Discrete Inverse Transport

Stanley Snelson

August 15, 2008

1 Introduction

The inverse problem for the stationary transport equation on a bounded domain in $\mathbb{R}^n$ has been well-studied; see for example [1] for $n \geq 3$ and [4] for $n = 2$. In this paper, we develop a problem on directed graphs with boundary that is a discrete analogue to the inverse transport problem. The discrete transport problem was first considered in [3] in a form slightly different from the problem we are studying.

We recall the formulation of the continuous problem. Let $X \subset \mathbb{R}^n$ be a bounded domain with C^1-boundary ∂X, and let $V \subset \mathbb{R}^n$ be open. The function $f(x, v)$, which represents the density of particle transport at a point $x \in X$ travelling with velocity $v \in V$, satisfies the *stationary linear transport equation,* given by

$$-v \cdot \nabla_x f(x, v) - \sigma_a(x, v) f(x, v) + \int_V k(x, v', v) f(x, v') dv' = 0, \qquad (1)$$

at every $(x, v) \in X \times V$. The coefficient $\sigma_a(x, v)$ represents the density of particles being aborbed into the surrounding medium, and the scattering kernel $k(x, v', v)$ represents the particles changing direction from v' to v at x. Let $n(x)$ be the outward unit normal to ∂X at $x \in \partial X$, and define $\Gamma_\pm = \{(x, v) \in \partial X \times V : \pm n(x) \cdot v > 0\}$. The boundary value problem, or forward problem, is to find f which satisfies (1) and such that

$$f|_{\Gamma_-} = f_-, \qquad (2)$$

where f_- is a given function defined on Γ_-. Under certain conditions on σ_a, k, and X (for the exact conditions, see [1] and [4]), the forward problem

(1)-(2) is uniquely solvable for $f(x, v)$. Let f_+ be the restriction of f to Γ_+, called the outgoing flux. On the boundary of our domain, we can measure the outgoing flux generated by some incoming flux f_-. The *albedo operator* $\mathcal{A}$, defined by

$$\mathcal{A} : f_- \mapsto f_+ = f|_{\Gamma_+}$$

encapsulates all the information that can be obtained from such measurements. The inverse problem is: Does $\mathcal{A}$ determine σ_a and k? It is proved in [1] that if $\sigma_a(x, v) = \sigma_a(x)$, i.e. the absorption depends only on position, then $\mathcal{A}$ determines σ_a and k for for $n \geq 3$. For $n = 2$, it is proved in [4] that if $\sigma_a = \sigma_a(x)$ and k is small in the appropriate norm, then $\mathcal{A}$ determines σ_a and k.

Our goal is to formulate a discrete version of the inverse transport problem, and if possible to derive analogous results. Instead of a region X, we imagine particles moving along the edges of a directed graph G. A possible physical interpretation of this discrete model is traffic flow, where the edges of G represent streets and vertices represent intersections. We may interpret f as a density of cars, scattering as cars changing streets, and absorption as cars being parked. However, our primary focus is the mathematical problem without a specific application in mind.

The paper is structured as follows. In Section 2, we derive a version of the transport equation for particles travelling through a discrete medium. The parameters σ_a and k, which depend on the medium, still represent absorption and scattering of particles respectively. In Section 3, we show that the forward problem as posed in Section 2 is uniquely solvable as long as the absorption dominates the scattering at each edge. This unique solution gives us access to the (linear) albedo operator A, which maps particle flux entering the medium to the measured outgoing flux of particles. The inverse problem, explored in Section 4, is to reconstruct σ_a and k from A, which encapsulates the information available from boundary measurements. This is possible only if the map

$$S : (\sigma_a, k) \mapsto A_{\sigma_a,k}$$

is one-to-one, where $A_{\sigma_a,k}$ is the albedo operator for a transport network with absorption coefficient σ_a and scattering kernel k. The inverse problem is highly nonlinear, and the main results of this paper deal with the linearization of the problem. Namely, in Section 4.2 we characterize the differential $D_{(\sigma_a,k)}$

of S and give conditions which ensure $D_{(\sigma_a,k)}S$ is injective. In Section 5, we discuss possible future work on this problem and related ones.

