideas in probability, functional analysis, and combinatorics. With these tools, students have shown that under certain conditions there is a unique, well-defined solution to this inverse problem. It's also been shown that there are finite networks with given boundary data whose inverse problem has *continuum many* solutions. In the last ten years, others have exhibited networks whose boundary conditions admit exactly n solutions to the inverse problem, for all natural numbers n. These results are all *constructive* in the sense that they includeded algorithms for constructing such networks or were explicit examples of such phenomena.

When I first studied inverse problems on electrical networks, I had wondered if there could be countably many solutions to the inverse problem (or, more generally, if there could be exactly κ many solutions for $\aleph_0 \leq \kappa < 2^{\aleph_0}$, not assuming the continuum hypothesis). After all, the constructive methods outlined above did not themselves rule out the existence of such networks. In order to study this problem, I turned to the methods of *model theory*, a purportedly esoteric branch of mathematical logic. In Summer 2013 I was able to show that the answer to this question was a resounding *no*. There is a strict dichotomy in the number of solutions to a inverse problem for *any* given finite electrical network: there are either finitely many solutions or there are continuum many.

After describing this to Tom Scanlon, he suggested that I extend these results to the computable setting. Again, using methods of model theory, I was able to prove the following strengthening of the dichotomy theorem: not only does the dichotomy theorem hold for finite electrical networks, but under natural assumptions on the type of number used to represent the problem there exists an *algorithm* that separates these two cases and actually computes how many solutions to the inverse problem there are. This result hinges upon two facts: that the property of having continuum many solutions to the inverse problem is in fact a first-order property in this setting (which is not at all obvious *a priori*), and that there exists an algorithm for quantifier elimination in the language of ordered fields.

Before we begin, I would like to make a few remarks about methodology. As briefly mentioned above, most investigations of inverse problems in electrical networks were highly constructive and led to strong results for specific (but very interesting) classes of electrical networks. This project, however, is completely general in its scope. Using the fact that I can interpret, for any finite electrical network, the data relevant to studying inverse problems in the language of real closed fields I am able to prove results about all finite electrical networks. This necessarily means that I am not proving completely optimized results. For instance, the algorithm that I use for enumerating solutions to electrical inverse problems on a given graph *does* necessarily halt, but the proof that it does so does not give any indication as to how long the algorithm will run. No doubt if I were to restrict attention to some nice class of graphs the computation would have explicit, bounded run time. It is my hope that future work will be done to produce reasonable algorithms for either solving or counting solutions of these problems in a more restricted setting. In this way the level of generality considered is both a blessing and a curse; while I show that there are effective methods for computing quantities related to the electrical inverse problem, they are *exceedingly* slow.

I hope that this work inspires an approach to the electrical inverse problem that, in dialectical tandem with the concrete and constructive methods that preceded it, stimulates further research in the algebro-geometric and structural components of discrete inverse problems.

Acknowledgements I would first like to thank my two undergraduate mentors, Jim Morrow and Tom Scanlon, for their patience and support spanning multiple years. I thank Jim for suggesting a general project of applying real algebraic geometry to the problems dealt with in his REU and to Tom for suggesting that I extend my results by considering issues of decidability and computability. In many ways this project reflects an amalgamation of the mathematical intuitions and ways of thinking that I have gleaned from them during my time as an undergraduate at the University of Washington in Seattle. I also acknowledge financial support from three sources: Jim Morrow's NSF REU grant, his Barbara Hand Sando and Vaho Rebasoo endowed professorship, as well as the UW's NSF RTG grant.

In a very real sense I would not have been able to do this work without the encouragement of my parents and siblings. I thank them for supporting me as I made my way through the University in the Early Entrance Program. The transition was certainly a shock at first, but their care ensured that I could make the most of my undergraduate education.

Chapter 2

The Inverse Problem for Electrical Networks

2.1 Definitions and Formulation of the Inverse Problem

Recall the informal setup of the inverse problem given in the introduction: you are presented with an electrical network along with some given *boundary* data (currents and voltages along the boundary) and, because of the fragility of the materials, trying to physically measure important electrical quantities (namely, the conductance of each edge) on the *interior* would cause the device to break. The *electrical inverse problem* is to deduce the conductances of each conductor on the inside of the black box (the *interior* of the network) from this boundary information so as to preserve the integrity of the electrical network

In this section we will build up a mathematical framework to precisely state the inverse problem. The central notion will be that of the *measurement map* of an electrical network L_G, a map that coherently assigns global data to boundary data. The map L_G takes the conductances of all of the edges on G, γ, and returns the linear *response map* Λ_γ which sends boundary voltages to boundary currents. The key insight will be that solutions to given inverse problems can be interpreted as fibers of this map L_G: given a candidate for a response map Λ, a solution to the inverse problem is simply an element of the fiber $L_G^{-1}(\Lambda)$.

We will give a formula for L_G and note that, when defined, it is given by a system of rational maps defined on the space of conductivities $(\mathbb{R}^+)^{|E|}$ and returns a map in $M_{\partial G}(\mathbb{R})$. This characterization of both L_G and the inverse problem will put us in the realm of model theory and real algebraic geometry, where we will be able to study the inverse problem in high generality and great detail.

Ordinary electrical networks can be modeled using *weighted graphs*, wherein each vertex represents a node of the network and each edge with edge weight γ corresponds to a measurement of the conductance between each node. However, as the informal statement suggests, we need to make a distinction between *interior* and *boundary* nodes. The mathematical device we use to model this is what Curtis and Morrow refer to as a *resistor network with boundary.*

Definition 1. An electrical network with boundary, G consists of the following data:

- A connected graph G with at most one edge between each node.
- A partition of the vertices V_G of G into disjoint sets ∂G and $\operatorname{int} G$ (standing for the "boundary vertices" and "interior vertices" respectively). We moreover assume that $|\partial G| \geq 1$ and that $|\operatorname{int} G| \geq 1$, to avoid degenerate cases.

In addition, we will consider any function from edges of G to $\mathbb{R}^+$, $\gamma : E_G \to \mathbb{R}^+$ to be a *conductance function*. □

Remark 1. Note that the function γ simply assigns nonzero weights to each edge of G. We *could* relax the condition that $\gamma(e) > 0$ for all e to $\gamma(e) \geq 0$, but for our purposes this data is equivalent to the data given by the graph G' with edges $E_{G'} = \operatorname{Supp} \gamma$ equipped with the restriction of γ to $E_{G'}$.

For our formulation of the inverse problem, we assume that we are given the "blueprint" of our network, which is simply a graph G. The other piece of data that we need is a precise notion of what we mean by "boundary information," which we take to be encapsulated by a certain linear map called the *response map*, which assigns arbitrary boundary voltages to boundary currents by way of Ohm's Law. We now construct the map Λ_γ in two ways, both due to Curtis and Morrow.

Begin by ordering the vertices of G in such a way that $\partial G = \{v_1, \cdots, v_n\}$ and $\operatorname{int} G = \{v_{n+1}, \cdots, v_{n+m}\}$.. Then, given a conductance function $\gamma : E_G \to \mathbb{R}^+$ and potential function $u : V_G \to \mathbb{R}$ we are able to define the *current at edge e_{pq} due to u* to be the quantity

$$c(e_{pq}) = \gamma(pq)[u(p) - u(q)]$$

and, for a given node p we define the *current into p* by the formula

$$\phi(p) = \sum_{q \sim p} c(e_{pq}) = \sum_{q \sim p} \gamma(pq)[u(p) - u(q)]$$

The map that we want to construct is the map that assigns a given potential u_∂ defined on all of ∂G to the current into each $p \in \partial G$. *A priori* there is a problem with this definition; the way we defined the net current flow $\phi(p)$ required us to know the values of potentials on the interior vertices as well. To show that this association of u_∂ to ϕ_∂ is in fact well defined, we need to show that there is a unique way to extend u_∂ to a potential u defined on all of V_G in an electrically coherent fashion. The way that to do this is to require that the electrical network G satisfies the following condition on interior nodes:

Principle 1. (Kirchhoff's Law) Let $p \in \operatorname{int} G$ be an interior node. Then

$$\phi(p) = \sum_{q \sim p} \gamma(pq)[u(p) - u(q)] = 0$$

In other words, the net current flow at an interior node is zero. □

This motivates the definition of a γ-harmonic function:

Definition 2. A γ-harmonic function is a potential function $u : V_G \to \mathbb{R}$ such that each interior node p satisfies Kirchhoff's Law. □

To show that this definition of the response matrix is well-defined, therefore, it suffices to show that given conductance function γ and boundary potential u_∂ there is a unique γ-harmonic function u defined on all of V_G extending u_∂. It is a result of Curtis and Morrow, using methods of "discrete complex analysis," that this is in fact the case.

Theorem 1 (Curtis-Morrow)**.** *Given an electrical network (G, γ) and boundary potential function u_∂, there exists a unique γ-harmonic extension u of u_∂ defined on the whole of V_G.*

Corollary 1. *The response map Λ_γ is a well-defined linear map.*

Moreover, Curtis and Morrow were able to give an explicit formula for the response matrix in terms of an auxiliary matrix called the *Kirchhoff matrix*.

Definition 3. Let (G, γ) be an electrical network. The Kirchhoff matrix of (G, γ) is defined to be the matrix K_γ with entries

- For $i \neq j$, $(K_\gamma)_{ij} = -\gamma(ij)$ for $i \neq j$ where, by convention, $\gamma(ij) = 0$ if $v_i \nsim v_j$
- For $i = j$ take $(K_\gamma)_{ij} = \sum_{j \neq i} \gamma(ij)$.

□

Note that K_γ is symmetric and has row (and column) sums equal to 0. Note also that our choice of indexing gives rise to the following decomposition of K_γ:

$$K_\gamma = \begin{bmatrix} A_\gamma & B_\gamma \\ B_\gamma^T & D_\gamma \end{bmatrix}$$

where

- A_γ is the top left $\partial G \times \partial G$ submatrix of K_γ, keeping track of weighted edges between vertices in ∂G
- B_γ (resp B_γ^T) is the $\operatorname{int} G \times \partial G$ (respectively $\partial G \times \operatorname{int} G$) submatrix of K_γ, keeping track of weighted edges between vertices in $\operatorname{int} G$ and vertices in ∂G.
- D_γ is the $\operatorname{int} G \times \operatorname{int} G$ submatrix of K_γ, keeping track of weighted edges between vertices in $\operatorname{int} G$ and $\operatorname{int} G$.

More generally, for an ordered subset $P, Q \subseteq V_G$ let $K_\gamma(P; Q) = (\gamma_{ij})_{(i,j) \in P \times Q}$.

Curtis and Morrow related K_γ and the above submatrices to the response matrix Λ_γ in a particularly neat way, giving an explicit formula for computing Λ_γ.

Theorem 2 (Curtis-Morrow, Lemma 3.8 and Theorem 3.9)**.** *Suppose that (G, γ) is a connected electrical network with boundary. Then K_γ is positive semidefinite and for any proper ordered $P \subseteq V_G$ the matrix $K_\gamma(P; P)$ is positive definite. Furthermore, the response matrix Λ_γ is given by the equation*

$$\Lambda_\gamma = A_\gamma - B_\gamma D_\gamma^{-1} B_\gamma^T$$

This formula gives us a connection between K_γ and Λ_γ. In fact, it shows us that we can express the entries of the Λ_γ as *rational* combinations of the γ_{ij}, since finding the inverse D_γ can be done with Gaussian elimination and matrix addition and multiplication are polynomial operations. It is also clear by this theorem that with our assumptions on γ (thinking of γ as an $|E|$-tuple of positive real numbers) this function is well-defined on all of $(\mathbb{R}^+)^{|E|}$ as it avoids the zero locus of $\det(D_\gamma)$. Hence the association of $\gamma \mapsto K_\gamma \mapsto \Lambda_\gamma$ is a well-defined rational function on the ordered tuple γ. We summarize thusly:

Corollary 2. *There is a well-defined, rational measurement map $L_G : (\mathbb{R}^+)^{|E|} \to M_{\partial G}(\mathbb{R})$ associated to any connected graph G given by mapping*

$$\gamma \mapsto \Lambda_\gamma$$

This means that there is a rational parametrization of response matrices that arise from G. A set of conductances γ induce a unique response matrix Λ_γ given by applying L_G. The inverse problem is precisely to go the other way:

Problem. *The inverse problem (G, Λ) is the following: Given a graph with boundary G and a candidate for a response matrix $\Lambda \in M_{\partial G}(\mathbb{R})$, does there exist a set of conductances γ on E_G such that $L_G(\gamma) = \Lambda$? A solution to the inverse problem (G, Λ) such a γ.*

Phrased in this way, a solution to the inverse problem (G, Λ) is simply some element of the fiber $L_G^{-1}(\Lambda)$. We consider both the *local* and *global* aspects of the problem, wherein we study problems of counting solutions to specific inverse problems (G, Λ) and problems where we study the behavior (G, Λ) as Λ ranges over all matrices in $M_{|\partial G|}(\mathbb{R})$. We typically identify the solutions of the inverse problem (G, Λ) with the set $L_G^{-1}(\Lambda)$.

2.2 Examples and Prior Results

Constructing problems and graphs that witness both unique and non-unique recoverability has been a major industry in Jim Morrow's Electrical Networks REU almost since its inception, often requiring a great deal of combinatorial ingenuity to produce or classify them. In this section we sketch the state of knowledge prior to this thesis. We first fix some terminology:

Definition 4. Let G be a (connected) graph with boundary and $\Lambda \in M_n(\mathbb{R})$.

- Call an inverse problem (G, Λ) κ-to-1 if $|L_G^{-1}(\Lambda)| = \kappa$.
- Call a network G *weakly* κ-to-1 if there is some inverse problem (G, Λ) with $|L_G^{-1}(\Lambda)| \geq \kappa$.
- Call G *strongly* κ-to-1 if $\max\limits_{A \in M_n(\mathbb{R})} |L_G^{-1}(A)| = \kappa$.

Finally, define the *fiber cardinals* of G, $\Psi\, G$, to be

$$\Psi\, G = \{\kappa \mid \exists \Lambda \in M_{|\partial G|}(\mathbb{R}) \; |(G, \Lambda)| = \kappa\}.$$

□

A corollary of the formula for the response matrix from the Kirchhoff matrix given in the previous section is that it immediately yields some *necessary* conditions that a matrix must satisfy in order to be a response matrix.

Proposition 1 (Curtis-Morrow, Section 5.1)**.** If G is a connected electrical network with boundary with measurement map L_G, then any response matrix $\Lambda_\gamma = L_G(\gamma)$ is necessarily symmetric with rows summing to zero.

An important consequence of this is the existence of inverse problems with no solutions whatsoever.

Corollary 3. *Let G be an arbitrary connected electrical network with boundary with measurement map L_G. Then if $\Lambda \in M_{|\partial G|}(\mathbb{R})$ is not symmetric, then the inverse problem (G, Λ) is* 0*-to-*1*.*

Moving to the case of *unique* recoverability, participants in the REU program isolated a certain class of graphs for which, if a solution to the inverse problem *existed*, then it was in fact *unique*. The graphs considered were so-called *critical circular planar graphs*. The definition of this class is very involved, but contains a great deal of combinatorial insight.

Definition 5 (Curtis-Morrow, Chapter 2, almost verbatim)**.** A circular planar graph G is an electrical network with boundary embedded in the disk $\mathbb{D}^1 \subseteq \mathbb{R}^2$ so that $\partial G \subseteq S^1$ and $\operatorname{int} G \subseteq \operatorname{int} \mathbb{D}^1$.

If $p, q \in \partial G$, a path $\alpha_{pq} : p \to q$ is a sequence of edges $e_0 = pr_1, e_1 = r_1 r_2, \cdots, e_{n-1} = r_{n-1} r_n, e_n = r_n q$ such that the $r_i \in \operatorname{int} G$ and are all distinct.

A sequence of points $r_1, \cdots, r_m \in \partial G \subseteq S^1$ is said to be in circular order if going from r_1 to r_m clockwise span an arc of S^1 and $r_1 < r_2 < \cdots < r_m$ in the linear order induced by going around the arc clockwise.

Let $P = (p_1, \cdots, p_k)$ and $Q = (q_1, \cdots, q_k)$ be finite sequences of length k of boundary nodes of G. We say that P and Q are connected through G if there is a permutation $\tau : [k] \to [k]$ and disjoint paths $\alpha_{p_i q_{\tau(i)}}$ such that the only boundary nodes each $\alpha_{p_i q_{\tau(i)}}$ passes through are p_i and $q_{\tau(i)}$. P and Q are said to be a *circular pair* if $(p_1, \cdots, p_k; q_k, \cdots, q_1)$ is in circular order.

The *admissible edge-removal operations* on a graph are

- *Deletion*, where given a graph $G = (V_G, E_G)$ and edge $e \in E_G$, we construct the new graph $G_e^d = (V_G, E_G \setminus \{e\})$. For graphs with boundary we do allow edges between boundary nodes to be deleted.
- *Contraction*, where given a graph $G = (V_G, E_G)$ and edge $e = e_{pq} \in E_G$ we construct the contraction graph G_e^c as a quotient graph as follows: let $\sim \subseteq V_G^2$ be the equivalence relation identifying p and q and no other vertices. Let $V_{G_e^c} = V_G/\sim$ and let $E_{G_e^c} = E_G/\sim$ be the induced quotient on edges. For a graph with boundary we do *not* allow the contraction of an edge between two boundary vertices (but allow all other contractions).

Let $P = (p_1, \cdots, p_k)$ and $Q = (q_1, \cdots, q_k)$ form a circular pair. We say that the removal of an edge e by any of the above admissible operations *breaks the connection* from P to Q if there is a k-connection from P to Q in G but no such k-connection in the modified graph G'.

A graph G is said to be *critical* if the removal of *any* edge by the above admissible processes breaks some connection in G. □

Criticality is a measure of how *efficient* the set of paths from p to q in a circular planar graph G is. This is an important concept intuitively, since Kirchhoff's Law gives a way of describing how electrical data travels across the graph from node to node.

Jeffrey Giansiracusa was able to show that

Theorem 3 (Giansiracusa)**.** *Let G be a critical circular planar graph. Then L_G is a diffeomorphism onto its image.*

In particular, this means that critical circular planar graphs are a class of graphs for which the $\Psi\, G = \{0, 1\}$. An example of a critical circular planar graph is the graph Σ_6 (due to Curtis, Ingerman, and Morrow) depicted below:

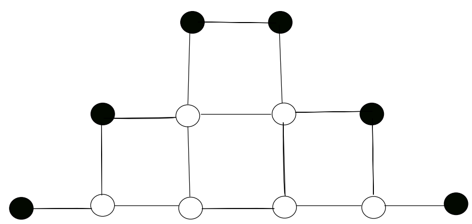

Critical Circular Planar
Network

where the filled-in vertices are boundary vertices and the empty vertices are interior vertices.

2^n-to-1 graphs were studied by Jenny French and Shen Pan, including the example of the triangle-in-triangle graph, which is strongly 2-to-1.

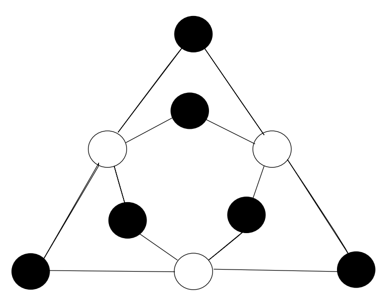

Triangle in Triangle
Network

More generally, Chad Klumb constructed n-to-1 graphs for all n in his senior thesis. The construction is very clever, and uses highly nontrivial combinatorial tools studied in the REU over the last twenty years.

Finally, it is very easy to see that not all graphs are strongly n-to-1 for $n < \omega$. A very simple example exists of a $2^{\aleph_0}$-to-1 graph, the *series network* S,

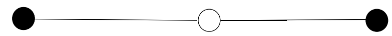

Series Network

Proposition 2. The electrical network S is strongly $2^{\aleph_0}$-to-1.

Proof. Ordering the rows with the two boundary vertices b_1 and b_2 first and the interior vertex i last and letting $\gamma_j = \gamma_{b_j,i}$, the Kirchhoff matrix of S is given by $K_\gamma = \begin{bmatrix} \gamma_1 & 0 & -\gamma_1 \\ 0 & \gamma_2 & -\gamma_2 \\ -\gamma_1 & -\gamma_2 & \gamma_1+\gamma_2 \end{bmatrix}$ and the response matrix Λ_γ by

$$\Lambda_\gamma = \begin{bmatrix} \gamma_1 & 0 \\ 0 & \gamma_2 \end{bmatrix} - \begin{bmatrix} -\gamma_1 \\ -\gamma_2 \end{bmatrix} \begin{bmatrix} \gamma_1+\gamma_2 \end{bmatrix}^{-1} \begin{bmatrix} -\gamma_1 & -\gamma_2 \end{bmatrix} = \begin{bmatrix} \frac{\gamma_1\gamma_2}{\gamma_1+\gamma_2} & -\frac{\gamma_1\gamma_2}{\gamma_1+\gamma_2} \\ -\frac{\gamma_1\gamma_2}{\gamma_1+\gamma_2} & \frac{\gamma_1\gamma_2}{\gamma_1+\gamma_2} \end{bmatrix} = \frac{\gamma_1\gamma_2}{\gamma_1+\gamma_2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

Now if we consider the inverse problem (S, Λ) for $\Lambda = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$ we find that we contain a whole *curve* of solutions in $(1,\infty) \times (1,\infty)$ carved out by the relation $\gamma_2 = \frac{\gamma_1}{\gamma_1 - 1}$. Hence S is strongly $2^{\aleph_0}$-to-1. $\square$

We have, therefore, a wide range of κ-to-1 behavior for the electrical inverse problem. All known graphs satisfied two properties: that κ was finite or $2^{\aleph_0}$, and the electrical spectrum ΨG was *finite*. In the course of this paper we show that this is in fact always the case. The precise questions we seek to resolve in this thesis are as follows:

1. Given an inverse problem (G, Λ), can we computably determine how many solutions it has; in other words, can we computably find the value of $|L_G^{-1}(\Lambda)|$?
2. Can we compute the largest such κ? Can we computably enumerate the elements of ΨG?
3. What can we say about the structure of $L_G^{-1}(\Lambda)$, independently of Λ? In other words, what is the fiberwise structure of L_G?

The answer to questions 1 and 2, it turns out that under the assumption that all quantities used to pose the inverse problem are algebraic reals, the answer is a resounding *yes*. Furthermore, we show that the only admissible cardinals κ for the inverse problem are $\kappa = n$ for some $n \in \omega$ or $\kappa = 2^{\aleph_0}$, and give a method of enumerating all elements of ΨG.

In the case of question 3, we will show that L_G is cellwise locally-trivial, meaning that there is a finite decomposition of $M_{|\partial G|}(\mathbb{R}) = C_1 \cup \cdots \cup C_n$ with C_n given by semialgebraic conditions such that $L_G \restriction_{L_G^{-1}(C_i)}$ is a fibration map. Moreover, we will be able to explicitly compute such a decomposition of $M_{|\partial G|}(\mathbb{R})$. Moreover, we will be able to computably present an implicit parametrization of the fibers $L_G^{-1}(\Lambda)$ over each cell C_i.

Finally, in the case that (G, Λ) is suitably presented and is finite-to-one, we will be able to produce approximations of *all* of its solutions by extracting from a certain algorithm a finite set of polynomials whose zero loci contain all of the coordinate valued attained as a solution of the inverse problem.

Chapter 3

A Crash Course in Model Theory

The methods we found most amenable to studying the structure of solution spaces of inverse problems in complete generality were the methods of *model theory*. A branch of mathematical logic, it may surprise some that model theory would have anything to say about an applied problem such as this one. This reaction is, however, ultimately unwarranted: the key insight of a certain strand of model theory is that if you know how a mathematical object is constructed in a first-order way in some “tame" structure, then you are immediately able to apprehend its fine structure.