The version of the discrete forward problem outlined in [3] was shown to have a unique solution. To the author's knowledge, the discrete inverse transport problem has not been previously studied.

2 The Discrete Transport Equation

Consider a symmetric digraph with boundary $G = (V, \partial V, E)$. V is the set of vertices, E is the set of edges, and ∂V is a nonempty subset of V whose elements are designated boundary vertices. The complement in V of ∂V is denoted $\operatorname{Int} V$, the set of interior vertices. Here, "symmetric" means that if a directed edge exists from p to q, then there is a directed edge from q to p. We assume G is finite and connected. Also, we assume that the in-degree (and hence also the out-degree) of each interior vertex is even, and that each boundary vertex has in-degree and out-degree equal to one. We can then consider edges incident to ∂V as *boundary edges* and denote the set of boundary edges ∂E. Since for each boundary vertex p, there is one directed edge originating at p and one directed edge ending at p, we can further divide ∂E into the set of *incoming boundary edges,* denoted ∂E_-, and the set of *outgoing boundary edges,* denoted ∂E_+. The set of interior edges, denoted $\operatorname{Int} E$, is the complement in E of ∂E.

At each interior vertex p, let $\mathcal{N}(p)$ be the set of neighbors of p. As with any directed graph, if qp ends at the vertex where pr begins, we say qp is a *direct predecessor* of pr and pr is a *direct successor* of qp. We assume that a fixed-point free permutation ϕ_p of $\mathcal{N}(p)$ has been specified for each interior vertex p, and that $\phi_p^2 = 1$. Interpret $\phi_p(q)$ as the "preferred direction" at qp, in the sense that a particle moving along an edge qp continues along $\phi_p(q)$ if it does not scatter. Sometimes, q' will denote $\phi_p(q)$ when the base vertex p is clear from the context. We assume that repeated application of ϕ always leads to ∂V eventually, i.e. that any sequence $p, q, \phi_q(p), \phi_{\phi_q(p)}(q), \ldots$ must terminate. Then, the collection of permutations ϕ_p induces a permutation Φ of the boundary vertices by letting $\Phi(r)$ be the unique boundary vertex reached by repeated application of ϕ_{p_i} for appropriate p_i.

Consider particles moving along the edges of G, and let $f(p\,q)$ denote the density of particles travelling along the edge $p\,q$ (in general, $f(p\,q) \neq f(qp)$). From an edge $p\,q$, a particle may continue along $\phi_q(p)$, scatter to a different edge, or be absorbed. We interpret the coefficients $\sigma_a(p\,q) \geq 0$ and $k(rp, p\,q) \geq 0$ as follows: $\sigma_a(p\,q)f(p\,q)$ is the rate of particle absorption at edge $p\,q$, and $k(rp, p\,q)f(rp)$ is the rate at which particles change direction from edge rp to edge $p\,q$ at p. The absorption coefficient σ_a is not defined on incoming boundary edges. We are now ready to make precise the discrete domain on which we will study the transport equation.

Definition 2.1. *A* transport network $(G, \{\phi_p\}, \sigma_a, k)$ *is a graph* G *with a collection of permutations* $\{\phi_p\}$ *satisfying the conditions described above, together with functions* $\sigma_a : E \setminus \partial E_- \to \mathbb{R}_{\geq 0}$ *and* $k : E \times E \to \mathbb{R}_{\geq 0}$ *(extend* $k(a, b)$ *to be 0 if* a *is not a direct predecessor of* b*, and* $k(e, \phi_p(e)) = 0$ *for all* $p \in V, e \in E$*).*

We introduce notation which will help us classify the transport networks under consideration.