In our case, the “tame structure" that we can interpret the electrical inverse problem in is the ordered field structure of $\mathbb{R}$, and the objects that we consider are constructed out of certain *semialgebraic* sets using the field operations and the order relation. This is a vague description, of course, but the elementary definitions of model theory make this intuition precise.

3.1 Languages and Structures

To begin with, we isolate what we mean by “mathematical objects constructed in a first-order way" by way of formal languages and structures.

Definition 6. A formal language (also called a *signature*) $\mathcal{L} = \langle \mathcal{C}, \mathcal{F}, \mathcal{R} \rangle$ consists of the following data:

- a set $\mathcal{C}$ of *constant* symbols with elements c,
- a set $\mathcal{F}$ of function symbols $\left\langle f, n_f \right\rangle$ with elements f of arity n_f,
- a set $\mathcal{R}$ of relation symbols $\langle R, n_R \rangle$ with elements R of arity n_R
- the logical connectives $\vee$ (or), $\wedge$ (and), $\neg$ (negation), $=$ (equality), $\forall$ (universal quantifier), $\exists$ (existential quantifier), and a set of variables $\langle x_i \rangle_{i \in \omega}$

An $\mathcal{L}$-*structure* $\mathcal{M}$ consists of the following data:

- a set M, called the *universe* of $\mathcal{M}$

- for each constant symbol c an element $c^{\mathcal{M}} \in M$
- for each function symbol f of arity n_f a total function $f^{\mathcal{M}} : M^{n_f} \to M$.
- for each relation symbol R of arity n_R a subset $R^{\mathcal{M}} \subseteq M^{n_R}$.

□

In and of itself this definition is very broad; there are almost no meaningful constraints limiting what $\mathcal{L}$-structures can occur. The main way of getting a handle on $\mathcal{L}$-structures that we wish to study is by way of specifying *theories* built up out of *first-order sentences*.

Definition 7. A first-order $\mathcal{L}$-formula ϕ is simply an expression constructed inductively as follows:

1. First of all, an $\mathcal{L}$-*term* is constructed inductively from the following:
    - a variable x_i
    - a constant symbol c
    - if f is a function symbol of arity n_f, and $t_1, \cdots, t_{n_f}$ are terms, then so is $f(t_1, \cdots, t_{n_f})$. In other words, substitutions of terms into functions are again terms.
2. An *atomic* $\mathcal{L}$-*formula* is given inductively by:
    - If t_1 and t_2 are terms, then $t_1 = t_2$ is an atomic formula
    - If R is an n_R-ary relation symbol and $t_1, \cdots, t_{n_R}$ are terms then $R(t_1, \cdots, t_{n_R})$ is an atomic formula.
3. An $\mathcal{L}$-*formula* is then given inductively as:
    - An atomic $\mathcal{L}$-formula is an $\mathcal{L}$-formula.
    - If ϕ, ψ are $\mathcal{L}$-formulas then $\neg\phi$, $\phi \vee \psi$, and $\phi \wedge \psi$ are also $\mathcal{L}$-formulas.
    - If x is a variable and ϕ a formula, then $\forall x\ \phi$ and $\exists x\ \phi$ are $\mathcal{L}$-formulas.

□

We make a distinction between formulas that have "free" variables and those that have no "free" variables. The formulas without free variables will allow us to give constraints on the class of $\mathcal{L}$-structures that we consider.

Definition 8. A variable x_i in a formula ϕ (with x_i occurring in it) is said to be *bound* if either $\exists x_i$ or $\forall x_i$ occurs in ϕ. If x_i occurs in ϕ and is not bound then we say that x_i is *free* in ϕ. We write ϕ as $\phi(x_1, \cdots, x_n)$ where the x_i are the free variables occurring in ϕ.

A formula ϕ is said to be a *sentence* if all of its variables are bound, or equivalently if it contains no free variables. □

Remark 2. The definitions of *terms*, *formulas*, and *sentences* are admittedly quite abstract and may seem unmotivated. The way to think of it is roughly that "terms" are operations on variables and constants admissible by $\mathcal{L}$, formulas are ways of constructing sets in an $\mathcal{L}$-structure out of terms and relations, and sentences are expressions that can have truth values in structures.

A set of constraints on $\mathcal{L}$-structures is called a *theory*. More precisely,

Definition 9. An $\mathcal{L}$-theory T is a set of $\mathcal{L}$-*sentences.*

An $\mathcal{L}$-structure $\mathcal{M}$ is said to model T, denoted $\mathcal{M} \vDash T$, if every sentence of T holds when interpreted in $\mathcal{M}$.

Given an $\mathcal{L}$-formula $\phi(x_1, \cdots, x_n)$ and $(m_1, \cdots, m_n) \in M^n$ we say that $M \vDash \phi(m_1, \cdots, m_n)$ just in case the resulting sentence in $\mathcal{L} \cup \{c_1, \cdots, c_n\}$, when intepreted in M by $c_i^{\mathcal{M}} = m_i$, holds true in M. □

Now that we have all of these model-theoretic definitions, we consider some examples to get a better sense of what they entail.

Example 4.

1. Let $\mathcal{L}_{ring} = \langle +, -, \cdot, 0, 1 \rangle$ be the language of rings with constants 0 and 1, function symbols $+$, $-$, and $\cdot$, and no relation symbols. We consider two $\mathcal{L}$-structures, $\mathcal{Q}_1$ and $\mathcal{Q}_2$, both with universe $\mathbb{Q}$. We let $+^1, -^1, \cdot^1, 0^1$, and 1^1 be interpreted in the standard way so that $\mathcal{Q}_1$ is a field. On the other hand, let $+^2 = -^2 = \cdot^2$ all be the constant function $f(x,y) = 0$ and let $0^2 = 1^2 = 12637845$. As it stands, $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are radically different $\mathcal{L}_{ring}$-structures, that have almost nothing to do with each other except for their shared universe $\mathbb{Q}$.

2. Let $\mathcal{Q}_1$ and $\mathcal{Q}_2$ be as above, and let T be the theory axiomatizing rings. That is, T consists of the sentences

   - $\forall x \forall y \; x + y = y + x$
   - $\forall x \forall y \forall z \; (x + y) + z = x + (y + z)$
   - $\forall x \forall y \forall z \; (x \cdot y) \cdot z = x \cdot (y \cdot z)$
   - $\forall x \forall y \forall z \; x \cdot (y + z) = x \cdot y + x \cdot z$
   - $\forall x \; x + 0 = 0 + x = x$
   - $\forall x \; x \cdot 1 = 1 \cdot x = x$
   - $0 \neq 1$

   Then $\mathcal{Q}_1 \vDash T$ but certainly $\mathcal{Q}_2 \not\vDash T$. In this way theories are natural constraints on structures: they distinguish properties of $\mathcal{L}$-structures by using the language at hand.

3. If $\mathcal{M}$ is an $\mathcal{L}$-structure then $T = Th_{\mathcal{L}}(\mathcal{M}) := \{\phi \in \mathcal{L} \mid \mathcal{M} \vDash \phi\}$ is an $\mathcal{L}$-theory. Note that for every sentence $\phi \in \mathcal{L}$, either $M \vDash \phi$ or $M \not\vDash \phi$, in which case $M \vDash \neg\phi$. Hence structures decide the truth or falsehood of *all* $\mathcal{L}$-sentences.

□

3.2 Definable Sets, Quantifier Elimination, and Decidability

Given the precise framework of the previous section we are finally able to describe what we mean when we talk about objects built in a first-order way with the notion of *definable set*.

Definition 10. Let $\mathcal{M}$ be an $\mathcal{L}$-structure. A set $D \subseteq M^n$ is said to be an A- *definable set* if there is a formula ϕ such that $D = \{(d_1, \cdots, d_n) \mid M \vDash \phi(d_1, \cdots, d_n)\}$ for some formula ϕ *perhaps involving* added constants $A \subseteq M$. In other words, an A-definable set is the locus of some first-order formula in the language $\mathcal{L} \cup \{a\}_{a \in A}$ adding $a \in A$ constants. We call a set *definable* in case it is $\varnothing$-definable.

We often identify a definable set D with *any* of the (possibly many) formulas ϕ that defines it. We also write, for a given formula ϕ, $\phi(M) := \{(m_1, \cdots, m_n) \in M^n \mid M \vDash \phi(m_1, \cdots, m_n)\}$

We call a function $f : D \to D'$ between definable sets a *definable function* if the graph

$$\Gamma_f = \{(x, f(x)) \mid x \in D\} \subseteq D \times D'$$

is definable.

We briefly list some closure properties of definable sets:

Proposition 3. Let $\mathrm{Def}^n_{\mathcal{L}}(\mathcal{M})$ be the set of definable subsets of M^n and let $\mathrm{Def}_{\mathcal{L}}(\mathcal{M}) = \bigcup_{n \in \omega} \mathrm{Def}^n_{\mathcal{L}}(M)$ be the set of $\mathcal{L}$-definable sets of M.

1. If $D, D' \in \mathrm{Def}^n_{\mathcal{L}}(\mathcal{M})$ then so are $D \cup D'$, $D \cap D'$ and $D^c = M^n \setminus D$.
2. If $D \in \mathrm{Def}^n_{\mathcal{L}}(\mathcal{M})$ and $D' \in \mathrm{Def}^m_{\mathcal{L}}(\mathcal{M})$ then $D \times D' \in \mathrm{Def}^{m+n}_{\mathcal{L}}(\mathcal{M})$
3. Projections of definable sets are definable.
4. If $f : D \to D'$ is definable then the image $f(D) \subseteq D'$ is also definable.

Proof.
1. Let $\phi_D(x_1, \cdots, x_n), \phi_{D'}(x_1, \cdots, x_n)$ define D and D'. Then $\phi_D(x_1, \cdots x_n) \wedge \phi_{D'}(x_1, \cdots, x_n)$, $\phi_D(x_1, \cdots x_n) \vee \phi_{D'}(x_1, \cdots, x_n)$, and $\neg\phi_D(x_1, \cdots, x_n)$ define $D \cup D'$, $D \cap D'$, and D^c respectively.
2. Let $\phi_D(x_1, \cdots, x_n)$, $\phi_{D'}(x'_1, \cdots, x'_m)$ define D and D'. Then making a dummy variable substitution let $\psi_{D'} : \phi_{D'}(y_1, \cdots, y_m)$ so that the sets of variables occuring in ϕ_D and $\psi_{D'}$ are disjoint. Then $\phi_D \wedge \psi_{D'}$ defines $D \times D'$.
3. Let $\phi(x_1, \cdots, x_n)$ define D. Then $\exists x_i\ \phi(x_1, \cdots, x_n)$ defines the projection
$$\pi_i(D) = \{(x_1, \cdots, \hat{x}_i, \cdots x_n) \in M^{n-1} \mid (x_1, \cdots, x_n) \in D\} \subseteq M^{n-1}.$$
4. Let $f : D \to D'$ be definable. Then $\Gamma_f \subseteq D \times D'$ is definable; take the projection onto the D' of Γ_f, $\pi_{D'}(\Gamma_f) \subseteq D'$. This coincides with $f(D)$, and so $f(D)$ is definable.

$\square$

Remark 3. There is a way to view any given definable set ϕ, thought of here as a formula, as a functor from an appropriate category of $\mathcal{L}$-structures with appropriately chosen morphisms to the category of sets given by assigning a structure $\mathcal{M} \mapsto \phi(M)$. $\square$

These closure properties hint at the way in which complexity may arise in the class of definable sets: they are closed under images of definable functions and, in particular, projections. Being closed under projections very often violates many nice properties of an otherwise “nice" family of mathematical objects, as the following examples indicate:

Example 5. 1. It is a well-known fact that projections of Borel subsets of a Polish space (for instance, $\mathbb{R}$) need not be Borel.

2. Analogously, one requires a strong set-theoretic hypothesis called *Projective Determinacy* if one wishes to show that the projection of a Lebesgue-measurable set is again Lebesgue-measurable.