Definition 2.2. *For an interior node* p*, let* $\mathcal{N}_{Int}(p) = \mathcal{N}(p) \cap Int V$ *be the set of* interior neighbors *of* p*, and let* $\mathcal{N}_\partial(p) = \mathcal{N}(p) \cap \partial V$ *be the set of* boundary neighbors *of* p*. If* $|\mathcal{N}_\partial(p)|$ *is greater than one (equivalently, if more than one boundary edge is incident to* p*), then* p *will be called a* corner,*and the boundary edges incident to* p *will be called* corner edges.

Because G is connected, a directed path $e_1, \ldots, e_r$ with $e_i \in E$ exists between any two vertices, and in particular between any two boundary vertices. However, in the context of particle transport, a particle can only travel from boundary vertex p to boundary vertex q if $q = \Phi(p)$ or if the particle scatters. Therefore, communication between most boundary nodes requires certain values of k to be nonzero.

Definition 2.3. *For* $p, q \in \partial V$*, we say that a* k-path *exists from* p *to* q *if* $q = \Phi(p)$ *or if a path* $e_1, \ldots, e_r$ *exists between* p *and* q *such that* $k(e_i, e_{i+1}) > 0$ *for all* $1 \leq i < r$*.*

The existence of a k-path from p to q means that some of the particles that enter the network at p will exit the network at q.

Next we derive a discrete analog to the stationary linear transport equation. Assume that particles are travelling throughout the network, and that the system has been allowed to reach equilibrium. In the absence of scattering and absorption, the number of particles travelling along an edge $p\,q$ will equal the number of particles travelling along $q'p = \phi_p(q)p$. Then we have $f(p\,q) = f(q'p)$, or

$$f(q'p) - f(p\,q) = 0.$$

Ignoring the particles scattering to edge $p\,q$ from other edges at p, the difference between the rates of particle transport on $p\,q$ and on $q'p$ equals the rate of absorbed particles:

$$f(q'p) - f(p\,q) = \sigma_a(p\,q)f(p\,q).$$

Finally, if we consider scattering, then the total particle density along $p\,q$ includes particles scattering from other directions,

$$f(q'p) - f(p\,q) + \sum_{\substack{r\in\mathcal{N}(p)\\ r\neq q'}} k(rp, p\,q)f(rp) = \sigma_a(p\,q)f(p\,q),$$

or, rearranging terms,

$$-[f(p\,q) - f(q'p)] - \sigma_a(p\,q)f(p\,q) + \sum_{\substack{r\in\mathcal{N}(p)\\ r\neq q'}} k(rp, p\,q)f(rp) = 0, \quad p\,q \in E\setminus\partial E_-. \tag{3}$$

Equation (3) is the *discrete transport equation* for a transport network (G, σ_a, k). It gives a formula for the particle density $f(p\,q)$ at each interior edge $p\,q$ in terms of the particle densities on the direct predecessors of $p\,q$. At an incoming boundary edge rs, where s is the unique neighbor of r, the value of $f(rs)$ is interpreted as the rate of particles entering the network at r, and it is not governed by (3). We refer to $f(rs)$ as the *incoming flux at* r*,* and the components of the vector f_- of incoming flux values are given as boundary conditions:

$$f(rs) = f_-(rs), \quad rs \in \partial E_-. \tag{4}$$

In this situation, $sr \in \partial E_+$, and $f(sr)$ represents the rate of particles escaping the network at r. $f(sr)$ will be called the *outgoing flux at* r , and the vector of outgoing flux values will be denoted f_+. Note that while f_- is an arbitrary vector given as a boundary condition, f_+ is part of the solution governed by the transport equation (3). Together, (3) and (4) comprise the boundary value problem, or forward problem, for the discrete transport equation.