3. Let $\mathbb{R}$ be considered as an $\mathcal{L}_{ring}$-structure under the usual ring operations. Then the subsets of $\mathbb{R}$ definable *without* quantifiers are simply finite Boolean combinations of polynomial equations and inequations, and it is easy to check that all such subsets are *finite* or *cofinite*. If one allows for quantifiers to be used to define subsets of $\mathbb{R}$, then the set

$$\mathbb{R}^{\geq 0} = \{x \in \mathbb{R} \mid x \geq 0\}$$

is defined by the formula $\phi(x) := \exists y \; y^2 = x$. This is an infinite-coinfinite set, and is therefore not expressible by a quantifier-free ring formula. Hence allowing quantifiers to define sets can add complexity to our class of definable sets.

With these examples and analogies in mind, we define a model theoretic notion that captures the notion that given sufficiently many constraints on *both* a language and a structure, we may be able to apprehend the structure of a definable set using only quantifier-free formulas. This notion is called *quantifier elimination*:

Definition 11. An $\mathcal{L}$-theory T is said to *eliminate quantifiers* if for all $\mathcal{L}$-formulas ϕ there exists a quantifier-free $\mathcal{L}$-formula ψ such that for all models $\mathcal{M} \vDash T$, $\mathcal{M} \vDash \forall \overline{x}(\phi(\overline{x}) \leftrightarrow \psi(\overline{x}))$. □

It is easy to construct examples of theories where quantifier elimination fails. For instance, take $T = Th_{\mathcal{L}}(\mathbb{R}) := \{\phi \in \mathcal{L}_{ring} \mid \mathbb{R} \vDash \phi\}$. Then $\mathbb{R} \vDash T$ and our above example shows that T does not eliminate quantifiers. Seeing as we hope to express the inverse problem in some first-order language that has all of the field operations, this counterexample is quite discouraging. In the next chapter we will actually *exhibit* a very simple language $\mathcal{L}$ for which $Th_{\mathcal{L}}(\mathbb{R})$ admits quantifier elimination that is also sufficient to interpret the inverse problems that we wish to consider. In many contexts, including the one we find ourselves in now, quantifier elimination is intimately related to the ability to algorithmically determine whether or not an $\mathcal{L}$-sentence ϕ holds true in a model $\mathcal{M}$ of some theory T (which we typically take to be *complete*, meaning that $T = Th_{\mathcal{L}}(\mathcal{M})$). We fix the notion of *decidability*:

Definition 12. An $\mathcal{L}$-theory T is said to be *decidable* if for every sentence $\phi \in T$, there is an algorithm which determines whether or not for all models $\mathcal{M} \vDash T$, $\mathcal{M} \vDash \phi$. □

The way that we will prove and use quantifier elimination in this context is to show the much stronger statement that given any formula ϕ (in a suitable language $\mathcal{L}$ describing sets of interest in $\mathbb{R}$) we can *algorithmically* produce a quantifier-free ψ equivalent to it modulo the $\mathcal{L}$-theory of $\mathbb{R}$. As we can computably replace ϕ with a logically equivalent quantifier-free sentence ψ the problem of deciding whether or not $\mathbb{R} \vDash \phi$ for arbitrary ϕ reduces to answering the question for *quantifier-free* ψ which themselves are simply Boolean combinations of *atomic* quantifier-free sentences.

Chapter 4

o-minimality and the First-Order Theory of $\mathbb{R}$

Throughout this section we consider $\mathbb{R}$ as an $\mathcal{L}_{RCF} = \langle +, -, <, \cdot, 0, 1 \rangle$ structure

4.1 Quantifier Elimination and Decidability of $\mathbb{R}$

In this section we exhibit an explicit algorithm for deciding sentences in the $\mathcal{L}_{RCF}$ theory of $\mathbb{R}$, realizing the promise of Section 3.2. This will split into two steps. First we will produce an algorithm for taking a formula ϕ and producing a quantifier-free formula ψ such that $\mathbb{R} \models \forall x(\phi(x) \leftrightarrow \psi(x))$. Then we will give an algorithm for deciding whether or not $\mathbb{R} \models \psi$ for any quantifier-free ψ.

We begin by giving the quantifier elimination algorithm based on the exposition in chapter 2 of Basu, Pollack, and Roy's *Algorithms in Real Algebraic Geometry* with the model theorist in mind. Throughout the section, if $p(\overline{y}, x) \in \mathbb{R}[\overline{y}, x]$ and $\overline{t} \in \mathbb{R}^{|\overline{y}|}$ write $p_{\overline{t}}(x) \in \mathbb{R}[x]$ to be the result of substituting $\overline{t}$ into $p(\overline{y}, x)$, viewed as a polynomial in x.

We first give an effective bound for roots of a given polynomial $p(x) \in \mathbb{R}[x]$.

Proposition 4 (Modulus of Roots)**.** Let $p(x) = \sum_{0 \leq i \leq n} a_i x^i \in \mathbb{R}[x]$ with $a_n \neq 0$. Then if $p(r) = 0$,

$$|r| \leq 1 + \max_{0 \leq i \leq n} \{ \left| \frac{a_i}{a_n} \right| \} = 1 + m_p$$

where $m_p = \max_{0 \leq i < n} \{ \left| \frac{a_i}{a_n} \right| \}$

Proof. If $|r| < 1$ with $p(r) = 0$ then surely $|r| < m_p$, so we assume that $|r| > 1$. Then if $p(r) = 0$, we have the equation

$$r^n = \frac{1}{a_n} \left(\sum_{i \leq n-1} a_i r^i \right)$$

Then each $\left|\frac{a_i}{a_n}\right| \leq m_p$ and so we have the inequality $r^n \leq m_p[\sum_{i \leq n-1} |r|^i]$. But $\sum_{i \leq n-1} |r|^i = \frac{|r|^n - 1}{|r| - 1}$, and so

$$|r|^n \leq m_p \frac{|r|^n - 1}{|r| - 1}$$

Rearranging terms, we yield

$$|r| \leq 1 + m_p \frac{|r|^n - 1}{|r|^n} \leq 1 + m_p$$

as desired. □

This bound allows us to express the number of solutions to a given univariate polynomial $p(x) \in \mathbb{R}[x]$ in terms of a sequence of polynomials known as the *Sturm sequence*.

Construction 1 (Sturm Sequences)**.** Let $f \in \mathbb{R}[x]$ be a squarefree polynomial. Define the *Sturm sequence* St_f of f, $\{f_i\}_{i \leq j}$, as follows:

1. Take $f_0 = f$, Let $f_1 = f'$
2. For $i \geq 0$ let f_{i+1} be $-\operatorname{rem}(f_{i-1}, f_i)$, as gotten from the Euclidean algorithm.
3. Let $j+1$ be the first index for which $f_{j+1} = 0$. Then end the Sturm sequence at $f_j = r \neq 0$ a constant.

In the case that f is not squarefree, run the procedure outlined above until you reach $f_j = \gcd(f, f')$. Then the Sturm sequence is defined to be $f_0 = \frac{f}{f_j}, \cdots, \frac{f_{j-1}}{f_j}, 1$. □

Sturm's key insight was that one could determine the number of solutions of a polynomial $p(x) \in \mathbb{R}[x]$ on a given interval (a, b) by measuring the sign changes in the Sturm sequence of p.

Theorem 6 (Sturm's Theorem, BPR 2.50)**.** *Let $f \in \mathbb{R}[x]$ and $(a, b) \subseteq \mathbb{R}$. Denote by $V(\mathrm{St}_f; a)$ the number of sign changes in the sequence of numbers $\mathrm{St}_f(a) = (f(a), f'(a), f_2(a), \cdots, f_j(a))$. Then the number of solutions of f in the interval $[a, b]$ is*

$$V(\mathrm{St}_f; a) - V(\mathrm{St}_f; b).$$

Proof. This is an unwinding of the proof of a special case of the result proven in [BPR 2.50].

We assume that f is squarefree for simplicity. We show that as a function $V(\mathrm{St}_f; t) : \mathbb{R} \to \omega$; $V(\mathrm{St}_f; t)$ is a weakly decreasing function such that $V(\mathrm{St}\, f; t) = V(\mathrm{St}\, f; u) - 1$ just in case there is exactly one root of f in the interval (t, u). Given this claim it is immediate from the number of roots of f in the interval $[a, b]$ is $V(\mathrm{St}_f; a) - V(\mathrm{St}_f; b)$. The two main properties of real polynomials that we use in this argument are that the zeroes of nonzero univariate polynomials are isolated and that f is squarefree implies that f' has no roots in common with f.

Let $Z = \{z \in \mathbb{R} \mid (\exists f_j \in \mathrm{St}_f)\, f_j(z) = 0\}$ and let $Z_i = \{z \in \mathbb{R} \mid f_i(z) = 0\}$. These are *finite*, hence discrete, sets, and $Z = \bigcup_{i \leq j} Z_i$ with each of the Z_i disjoint (by the squarefreeness of f and the construction of the f_i as remainders). Note that it is also clear that the *only* places where $V(\mathrm{St}_f;t)$ can possibly vary is at elements of Z.

First we show that if $f(x)=0$ then $\lim_{t\to x^-} V(\mathrm{St}_f;t) = \lim_{u\to x^+} V(\mathrm{St}_f;u)+1$; in other words that crossing x results in the decrease of $V(\mathrm{St}_f)$ by 1. Now as f is squarefree there is a neighborhood of x, (a_x,b_x) contained in the connected components of x in $\mathbb{R}\setminus\{Z\setminus\{x\}\}$ such that $\mathrm{sgn}(f'(y))$ is constant on (a_x,b_x) and such that $f(a_x)<0$ and $f(b_x)>0$ or visa versa. This induces a decrease by one of $V(\mathrm{St}_f)$ as one crosses x in (a_x,b_x) and since we avoided all the points $Z\setminus\{x\}$ no further sign changes occur in the f_i for $i>0$ on (a_x,b_x). Hence if $f(x)=0$ the equation

$$\lim_{t\to x^-} V(\mathrm{St}_f;t) = \lim_{u\to x^+} V(\mathrm{St};u)+1$$

holds.

Now, repeating almost that exact same argument for f_i with $0<i<j$ we can see that if $f_i(x)=0$ then there is a neighborhood (a_x,b_x) in a connected component of $\mathbb{R}\setminus\{Z\setminus\{x\}\}$ on which f_{i-1} and f_{i+1} are constant sign and with, by construction, $\mathrm{sgn}(f_{i-1}(w)) = -\mathrm{sgn}(f_{i+1}(w))$ for $w\in(a_x,b_x)$. Because f_{i-1} and f_{i+1} have opposite signs on this neighborhood a sign change induced by crossing x results in *no* change in $V(\mathrm{St}_f)$. Finally $f_j=c\neq 0$ is a nonzero constant, so it never contributes any change to $V(\mathrm{St}_f)$ and the result is proven. □

Combining the two previous results, we yield the following effective method for counting real roots of a real polynomial:

Corollary 4. *Let $f\in\mathbb{R}[x]$ and $a>m_p$. The total number of real roots of a polynomial $f(x)\in\mathbb{R}[x]$ is given by $V(\mathrm{St}_f;-a)-V(\mathrm{St}_f;a)$.*

In particular, to determine whether or not the set $\{x\in\mathbb{R}\mid f(x)=0\}$ is empty can be computed by the above formula. □

In order to prove quantifier elimination we will need as *uniform* a way of making sense of what we might mean by a “Sturm sequence" for multivariate polynomials. We do this by taking a polynomial $p(\overline{y},x)$ and considering a certain *tree* of polynomials obtained in a similar way as the Sturm sequence applied to $p_{\overline{y}}(x)$ in such a way that we account for the possibility of the certain cancellations arising from substituting $\overline{y}$ into p. Before we do this, we will need the following proposition:

Proposition 5. The property of a polynomial $p_{\overline{y}}\in\mathbb{R}[\overline{y}][x]$ having degree equal to i (for $i\in\{-\infty,0,1,\cdots\}$) (in the variable x) is quantifier free definable.

Proof. (BPR) Express $p_{\overline{y}}(x)=\sum_{i\le n_p} a_i x^i$ with each $a_i=a_i(\overline{y})$ a term in the variables $\overline{y}$. Then let $\deg_x(p)=i\subseteq\mathbb{R}^k$ be the formula in the variables $\overline{y}$ given by

$$\text{“}\deg_x(p)=i\text{”} := \begin{cases} \bigwedge_{n_p>j>i} (a_j(\overline{y})=0)\wedge a_i\neq 0 & \mid 0\le i<n_p\in\mathcal{O} \\ a_j(\overline{y})\neq 0 & \mid i=n_p \\ \bigwedge_{j\le n_p} a_j(\overline{y})=0 & \mid i=-\infty \end{cases}$$

□

Given this first order characterization, it's clear that $\bigcup_{i\in\{-\infty,0,\cdots,n_p\}} (\deg_x(p)=i)(\mathbb{R}^k)$ partition $\mathbb{R}^k$ into finitely many sets, and on each set the degree of $p_{\overline{y}} \in \mathbb{R}[x]$ is constant.

Construction 2 (Pseudo-remainder Tree; BPR section 1.3)**.** Let $p,q \in \mathbb{R}[x]$ with $p = \sum_{i=1}^{n_p} a_i x^i$ and $q = \sum_{i=1}^{n_q} b_j x^j$. We assume that $p,q \in \mathcal{O}$ some subring of $\mathbb{R}$. The *signed pseudoremainder* $\mathrm{pr}(p,q)$ is the remainder of $b_{n_q}^d p$ in q under the Euclidean algorithm, with d the minimal even integer $\geq n_p - n_q + 1$.

For a polynomial $p(x) \in \mathbb{R}[x]$ let $\mathrm{lc}(p)$ be its leading coefficient.

If $q \in \mathbb{R}[x]$ let the i^{th} truncation of q, $\mathrm{tr}_i(q) = \sum_{j\leq i} b_j x^j$.

Given $q \in \mathbb{R}[\overline{y},x]$ let $Tr(q)$, the set of truncations of q, be defined recursively by

$$Tr(q) = \begin{cases} \{q_{\overline{y}}\} & \mathrm{lc}(q_{\overline{y}}) \in \mathcal{O} \\ \{q_{\overline{y}}\} \cup Tr(\mathrm{tr}_{\deg(q_{\overline{y}}-1)}(q_{\overline{y}}) & \mathrm{lc}(q_{\overline{y}})) \notin \mathcal{O} \end{cases}$$

The *pseudo-remainder tree* of $(p,q) \in \mathcal{O}[\overline{y}][x]$, $T_{\mathrm{pr}}(p,q)$ is the tree with nodes N each assigned a polynomial f_N as specified by the following data:

- The root R with $f_R = p$
- The children of p are nodes corresponding to elements $Tr(q)$.
- A node N is a leaf just in case the polynomial f_N is 0.
- If N is not a leaf then the children of N are precisely $Tr(-\mathrm{pr}(f_{pN}, f_N))$ where pN is the parent of N.

□

By construction, for each leaf $L \in T_{\mathrm{pr}}(p,q)$ there is a unique path B_L from the root to L. If $N \in B_L \setminus \{L\}$ then by construction there is a unique child of N, $c_L(N)$, also in B_L. For each leaf $L \in T_{\mathrm{pr}}(p,q)$ consider the first order formula $\theta_L(\overline{y})$ defined by

$$\theta_L(\overline{y}) := \deg_x(q) = \deg_x(f_{c_L(R)}) \wedge \bigwedge_{N\in B_L\setminus\{R\}} (\deg_x(-\mathrm{pr}(f_{p(N)}, f_N)) = \deg_x(f_{c_L(N)})$$

This tree plays the rôle that the Sturm sequence played in the univariate case in the proof of quantifier elimination. We require a bit more machinery before we describe the quantifier elimination algorithm:

Theorem 7 (Effective Quantifier Elimination)**.** *There is an algorithm for taking a given $\mathcal{L}_{\mathrm{RCF}}$-formula ψ and returning a quantifier free ϕ equivalent to it over $\mathbb{R}$.*

Proof. (BPR 2.4) We suppose at the outset that χ is in *prenex normal form*, meaning that $\chi := Q_1 x_1 \cdots Q_n x_n\, \phi(\overline{y};\overline{x})$ with ϕ quantifier free and $Q_i \in \{\exists, \forall\}$. Furthermore, we assume that ϕ is an $\mathcal{L}_{\mathrm{RCF}}$ formula of the form

$$[p(\overline{y},\overline{x}) = 0] \wedge \bigwedge_{i\in[\ell]} (\mathrm{sgn}(q_i(\overline{y},\overline{x})) = \sigma(q_i))$$

where $\sigma : [j] \to \{-1, 0, 1\}$ determines the required sign of $q_i(\overline{y}, \overline{x})$. Certainly all $\mathcal{L}_{\mathrm{RCF}}$ formulas are expressible as Boolean combinations of such sets.

First, it suffices to show effective quantifier elimination in the simple case of formulas with a single existential quantifier $\exists x\, \phi(\overline{y}, x)$. It suffices as $\forall z\, \psi(z; w)$ is logically equivalent to $\neg\exists z\, \neg\psi(z; w)$ and because we may successively iterate such an algorithm as follows to yield the desired quantifier free formula:

1. Start "inside out" on the formula $Q_n x_n\, \phi(\overline{y}, x_1, \cdots, x_{n-1}; x_n)$ and output an equivalent quantifier-free formula $\psi_{n-1}(\overline{y}; x_1, \cdots, x_{n-1})$.
2. Suppose we're given ψ_i as above, run the algorithm again on $Q_i\, \psi_i$ after ensuring that it is in the form prescribed above. Continuing in this way, $\psi_0(\overline{y})$ will be the desired quantifier-free formula.

Therefore we need only show the base case of the induction, the case of a single projection. For simplicity we assume that $\chi := \exists x\, \phi(\overline{y}, x)$ for ϕ of the form $[p(\overline{y}, \overline{x}) = 0]$ with $|\overline{y}| = k$. That is, we wish to show that the projection of an algebraic set is *semialgebraic*. The quantifier elimination algorithm for showing that the projection of a semialgebraic set is again semialgebraic is very similar to this one, but would take too long and distract from the rest of this paper. The full argument can be found in [BPR 2.4].

We fix some notation for useful objects in the algorithm:

- Let $Z = \phi(\mathbb{R}) \subseteq \mathbb{R}^{k+1}$ be the set defined by ϕ.
- For given $\overline{y} \in \mathbb{R}^k$ let $Z_{\overline{y}} = \{x \in \mathbb{R} \mid p_{\overline{y}}(x) = 0\}$.
- Let $C = \{(\overline{y}, x) \in \mathbb{R}^{k+1} \mid [p(\overline{y})](v) \not\equiv 0 \in \mathbb{R}[v]\}$ be the set of $(\overline{y}, x)$ such that $[p(\overline{y})](v) \in \mathbb{R}[v]$ is *not* the zero polyomial (this is *not* a typo!)
- Let $Z' = Z \cap C$.

Let Θ_L be the set of polynomials occuring in the formula $\theta_L(\overline{y})$ of a given leaf of $T_{\mathrm{pr}}(p, p')$ (explicitly, this is the set of pseudoremainders $-\mathrm{pr}(f_{p(N)}, f_N)$ and the polynomials f_N for $N \in B_L$), and let

$$\Theta = \bigcup_{L \in T_{\mathrm{pr}}(p,p') \text{ a leaf}} \Theta_L.$$

For $\sigma : \Theta \to \{-1, 0, 1\}$ a sign condition, the set of *realizations of* σ is the definable set

$$R^{\Theta}_{\sigma}(\mathbb{R}) = \{\overline{y} \in \mathbb{R}^k \mid \bigwedge_{f \in \Theta} [\mathrm{sgn}(f(\overline{y})) = \sigma(f)]\}$$

Note that for any family of polynomials Θ and any sign condition σ whatsoever, R^{Θ}_{σ} is definable by a quantifier free formula.

By Sturm's theorem, the effective method for calculating the number of real solutions of a given polynomial, and the observation that the possible sign conditions on Θ yield *all* possible sign conditions that can occur for substituted polynomials $p_{\overline{y}}$, we have that simply knowing σ

lets us computably determine whether or not an $\overline{y} \in \mathbb{R}^k$ satisfying R^{Θ}_{σ} corresponds to a $Z_{\overline{y}}$ that is empty or not. In summary, we have for all sign conditions σ that **either** $R^{\Theta}_{\sigma}(\mathbb{R}) \subseteq \{\overline{y} \mid Z_{\overline{y}} \neq \varnothing\}$ **or** $R^{\Theta}_{\sigma}(\mathbb{R}) \subseteq \{\overline{y} \mid Z_{\overline{y}} = \varnothing\}$.

Now let

$$\Sigma = \{\sigma : \Theta \to \{-1, 0, 1\} \mid \forall y \in R^{\Theta}_{\sigma}, Z_{\overline{y}} \neq \varnothing\}$$

Then clearly $\bigcup_{\sigma \in \Sigma} R^{\Theta}_{\sigma}$ defines $\pi(Z')$. This means that $\pi(Z')$ is quantifier-free definable as it is the finite union of the the R^{Θ}_{σ}, which were themselves quantifier-free definable.

On the other hand, $\pi(Z \setminus Z')$ is the projection of the set of $(\overline{y}, x) \in \mathbb{R}^{k+1}$ such that $p(\overline{y}, x) = 0$ with $p(\overline{(y)}, v) \equiv 0 \in \mathbb{R}[v]$. But then $\exists x([p_{\overline{y}}(x) = 0] \wedge Z_{\overline{y}} = \mathbb{R})$ is equivalent to saying that $p_{\overline{y}} = 0$ as a function, which is equivalent to the condition that the terms $a_i(\overline{(y)}) = 0$ for all i. This is manifestly a quantifier-free condition and so $\pi(Z \setminus Z')$ is quantifier-free definable.

Therefore $\pi(Z) = \pi(Z') \cup \pi(Z \setminus Z')$ is quantifier-free definable by the formula

$$\bigcup_{\sigma \in \Sigma} R^{\Theta}_{\sigma} \vee \bigwedge_{i \leq \deg(p)} [a_i(\overline{y}) = 0]$$

and we've produced an algorithm for *effective* quantifier elimination. □

Given this proof of quantifier elimination we can almost immediately conclude that the theory of $\mathbb{R}$ as an ordered field is in fact decidable.

Corollary 5 (Decidability of $\mathbb{R}$)**.** *Let ϕ be a sentence in $\mathcal{L}_{\mathrm{RCF}}$. Then there is an algorithm for determining whether or not $\mathbb{R} \vDash \phi$.*

Proof. By applying the quantifier elimination algorithm to ϕ, we may assume without loss of generality that ϕ is a quantifier-free sentence. But quantifier-free sentences are Boolean combinations of *terms* and are therefore Boolean combinations of diophantine equalities and inequalities applied to integers. Checking the truth or falsehood of such a sentence is a familiar exercise in elementary school arithmetic and does the job. □

As such, we've furnished a sketch of the methods used to prove effective quantifier elimination for and decidability of the ordered field theory of the real numbers, which we will use extensively in our applications to the electrical inverse problem.

4.2 *o*-minimality

The abstract model theoretic notion of *o-minimality* provides a way of extending the above results to theories in the incomputable realm. For instance, if one considers the language $\mathcal{L}'$ extending $\mathcal{L}_{\mathrm{RCF}}$ by adding a constant c_r for each $r \in \mathbb{R}$, we would have that $Th_{\mathcal{L}'}(\mathbb{R})$ is an *uncountable* set and hence cannot possibly be decidable. Nevertheless, one would not expect that adding names for constants would substantially alter the geometry of definables subsets of $\mathbb{R}^n$ from $\mathcal{L}$ to $\mathcal{L}'$. *o*-minimality confirms this intuition.