3 The Forward Problem

Here we solve the forward problem for a general transport network. We assume that σ_a and k are known everywhere in the network, and we wish to find a function $f : E \to \mathbb{R}_{\geq 0}$ that satisfies (3) at each edge in $E \setminus \partial E_-$, and such that $f|_{\partial E_-} = f_-$. Since (3) is linear in f, we will cast the problem as a matrix equation, and we derive some properties of the matrices involved that, while irrelevant to the forward problem, will be useful in studying the inverse problem.

To solve the forward problem, it will be necessary to impose the following condition on σ_a and k:

$$\sum_{\substack{r \in \mathcal{N}(p) \\ r \neq q'}} k(p\,q, qr) \leq \sigma_a(p\,q) \quad \text{for all } p\,q \in E. \tag{5}$$

This condition ensures that, at each edge, the amount of particles scattering to other directions from that edge is no more than the amount absorbed. It is analogous to a condition imposed in [4] and elsewhere to ensure solvability of the continuous transport equation; that condition guarantees subcritical dynamics (i.e., uniformly bounded energy).

Assume that G has n boundary vertices and m (directed) interior edges. In the transport problem represented by (3)-(4), there are $m + n$ unknowns. Let $f_- \in \mathbb{R}^n$ be the vector of inward flux data, and let $f_+ \in \mathbb{R}^n$ be the outward flux values. Finally, let $x \in \mathbb{R}^m$ be the interior particle densities, and let $f = (f_+, x)^T \in \mathbb{R}^{m+n}$. We construct a system of $m+n$ linear equations for the $m + n$ unknown pieces of information. Since (3) holds at each edge $e \in E \setminus \partial E_-$, we associate each edge (excluding incoming boundary) with one of the $m + n$ equations in the obvious way.

For each edge $p\,q$, the value of $f(p\,q)$ depends on the values of f on all of the direct predecessors of $p\,q$. If $p\,q$ is incident to ∂E (in particular, if $f(p\,q)$ is a component of f_+), one or more of the scattering terms in (3) involve f_-, and are therefore known. We move these terms to the right-hand side. If p has degree j, and l of its neighbors are boundary vertices, equation (3) becomes

$$\begin{aligned}(-1 - \sigma_a(p\,q))f(p\,q) + f(q'p) + k(r_1p, p\,q)f(r_1p) + \cdots + k(r_{j-l}p, p\,q)f(r_{j-l}p) \\ = -k(r_{j-l+1}p, p\,q)f(r_{j-l+1}p) - \cdots - k(r_jp, p\,q)f(r_jp).\end{aligned} \tag{6}$$

For an edge $p\,q \in \partial E_+$, if the preferred direction $q' = \phi_p(q) \in \partial V$, the term $f(q'p)$ is a component of f_- and is also sent to the right-hand side. If $p\,q$ is an element of $\operatorname{Int} E$ that is not incident to ∂E, the equation will not involve f_- and all terms will remain on the left-hand side.

We now construct a matrix equation for our vector $f \in \mathbb{R}^{m+n}$ of unknown quantities, which shall be indexed as follows: the first n components of f are comprised of f_+, the vector of outgoing fluxes, and the last m components are the interior particle densities x. Since the right-hand side consists of a linear combination of incoming flux values, our equation will be of the form $Qf = Jf_-$, with block form

$$\begin{pmatrix} Q_1 & Q_2 \\ Q_3 & Q_4 \end{pmatrix} \begin{pmatrix} f_+ \\ x \end{pmatrix} = \begin{pmatrix} J_1 \\ J_2 \end{pmatrix} \left(f_- \right) \tag{7}$$

Here we derive the properties of Q needed to solve the forward problem. Further properties of Q and J are postponed to the end of this section.

Lemma 3.1. ***Properties of*** Q***.*** *Let $f(p\,q)$ be the ith component of f.*
(a) Each diagonal entry $q_{ii} = -1 - \sigma_a(p\,q)$.
(b) The block Q_1 is diagonal.
(c) The block $Q_3 = 0$.
(d) All off-diagonal entries are nonnegative, and in each column, the sum of the off-diagonal entries is at most $1 + \sum_{r \neq q'} k(p\,q, qr)$.