We begin by recalling a corollary of quantifier elimination of $\mathbb{R}$ in $\mathcal{L}_{\mathrm{RCF}}$:

Corollary 6. *Every definable subset $D \subseteq \mathbb{R}^1$ is a finite union of points and intervals.*

It is this property that motivates the definition of *o*-minimality.

Definition 13. Let $\mathcal{L} = (<, \cdots)$ be any language containing a symbol for a linear order and let $\mathcal{M}$ be an $\mathcal{L}$-structure. We say that $\mathcal{M}$ is *o-minimal* just in case every definable $D \subseteq \mathcal{M}^1$ is a finite union of points and intervals.

Hence $(\mathbb{R}, <, \cdot, +, -, 0, 1)$ provides a natural example of an *o*-minimal structure.

It turns out that the geometry of *arbitrary* definable sets $D \subseteq \mathcal{M}^n$ is rather simple if $\mathcal{M}$ is an *o*-minimal structure: a version of cell decomposition holds, there is a well-defined dimension theory, and all definable maps are *definably* piecewise trivial. Moreover, *o*-minimality guarantees certain *uniformities* when parametrizing definable sets by other definable sets. We recall some of the central results in the theory of *o*-minimal structures.

Definition 14 (van den Dries 2.2.3, 2.2.10)**.** We define the notion of a cell in an *o*-minimal structure $\mathcal{M}$ inductively:

- A point $m \in M$ is a 0-cell.
- An interval $(m_1, m_2) \subseteq M$ is a 1-cell.
- Given an i-cell $C \subseteq M^\ell$, we can construct cells in $M^{\ell+1}$ as follows:

  - If $f : C \to M$ is a definable function then the graph $\Gamma_f = \{(x, f(x)) \subseteq M^\ell \times M \mid x \in C\}$ is an i-cell.

  - If $f, g : C \to M$ are definable functions (or $f = \pm\infty$) with $f < g$ then the set

  $$\{(x, s) \mid x \in C, f(x) < s < g(x)\}$$

  is an $(i+1)$-cell

A cell decomposition of M^ℓ is a partition of $M^\ell = C_1 \cup \cdots \cup C_n$ where the C_i are disjoint cells satisfying the following inductive definition:

- A cell decomposition of M is a partition $M = C_1 \cup \cdots \cup C_n$ where the C_i are disjoint intervals or points.
- A cell decomposition of $M^\ell = C_1 \cup \cdots \cup C_n$ into cells such that $\pi(C_1) \cup \cdots \cup \pi(C_n)$ (where $\pi : M^n \to M^{\ell-1}$) is projection) is a cell decomposition.

A cell decomposition subordinate to a definable $S \subseteq M^\ell$ is a cell decomposition of M^ℓ such that each cell C_i occuring in the decomposition is such that either $C_i \subseteq S$ or $C_i \subseteq M^n \setminus S$. $\square$

The cell decomposition theorem for *o*-minimal structures states that

Theorem 8. *Let $\mathcal{M}$ be an o-minimal structure. Then for all definable $D \subseteq M^\ell$ there exists a cell decomposition of M^ℓ subordinate to D.*

A key lemma used in the proof of this result, interesting in its own right, is the *uniform finiteness lemma*. We call a family of definable sets $\{D_a \subseteq X \mid a \in Y\}$ a *uniformly definable family* where X and Y are definable and there is a set $D \subseteq X \times Y$ such that

$$D_a = \{x \in X \mid \mathcal{M} \models \phi_D(x, a) \text{ for the given } a \in Y\}.$$

A typical example of such a family is the set of fibers of a definable map: if $f : X \to Y$ is a definable map then $\{f^{-1}(a) \mid a \in Y\}$ is a uniformly definable family. The lemma says that

Lemma 9. *Let $\{D_a \subseteq X \mid a \in Y\}$ be a uniformly definable family. Then there exists an $M < \omega$ such that if $|D_a| > M$ then $|D_a|$ is infinite.*

This lemma *can* be used to reduce the proof of one key result in this paper to a triviality, although we will see that it yields somewhat weaker information than we get with computable methods.

Finally, there is a refinement of the cell-decomposition theorem that applies to definable maps called the *definable trivialization* theorem. It states roughly that given a surjective definable map $f : S \to T$ we can partition T into definable subsets T_i such that we can *definably* parametrize the fibers.

We begin with the definition of a *trivial* definable map.

Definition 15. Let $f : X \to Y$ be a definable map. f is said to be *definably trivial* if the following diagram commutes

$$\begin{tikzcd} X \arrow[r, "(f,\lambda)", "\sim"'] \arrow[d, "f"] & Y \times F \arrow[dl, "\pi_Y"] \\ Y & \end{tikzcd}$$

Where $F \subseteq \mathcal{M}^\ell$ and $\lambda : X \to F$ are definable.

Note, in particular, that $\lambda : X \to F$ is a surjective definable map, meaning that λ can be used to determine a first-order formula for generating the fibers. To see this, let $f^{-1}(t)$ be a fiber. Then as (f, λ) is a definable bijection, its inverse $(f, \lambda)^{-1}$ is also a definable bijection. But then as the trivialization diagram commutes, we have that we have that $\pi_Y^{-1}(t) = \{t\} \times F$ is bijective with $f^{-1}(t)$ and so $(f, \lambda)^{-1} : \{t\} \times F \to f^{-1}(t)$ is a definable bijection between fibers. Hence we can view the points of Y as definably (which, in the case of the real ordered field $\mathbb{R}$, means *semialgebraically*) parametrizing a family of fibers. The trivialization theorem says that there's always a definable way to reduce to this case:

Theorem 10 (van den Dries 9.1.2)**.** *Suppose that $X \subseteq \mathcal{M}^n$, $Y \subseteq \mathcal{M}^m$ are definable sets and $f : X \to Y$ a continuous definable map. Then there is a partition of Y into cells, $Y = C_1 \cup \cdots \cup C_\ell$, such that $f : f^{-1}(C_i) \to C_i$ is a definably trivial map.*

Chapter 5

Applications to n-to-1 Graphs

5.1 Dichotomy Theorems and a Hierarchy of Graphs

In this section we prove two dichotomy theorems, which will be very important in our study of both the local and global inverse problems.

Theorem 11 (Dichotomy Theorem for Inverse Problems)**.** *For a given inverse problem (G, Λ), we have that either $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$ or $|L_G^{-1}(\Lambda)| = M < \omega$ for some M.*

Proof. It is a straightforward model-theoretic fact that if an $\mathcal{L}$-structure $\mathcal{M}$ is o-minimal then $\mathcal{M}$ is o-minimal as an $\mathcal{L}' = \mathcal{L} \cup \{c_i\}$ structure where the c_i are constants when interpreted by $c_i^{\mathcal{M}}$. In particular, given an inverse problem (G, Λ) we may add to $\mathcal{L}_{\mathrm{RCF}}$ the constants $\{\lambda_{ij}\}$ for the entries of Λ. Then the set $L_G^{-1}(\Lambda) \subseteq (\mathbb{R}^+)^{|E|}$ is $\mathcal{L}_{\mathrm{RCF}}(\{\lambda_{ij}\})$-definable. Therefore we may decompose $L_G^{-1}(\Lambda) = C_1 \cup \cdots \cup C_\ell$ into a finite union of cells. If $\dim C_i > 0$ for some i then, by the definition of cells, C_i has cardinality $2^{\aleph_0}$. If, to the contrary $\dim C_i = 0$ for all i then $L_G^{-1}(\Lambda)$ is a finite union of points. Hence either $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$ or $|L_G^{-1}(\Lambda)| = M < \omega$. $\square$

Theorem 12 (Uniform Dichotomy Theorem for Graphs)**.** *For a given graph G, we have that either $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$ for some A or that there exists a constant natural number M such that for all Λ, $|L_G^{-1}(\Lambda)| \leq M$.*

Proof. I will present two proofs of this theorem: one using the fact that definable functions in o-minimal structures are piecewise fibrations, and the other using the fact that o-minimality is a property preserved under elementary equivalence. Both proofs are included as they appeal to different intuitions.

Let $X = L_G((\mathbb{R}^+)^{|E|})$. As L_G is a definable map applied to the definable space $(\mathbb{R}^+)^{|E|}$, X is definable and the local trivialization theorem of o-minimality guarantees that we can partition $X = X_1 \cup \cdots \cup X_n$ into finitely many definable sets such that the maps $L_G \restriction_{L_G^{-1}(X_i)}$ is a fibration. More precisely, the following diagram commutes:

$$\begin{tikzcd} L_G^{-1}X_i \arrow[r, "(L,h_i)", "\sim"'] \arrow[d, "L_G"] & X_i \times F_i \arrow[dl, "\pi_{X_i}"] \\ X_i & \end{tikzcd}$$

where F_i is a definable representative fiber of X_i under L_G and $h_i : L_G^{-1}X_i \to F_i$ is a definable map such that $(L_G, h_i) : L_G^{-1}X_i \to X_i \times F_i$ is a definable bijection. But then for each $\Lambda \in X_i$ we have that $|L_G^{-1}(\Lambda)| = |\{\Lambda\} \times F_i| = |F_i|$, so that the cardinality of the fiber depends only on i. As we decomposed X into finitely many such X_i, we have that the set $\{\kappa \mid \exists\Lambda\ |L_G^{-1}(\Lambda)| = \kappa\}$ is finite. Furthermore, by the Dichotomy Theorem for Inverse Problems we have that the κ appearing above are either $\kappa = n < \omega$ or $\kappa = 2^{\aleph_0}$. Now if $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$ for some Λ then we satisfy the hypothesis of the theorem and we are done. Otherwise, let $M = \max\{\kappa \mid \exists\Lambda\ |L_G^{-1}(\Lambda)| = \kappa\}$. Then clearly for all Λ we have that $|L_G^{-1}(\Lambda)| \leq M$ and we are done.

For the more clearly model-theoretic proof, suppose for the contrary that $|L_G^{-1}(\Lambda)| < \omega$ for all Λ. Then by Uniform Finiteness for *o*-minimal structures we may immediately conclude that there is some $M < \omega$ such that $|L_G^{-1}(\Lambda)| < M$. ☐

5.2 Finding One's Place in the Hierarchy

Using the theorem in the preceding section we are able to show that the place of a finite partitioned graph G in the hierarchy of the electrical inverse problem is in fact *computable*. To this end we first show that we can translate the problem into questions about first-order statements in the theory of real-closed fields and then apply the decidability algorithm for real-closed fields to a sequence of first order formulas. The main worry *a priori* is that the algorithm described would not halt. The dichotomy theorems, however, guarantee that this does in fact happen.

To begin with, we show that the properties that arise in the hierarchy are, in fact, first order definable.

Proposition 6. Let G be a partitioned graph, L_G its measurement map, definable over a subset $C \subseteq \mathbb{R}$. Then the following properties are first-order expressible in $\mathcal{L}_{\mathrm{RCF}}(C)$:

1. G is strongly $2^{\aleph_0}$-to-1.
2. G is weakly n-to-1 for a given n.
3. G is strongly n-to-1 for a given n.

Proof. We prove by cases.

1. We first show that for a *fixed* inverse problem(G, Λ), being $2^{\aleph_0}$-to-1 is first-order definable. To this end we invoke the dichotomy theorem and cell decomposition of sets definable in an *o*-minimal structure. Certainly the set $L_G^{-1}(\Lambda) = \{\overline{\gamma} \mid L_G(\overline{\gamma}) = \Lambda\}$ is first-order definable with finitely many parameters. Now, suppose that $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$. Then one of the coordinate projections $\pi_j(L_G^{-1}(\Lambda)$ onto the coordinates not equal to j, $(1, \cdots, \hat{j}, \cdots, n)$, must also

be of size $2^{\aleph_0}$, for otherwise the fibers would all be finite and $|L_G^{-1}(\Lambda)| \leq \prod |\pi_j(L_G^{-1}(\Lambda))| < \aleph_0$, a contradiction. Repeating this process, we get to a step where at least one sequence of projections is a *definable* subset S of $\mathbb{R}^1$ of size $2^{\aleph_0}$. Then as $\mathbb{R}$ is o-minimal, S necessarily contains an interval (a, b). Conversely, if such a series of projections exist, by taking inverse images we must have that $2^{\aleph_0} \leq |L_G^{-1}(\Lambda)| \leq 2^{\aleph_0}$, whence $|L_G^{-1}(\Lambda)| = 2^{\aleph_0}$. Hence the existence of such a filtration of projections containing an interval is equivalent to the inverse problem having $2^{\aleph_0}$ solutions. To check that G is strongly $2^{\aleph_0}$-to-1, we merely need ask for the *existence* of such a Λ, which is again first order.