Proof. (a) is clear from the form of equation (6).

Note that for a given $p\,q \in E$, all the terms appearing in (6) depend on values of f on direct predeccessors of $p\,q$. Thus, the value of $f_+(rs)$, the outgoing flux at $s \in \partial V$, does not directly influence any equation in (7 other than the equation for $f_+(rs)$. This proves (b) and (c).

By equation (6), nonzero off-diagonal entries of Q must be either 1 or $k(a_j, b_l)$ for some $a_j, b_l \in E$. If $f(p\,q)$ is a component of x, it appears in $d = \deg(q)$ of the equations in (7): once with a coefficient of 1, $d-2$ times with $k(p\,q, qr_j)$ for some r_j, and once with $(-1-\sigma_a(p\,q))$ (on the diagonal). Thus, in the last m columns of Q, the sum of off-diagonal entries is exactly $1+\sum_{r \neq q'} k(p\,q, qr)$. By (b) and (c), all off-diagonal entries in the first n columns are zero. □

Our goal is to invert Q and solve equation (7) for f. To do this, we must impose condition (5), and then we are able to solve the forward problem.

Theorem 3.2. *Let $(G, \{\phi_p\}, \sigma_a, k)$ be a transport network. If (5) is satisfied for every $p\,q \in E$, the discrete transport equation (3) with boundary data (4) has a unique solution $f(p\,q)$.*

Proof. Equation (7) contains the information from (3) evaluated at each edge. The unknown values of $f(p\,q)$ comprise the vector f. By Lemma 3.1(a) and (d), and condition (5), Q is weakly column diagonally dominant. By Lemma 3.1(b) and (c), the first n columns are strictly diagonally dominant. Hence, Q is invertible. □

Finally, we prove more properties of Q and J which will be useful in analyzing the inverse problem. We recall the definition of a so-called *M-matrix* : E is an M-matrix if all the principle minors of E are positive (see [2]).

Lemma 3.3. *The matrix Q_4^{-1} has nonpositive entries.*

Proof. By Lemma 3.1, $-Q_4$ has positive diagonal entries and nonpositive off-diagonal entries. Also, condition (5) implies the column sums of $-Q_4$ are nonnegative with at least one positive column sum. Together, these conditions are equivalent to the above definition of an M-matrix as well as to the condition that a matrix's inverse be nonnegative (see [2]). We need the second equivalence, so we recall its proof.

Letting $\alpha = \max(\{-Q_4\}_{ii})$ be the maximum diagonal entry of $-Q_4$, we can write $-Q_4 = \alpha I - P$ with P nonnegative. Equivalently, $-Q_4 = \alpha(I - P')$, where P' is nonnegative. By the conditions on the column sums of $-Q_4$, it follows that the column sums of P' are all less than 1, except possibly in column i, where the column sum may equal 1. Hence, $I - P'$ is invertible, and

$$-Q_4^{-1} = \frac{1}{\alpha}(I + P' + (P')^2 + \cdots)$$

Since P' is nonnegative, so is $-Q_4^{-1}$. □

For the next three properties, we must exclude certain cases. For $p \in V$, let $\delta'_{ij} = 1$ if $q_i = \phi_p(q_j)$ and $\delta'_{ij} = 0$ otherwise.

Lemma 3.4. ***Properties of** Q **and** J**.***
(a) If $k(rs, sr) \neq 0$ for each $rs \in \partial E_-$, and if at each corner s with boundary neighbors $r_1, \ldots, r_d \in \partial V$, we have that the $d \times d$ matrix $\{m_{ij}\} = -(\delta'_{ij} + k(r_i s, s r_j))$ is nonsingular, then J_1 is invertible.