Now we formalize this. Let ψ_G be the formula

$$\exists(\Lambda \in M_{|\partial V|}(\mathbb{R}^+))\exists(a, b \in \mathbb{R}^+)$$
$$\left[\forall x(a < x < b \to \bigvee_{\substack{S_i \subseteq [1,\cdots,n] \\ |S_i| = n-1}} \exists(\gamma_{s_{i1}}, \cdots, \gamma_{s_{i(n-1)}}) L_G(\gamma_{s_{i1}}, \cdots, x, \cdots, \gamma_{s_{i(n-1)}}) = \Lambda)\right]$$

where we insert x into the sequence of γ_{ij}'s (for a fixed i) in the m_i^{th} place where $m_i \notin S_i$ place. By inspection, ψ_G is the first order sentence corresponding to the above construction. Hence, we have that the graph G is strongly $2^{\aleph_0}$-to-1 if and only if ψ_G holds in $\mathbb{R}$.

2. Let $\chi_n(G)$ be the formula

$$\exists(\Lambda \in M_{\partial V}(\mathbb{R}^+))\exists\overline{\gamma}_1, \cdots \overline{\gamma}_n \in (\mathbb{R}^+)^{|E|} \bigwedge_{i \neq j} \overline{\gamma}_i \neq \overline{\gamma}_j \wedge \bigwedge_{i \leq n} L_G(\overline{\gamma}_i) = \Lambda$$

for $n > 0$. Then $\chi_n(G)$ says precisely that there is an A with at least n distinct solutions to the electrical inverse problem on G, which exactly says that G is weakly n-to-1.

3. Let $\phi_n(G)$ be the formula $\chi_n(G) \wedge \neg\chi_{n+1}(G)$ for $n > 0$. Then $\phi_n(G)$ expresses that G is strictly n-to-1.

□

Remark 4. I do not know of a topological proof that if G is strongly $2^{\aleph_0}$-to-1 then there is an interval in some projection of the domain on which $L_G(\gamma_{s_{i1}}, \cdots, x, \cdots, \gamma_{s_{i(n-1)}}) = \Lambda$. The main difficulty is that there exist uncountable nowhere dense subsets of $\mathbb{R}$, namely Cantor sets. Hence this proof relies on something much stronger than continuity- it relies on L_G being definable in an o-minimal structure.

By essentially the same argument as for the global case, the analogous properties for a fixed inverse problem (G, Λ) are first order:

Proposition 7. Let G be a partitioned graph and suppose that L_G and $\Lambda \in M_n(\mathbb{R})$ are definable over a subset $C \subseteq \mathbb{R}$. Then the following properties are first-order expressible in $\mathcal{L}_{\text{RCF}}(C)$:

1. (G, Λ) is $2^{\aleph_0}$-to-1.

2. (G, Λ) is n-to-1 for a given n.

3. $\dim_{\mathbb{R}}(L_G^{-1}) = n$

Proof. Clearly $\exists \overline{\gamma} L_G(\overline{\gamma}) = \Lambda$ is a C-definable first-order sentence. Repeating the argument of the proposition above, we see that the statement "$L_G(\gamma) = \Lambda$ has $2^{\aleph_0}$ solutions" is also first-order expressible in $\mathcal{L}_{\mathrm{RCF}}(C)$; defined by a formula $\psi_{(G,\Lambda)}$, as are the statements "$L_G(\gamma) = \Lambda$ has n distinct solutions" (let this be $\chi_{n,(G,\Lambda)}$)) and "$L_G(\gamma) = \Lambda$ has exactly n-distinct solutions" $(\phi_{n,(G,A)})$.

For computing dimension, we adapt an exercise in [van den Dries, 4.1.17.1] and use the fact that the real dimension $\dim_{\mathbb{R}}(L_G^{-1}(\Lambda)) \geq n$ just in case there is some subset of n coordinates such that the projection map $\pi : L_G^{-1}(\Lambda) \to \mathbb{R}^n$ given by mapping a point $\gamma \in L_G^{-1}(\Lambda) \mapsto (\gamma_{i_1}, \cdots, \gamma_{i_n})$ has non-empty interior, which is a first-order property. Then $\dim_{\mathbb{R}}(L_G^{-1}) = n$ just in case $[\dim_{\mathbb{R}}(L_G^{-1}(\Lambda)) \geq n] \wedge \neg[\dim_{\mathbb{R}}(L_G^{-1}(\Lambda)) \geq n+1]$. Hence the property of having given dimension n is first-order expressible. □

We now provide sufficient conditions for electrical inverse problems to be $\varnothing$-definable, allowing us to apply the decision procedures of the preceding chapter to the problems at hand:

Proposition 8. 1. If $\Lambda \in M_{|\partial V|}((\mathbb{R}^{alg})$ then Λ is $\varnothing$-definable in $\mathcal{L}_{\mathrm{RCF}}$.

2. For all electrical networks G, L_G is $\varnothing$ definable.

Proof. 1. It suffices to show that each entry Λ_{ij} is $\varnothing$-definable. As $\Lambda_{ij} \in \mathbb{R}^{alg}$, there is some polynomial $p_{ij}(x) \in \mathbb{Q}[x]$ such that $p_{ij}(\Lambda_{ij}) = 0$. This specifies Λ_{ij} up to at most $\deg(p_{ij})$-many other reals. Assume that Λ_{ij} is the k^{th} real solution of p_{ij}, ordered by the induced ordering on $\mathbb{R}$. Let $\xi_{ij}(x) :=$ "$p_{ij}(x) = 0$". The formula

$$\rho_{ij}^k(x) := \exists y_1 \cdots \exists y_{k-1} \left(\bigwedge_{\ell < k} (x > y_\ell \wedge p_{ij}(y_\ell) = 0) \bigwedge_{k \neq k'} (y_k \neq y_{k'}) \wedge p_{ij}(x) = 0 \right) \wedge \\ \nexists y_k([y_{k-1} < y_k < x] \wedge p_{ij}(y_k) = 0)$$

says that x is the k^{th} real solution to p_{ij}. Then Λ_{ij} is the unique element specified by ρ_{ij}^k, and hence each entry of Λ is $\varnothing$-definable, whence Λ is $\varnothing$-definable in $\mathcal{L}_{\mathrm{RCF}}$.

2. Immediate from the formula $L_G(\gamma) = A_\gamma - B_\gamma^T D_\gamma^{-1} B_\gamma$ as described in Chapter 2.

□

If we implemented this algorithm in an actual computer, we would require that the entries of Λ were presented in the form of ρ_{ij}^k above, which can certainly contribute quite a bit of complexity to the actual representation of the formulas we consider in the following algorithms.

Now that we know that all of the relevant data is in fact *first-order* expressible over $\varnothing$, we are in prime position to give algorithms computing quantities related to fiber cardinals, ΨG. We insist that our inverse problem Λ has algebraic-real entries for the purpose of computability. If you were ever to calculate this for a physical electrical network, you would use only rational approximations of boundary measurements anyways, so allowing for real algebraic coefficients gives us slightly *more* generality than we actually need.

Theorem 13. *The problems* (1) *and* (2) *below are answerable when* Λ *has real-algebraic coordinates.*

1. *Given a particular inverse problem* (G, Λ) *with* $\Lambda \in M_{|\partial V|}((\mathbb{R}^{alg})$ *on a graph* G *with* L_G *being* $\varnothing$*-definable in* $\mathcal{L}_{\mathrm{RCF}}$*, how many solutions does it have? If it is not* n*-to-*1*, what is the dimension of* $L_G^{-1}(\Lambda)$*?*

2. *Given a graph* G*, where in the hierarchy of* n*-to-*1 *graphs does it fall?*

Proof. This is a simple application of the model theoretic black-boxes that we've proven.

For problem (1), let (G, Λ) be an inverse problem with $\Lambda \in M_{|\partial V|}((\mathbb{R})^{alg})$. As $\mathbb{R}$ is o-minimal, algebraic closure is the same as definable closure, and so as $\Lambda \in M_{|\partial V|}((\mathbb{R}^{alg})$, all its coefficients a_{ij} live in $(\mathbb{R}^{alg})$ and are therefore $\varnothing$-definable in the language of ordered fields. Hence Λ is $\varnothing$-definable in $\mathbb{R}$ and so all the properties "$L_G(\gamma) = \Lambda$ has $2^{\aleph_0}$ solutions, " as is "$L_G(\gamma) = \Lambda$ has exactly n distinct solutions." Denote the former formula by $\psi_{(G,\Lambda)}$ and the latter family of formulas by $\phi_{n,(G,\Lambda)}$ for each n (including $n = 0$). Using the decision procedure for $\varnothing$-definable first-order formulas, consider the following algorithm:

1. Run the decision algorithm on $\psi_{(G,\Lambda)}$. If true, run the decision algorithm on the formulas $\{\dim(L_G^{-1}(\Lambda)) = n\}_{n\in\omega}$. This returns that $\mathbb{R} \vDash [\dim_{\mathbb{R}}(L_G^{-1}(\Lambda)) = m]$ for some number $m < E + 1$ and we know that (G, Λ) is $2^{\aleph_0}$-to-1 and of real dimension m.

2. If $\mathbb{R} \vDash \neg\psi_{(G,\Lambda)}$, then run the decision procedure on the family of first order formulas $\{\phi_{n,(G,\Lambda)}\}_{n\in\omega}$. Terminate when the decision procedure determines that $\mathbb{R} \vDash \phi_{M,(G,\Lambda)}$. Then (G, Λ) is M-to-1.

By the Dichotomy Theorem for Inverse Problems, we have that either the inverse problem (G, Λ) has $2^{\aleph_0}$ solutions or it has $M < \omega$ solutions. Since the decidability algorithm for sentences in RCF terminates and since we are simply applying the decidability algorithm to finitely many formulas, the algorithm above must terminate after finitely many steps.

Now we consider Problem (2). The algorithm is almost identical to the one for a fixed inverse problem (G, Λ). It is:

1. Run the decision algorithm on ψ_G. If true, terminate. G is then strongly $2^{\aleph_0}$-to-1.

2. If $\mathbb{R} \vDash \neg\psi_G$, then run the decision procedure on the family of first order formulas $\{\phi_{n,G}\}_{n\in\omega}$. Terminate when the decision procedure determines that $\mathbb{R} \vDash \phi_{M,G}$. Then G is strongly M-to-1.

This time, by the Uniform Dichotomy Theorem for Graphs, we have that a graph G is either strongly $2^{\aleph_0}$-to-1 or strongly M-to-1 for some $M < \omega$. Again, the algorithm for sentences in RCF terminates after finitely many steps and, by the Uniform Dichotomy Theorem, we need only apply this algorithm to finitely many sentences before it terminates. □

5.3 Finer Analysis

The previous section gives an algorithm for computing exactly how many solutions a given (algebraically presented) inverse problem has and, moreover, for where in the hierarchy a given graph falls. In this section we very briefly sketch how more sophisticated algorithms can yield finer information about the solution spaces.