(b) If at each corner s with boundary neighbors $r_1, \ldots, r_d \in \partial V$ and interior neigbors $t_1, \ldots, t_h \in Int E$, we have that the $d \times h$ matrix $\{n_{ij}\} = -(\delta'_{ij} + k(r_i s, st_j))$ has full column rank, then J_2 has full column rank.
(c) If at each corner s with boundary neighbors $r_1, \ldots, r_d \in \partial V$ and interior neigbors $t_1, \ldots, t_h \in Int E$, we have that the $h \times d$ matrix $\{o_{ij}\} = \delta'_{ij} + k(r_j s, st_i)$ has full row rank, then Q_2 has full row rank.

Proof. (a) Each row of J_1 corresponds to $f(sr)$, an element of f_+, and to the associated edge sr. If sr is not a corner edge, then the corresponding row of J_1 will be nonzero only on the diagonal, where it takes the value $k(rs, sr)$, which is nonzero by assumption. A corner s with d boundary neighbors will correspond to d rows of J_1, and since the only boundary-to-boundary scattering at s is from one of these d boundary neighbors to another, the $d \times d$ matrix $\{m_{ij}\} = \delta'_{ij} + k(r_i s, sr_j)$ is the submatrix of J_1 consisting of the rows and columns corresponding to the d corner edges incident to s. δ'_{ij} accounts for the coefficient of 1 which appears next to $f(\phi_s(r))$ in the equation associated to $f(sr)$. If all the submatrices $\{m_{ij}\}$ are invertible, then J_1 will be as well, since the sets of rows of J_1 corresponding to each corner and to each other boundary edge are disjoint.

(b) Each column of J_2 multiplies a certain incoming flux value (a certain component of f_-), and thus corresponds to a certain incoming boundary edge. Excluding corners, the sets of direct successors of two distinct incoming boundary edges are disjoint, so the rows where the corresponding columns have nonzero entries are disjoint as well. If the preferred direction from an incoming boundary edge is not an interior edge, that edge is a corner edge (this follows from $\phi_p^2 = 1$). Thus, each column in J_2 not corresponding to a corner edge contains exactly one 1. This means that, regardless of the other entries , those columns of J_2 not corresponding to corner edges are independent. Our condition on corners ensures independence of all the columns of J_2, since the set of direct successors to any corner edge will not intersect the set of direct successors to any non-corner incoming edge or that of any edge incident to a different corner.

(c) The argument here is similar to the proof of (b), except that now we are keeping track of rows of Q_2, which come from equations associated to outgoing flux values, components of f_+. If $sr, tu \in \partial E+$ are not corner edges, the sets of their direct predecessors are disjoint, so the sets of columns in which the corresponding rows have nonzero entries will also be disjoint. As before, every such row has exactly one 1 because the associated edge is not a corner edge, and the rows of Q_2 not corresponding to corner edges are independent. Again, the set of direct predecessors of any corner edge will not intersect the set of direct predecessors of any edge not incident to the same corner, so our condition on $\{o_{ij}\}$ ensures that the rows of Q_2 are independent. □

Remark 1. Note that Lemma 3.4(b) and (c) fail if G contains a corner s with more interior neighbors than boundary neighbors, since the submatrix corresponding to s is the wrong shape to have full column rank. In that case, J_2 and Q_2 are in fact not of full rank. Of all the assumptions in Lemma 3.4, this is the only one which excludes a useful class of networks.

Remark 2. In Lemma 3.4, we give sufficient conditions for Q_2, J_1, and J_2 to have full rank. Later, when we calculate the Jacobian of the map S, these conditions will be imposed to try to maximize the number of linearly independent columns of this Jacobian. We could obtain a more general estimate of the rank of the Jacobian if we derive lower bounds on the ranks of Q_2, J_1, and J_2 rather than insisting they be of full rank, but to simplify the discussion, we will exclude the cases where Lemma 3.4 is not satisfied.