First off, we cite a very powerful result:

Theorem 14. ***Find Citation*** *Given any first-order expressed semialgebraic map $f : X \to Y$ we can computably produce a cell decomposition of $Y = C_1 \cup \cdots \cup C_n$ on which $f : f^{-1}(C_i) \to C_i$ is a piecewise trivial map.*

Proof. $\square$

Using this theorem it is clear that we can *computably* partition $M_n(\mathbb{R})$ into finitely many nice cells C_i on which L_G is a locally trivial map. In particular, on each of these cells the cardinality of fibers will be constant. This allows us to computably break up $(\mathbb{R}^+)^{|E|}$ into finitely many sets on which we know the behavior of L_G. Moreover, we will be able to computably produce on each $L_G^{-1}(C_i)$ an implicitly defined bijection $(L_G, h) : L_G^{-1}(C_i) \to C_i \times F_i$, giving us an implicit parametrization of the fibers as we vary Λ along C_i.

In the local setting we consider the case where (G, Λ) is n-to-1. In this former case we are able to computably produce polynomials for the coordinates are solutions as follows:

Start with the definable set $\{\gamma \mid L_G(\gamma) = \Lambda\} \subseteq \mathbb{R}^{|E|}$. If (G, Λ) is finite-to-1, then applying the quantifier elimination algorithm to each set

$$S_i = \{x \in \mathbb{R} \mid \exists(\gamma_1, \cdots, \gamma_{i-1}, \gamma_{i+1}, \gamma_{|E|}) L_G(\gamma_1, \cdots, \gamma_{i-1}, x, \gamma_{i+1}, \gamma_{|E|}) = \Lambda$$

we end up with a quantifier-free $\mathcal{L}_{\mathrm{RCF}}$ formula defining S_i of the form

$$\eta_i = \bigvee_{j_i \in [k_i]} [p_{j_i}(x) = 0] \wedge \bigvee_{j' \in [k']} [q_{j'}(x) \square_i 0]$$

for $\square_i \in \{\neq, >, <\}$. Then let $F_i = \prod_{j_i \in [k_i]} p_{j_i}$. Then the solutions of $L_G^{-1}(\Lambda)$ are contained in the set

$$\{(\gamma_1, \cdots, \gamma_n) \mid \bigwedge_{i \in [|E|]} [F_i(\gamma_i) = 0]\}$$

Then, by approximating solutions to these polynomials (by using Newton's method or any other) and ruling out combinations of solutions of the F_i's that aren't solutions of (G, Λ) (by using the integral form of the Taylor Remainder theorem, for instance) we are able to computably approximations to *all* solutions to (G, Λ).

5.4 Conjectures and Further Research

While this thesis provides answers to many open problems in the study of inverse problems on electrical networks, the *asymptotic* theory of inverse problems is still very mysterious. Recall that the *fiber cardinals* of a electrical network with boundary G, ΨG, which was defined as

$$\Psi G = \{\kappa \mid \exists \Lambda \in M_{|\partial G|}(\mathbb{R}) \; |(G, \Lambda)| = \kappa\}.$$

One of the key results in this paper was that for all finite electrical networks with boundary, ΨG was finite. Moreover, for a given network G we are able to compute ΨG by producing a computable cell decomposition of $M_{|\partial G|}(\mathbb{R})$ making L_G *cellwise* a fiber bundle and then running our decision procedure to find the cardinality of each fiber. However, this sheds no light on the dependence of ΨG on simple numerical quantities associated to G, such as $|E_G|$, $|V_G|$, and $|\partial G|$. In my mind the most pressing question in this regard is

Question 15. *(Strong Form) Let $\kappa_G = \max\{\kappa \mid \kappa \in \Psi G \wedge \kappa \neq 2^{\aleph_0}\}$ be the largest finite cardinal in the fiber cardinal of G. Given any graph-theoretic quantity (or tuple of quantities) ν_G associated to finite connected electrical networks with boundary G, is there an asymptotic relationship between κ_G such that as $\nu_G \to \infty$, $\kappa_G \sim \theta(\nu_G)$ for some explicit function θ?*

(Weak Form) Given a graph-theoretic quantity ν_G associated to ***some subclass*** *$\mathcal{G}$ of finite connected electrical networks with boundary G, is there an asymptotic relationship between κ_G such that as $\nu_G \to \infty$, $\kappa_G \sim \theta(\nu_G)$ for some explicit function θ?* □

It would be of great interest to answer this in the special case of $\nu_G = |E_G|$ or $\nu_G = |V_G|$, since all known constructions of n-to-1 graphs require high numbers of edges and vertices relative to n.

In a similar vein, it is unclear whether or not we should *generically* expect strictly finite-to-1 behavior in a graph or whether we should expect continuum-to-1 behavior. An asymptotic study of this would be very enlightening, for it would help get a better picture of exactly along what dividing lines ground the dichotomy theorem proved in this thesis. One possible framing of this question is as follows:

Question 16. *(Strong Form) Let ν_G be a given graph-theoretic quantity such that for all n there are only finitely many finite connected electrical networks with boundary G with $\nu_G \leq n$, and let κ be a cardinal number. Is there any asymptotic estimate for the probability that a given G of bounded ν_G admits κ-to-1 behavior, in the sense that*

$$\frac{|\{G \mid G \text{ is strongly } \kappa\text{-to-}1 \text{ and } \nu_G \leq n\}|}{|\{G \mid \nu_G \leq n\}|} \sim \theta(n)$$

for some explicit $\theta : \mathbb{R} \to [0,1]$ as $n \to \infty$?

(Weak Form) Given any graph-theoretic quantity ν_G associated to ***some subclass*** *$\mathcal{G}$ of finite connected electrical networks with boundary G, such that for all n there are only finitely many finite connected electrical networks with boundary $G \in \mathcal{G}$ with $\nu_G \leq n$. Is there an explicit $\theta : \mathbb{R} \to [0,1]$ such that, asymptotically,*

$$\frac{|\{G \in \mathcal{G} \mid G \text{ is strongly } \kappa\text{-to-}1 \text{ and } \nu_G \leq n\}|}{|\{G \in \mathcal{G} \mid \nu_G \leq n\}|} \sim \theta(n)$$

□

In the previous section I presented a method of producing approximations to *all* solutions for a finite-to-1 inverse problem (G, Λ), not *exact* solutions. I expect this to be optimal:

Conjecture 17. *There is no algorithm applying to all inverse problems (G, Λ) that returns for a given finite-to-1 problem (G, Λ) all of its solutions exactly, presented in radicals.* □

The intuition behind this is that, for Galois theoretic reasons, we cannot solve a general polynomial of degree 5 or greater. The algorithm given above to compute minimal polynomials for the coordinates of solutions definitely will return univariate polynomials of high degree, so it would be quite shocking if one were to find a way to avoid this when computing all solutions exactly. I believe that one should be able to easily code the Galois-theoretic problem directly into some inverse problem, thus proving the result, but I have yet to do so.

The algorithm that I gave for deciding where in the κ-to-1 hierarchy both a given inverse problem (G, Λ) and a graph G land in relies on repeated iterations of a very time-consuming decision algorithm. In fact, it has been proven by Davenport and Heintz that any decision procedure for $\langle \mathbb{R}, +, -, \cdot, <, 0, 1 \rangle$ is of at least 2^{2^n}-complexity; Ben-Or, Kozen, and Reif showed that this bound is sharp by exhibiting a decision algorithm of doubly-exponential complexity. In many ways, this is the price of approaching the problem in full generality. It's very feasible that by restricting attention to a class of graphs satisfying certain nice graph theoretic properties we could construct far faster algorithms in these cases. With this in mind we can ask the following open-ended question:

Question 18. *Which classes of graphs can we find algorithms of low computational complexity for enumerating and, if applicable, approximating solutions of inverse problems in both the local and global settings?* □

The methods used to abstractly study the solution spaces of graphs crucially relied on *o*-minimality and the ability to rephrase the problem in terms of finite-dimensional subsets of $\mathbb{R}^N$ for some large enough N. This renders these methods useless for studying *countable* electrical networks. As such, another open-ended problem would be to find an appropriate model-theoretic framework for studying the enumeration problem in the case of coubtable electrical networks.

Question 19. *Is there a model-theoretic framework amenable to approaching the problems considered in this thesis for countable electrical networks?* □

Finally, it would be lovely to see an account of some other discrete inverse problem investigated using methods of *o*-minimality that go beyond the simple case of the structure $\langle \mathbb{R}, +, -, \cdot, <, 0, 1 \rangle$. This leads to the following suggestive, yet very ill-posed, question:

Question 20. *Is o-minimality an appropriate framework for studying the discrete inverse problems that occur in the practical world?* □

Chapter 6

Appendix: The Characterization of Response Matrices

One of the major problems in the program on electrical inverse problems has been to give (for a fixed but arbitrary electrical network G) necessary and sufficient criteria for a matrix $\Lambda \in M_{\partial G}(\mathbb{R})$ to be the response matrix of some set of conductivities $\overline{\gamma} \in (\mathbb{R}^+)^{|E|}$. In fact, the effective quantifier elimination for the real field yields such a characterization instantly.

Theorem 21. *Given a finite electrical network G, there is an effective characterization of its response matrices by semialgebraic constraints.*

Proof. Note that the set of $\Lambda \in M_{\partial G}(\mathbb{R})$ that are the response matrices of some set of conductivities $\overline{\gamma} \in (\mathbb{R}^+)^{|E|}$ is precisely the image of $(\mathbb{R}^+)^{|E|}$ under the measurement map L_G. This set is described explicitly by the definable set

$$\{\Lambda \in M_{\partial G}(\mathbb{R}) \mid \exists \overline{\gamma}\, L_G(\overline{\gamma}) = \Lambda\}$$

The $\mathcal{L}_{\mathrm{RCF}}$ formula "$\exists \overline{\gamma}\, L_G(\overline{\gamma}) = \Lambda$" is defined over $\varnothing$, and so we are free to apply *effective* quantifier elimination to yields an equivalent quantifier-free formula ϕ in $\mathcal{L}_{\mathrm{RCF}}$. This is simply a finite Boolean combination of explicit semialgebraic conditions, and so the result is proven. $\square$

This theorem in principle allows one to test whether an arbitrary Λ is in the image of the measurement map of G or not simply by checking the truth of certain polynomial equalities and inequalities *dependent only on G*.

Chapter 7

Bibliography

1. Basu, Saugata; Richard Pollack, Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. 2^{nd} Edition. Springer 2006.

2. Ben-Or, Michael; Dexter Kozen, and John Reif, *The complexity of elementary algebra and geometry*, Journal of Computer and Systems Sciences 32 (1986), no. 2, pp. 251 - 264.

3. Chang, C.C. and Jerome Keisler. *Model Theory: Third Edition*. Dover 2012.

4. Coste, Michel. *An Introduction to Semialgebraic Geometry*. 2002.
<perso.univ-rennes1.fr/michel.coste/polyens/SAG.pdf>

5. Curtis, Ingerman, and Morrow. "Circular Planar Graphs and Resistor Networks." Linear Algebra and its Applications 283 (1998), 115-150.

6. Curtis, Edward B. and James Morrow. *Inverse Problems for Electrical Networks*. World Scientific. 2000.

7. Dale, Reid. "*o*-minimality and Uniformity in $n-1$ Graphs." Memo. 2013.
<http://www.math.washington.edu/∼reu/papers/current/reid/o-minimality.pdf>

8. Davenport, James H.; Heintz, Joos (1988). *Real quantifier elimination is doubly exponential*. J. Symb. Comput. 5 (1-2): 29 - 35.

9. French, Jenny and Shen Pan. "2^n To 1 Graphs." 2004.
<http://www.math.washington.edu/∼reu/papers/2004/jenny/firstpaper.pdf>

10. Giansiracusa, Jeffrey. "The Map L and Partial Recovery in Circular Planar Non-Critical Networks." 2003.
<http://www.math.washington.edu/∼reu/papers/2002/giansiricusa/giansiricusa_diffeo.pdf>

11. Klumb, Chad. "Discrete Unsolvability for the Inverse Problems for Electrical Networks." 2011.
<http://www.math.washington.edu/∼morrow/papers/chad-thesis.pdf>

12. Marker, Dave. *Model Theory: An Introduction*. Springer 2003.
13. van den Dries, Lou. *Tame Topology and o-minimal Structures*. Cambridge University Press 1998.