4 The Inverse Problem

Throughout this section, we suppose that our graph G has n boundary edges and m directed interior edges. Letting $M = \sum_{p \in \operatorname{Int} V} \deg(p)(\deg(p) - 1)$, there are $m + n$ values of σ_a and M values of k in the network.

4.1 The Albedo Operator

Here we define the albedo operator, which maps incoming flux to outgoing flux. Contrary to the albedo operator in the continuous transport problem, this albedo operator is linear.

Definition 4.1. *Let $(G, \{\phi_p\}, \sigma_a, k)$ be a transport network satisfying (5). The unique solution to (3) with boundary data f_- (which exists, by Theorem 3.2) determines the outgoing flux f_+. The* albedo operator *is the linear operator that maps f_- to f_+ for this network.*

Since G has n boundary nodes, the albedo operator can be represented by an $n \times n$ matrix $A = \{a_{ij}\}$ (the *albedo matrix*):

$$f_+ = Af_-.$$

If we index the boundary vertices from $1, \ldots, n$, then a_{ij} is equal to the outgoing flux at node i resulting from an incoming flux of 1 at node j and 0 at every other boundary node.

Equation (7) allows us to obtain an explict expression for A. By Lemma 3.1(b) and the invertibility of Q, we have $x = Q_4^{-1} M_2 f_-$, and thus

$$Q_1 f_+ + Q_2 Q_4^{-1} J_2 f_- = J_1 f_-,$$

or

$$f_+ = Q_1^{-1}(J_1 - Q_2 Q_4^{-1} J_2) f_-,$$

from which we obtain

$$A = Q_1^{-1}(J_1 - Q_2 Q_4^{-1} J_2). \tag{8}$$

Next we give conditions for A to be invertible. For this, we need the matrix construction known as the Schur complement. For a matrix V with block structure

$$V = \begin{pmatrix} W & X \\ Y & Z \end{pmatrix}$$

such that Z is invertible, the *Schur complement of Z in V*, denoted V/Z, is given by

$$V/Z = W - XZ^{-1}Y.$$

The rank of a Schur complement can be computed by the Guttman rank additivity formula (see [5]),

$$\operatorname{rank}(V/Z) = \operatorname{rank}(V) - \operatorname{rank}(Z). \tag{9}$$

That formula follows from this factorization of V:

$$V = \begin{pmatrix} I_p & XZ^{-1} \\ 0 & I_q \end{pmatrix} \begin{pmatrix} W - XZ^{-1}Y & 0 \\ 0 & Z \end{pmatrix} \begin{pmatrix} I_p & 0 \\ Z^{-1}Y & I_q \end{pmatrix}$$

The first and third factors are clearly invertible, so the rank of V equals the rank of the central factor, or $\operatorname{rank}(V) = \operatorname{rank}(W - XZ^{-1}Y) + \operatorname{rank}(Z)$, which implies (9).

Formula (9) gives us a lower bound on the rank of A.

Lemma 4.2.

$$\text{rank}(A) = \text{rank}(J_1)$$

In particular, if the conditions of Lemma 3.4(a) are satisfied, so that J_1 is invertible, then A is invertible.

Proof. Introduce the matrix T with the following block structure:

$$T = \begin{pmatrix} Q_1^{-1}J_1 & Q_1^{-1}Q_2 \\ J_2 & Q_4 \end{pmatrix}.$$

Then the Scur complement of Q_4 in T, denoted T/Q_4, is given by

$$\begin{aligned} T/Q_4 &= Q_1^{-1}J_1 - Q_1^{-1}Q_2Q_4^{-1}J_2 \\ &= A. \end{aligned}$$

Since $\text{rank}(Q_4) = m$, we have that $\text{rank}(T) = m + \text{rank}(J_1)$. It follows from (9) that $\text{rank}(A) = \text{rank}(T) - m = \text{rank}(J_1)$. □

The positivity of A provided by the next lemma will also be useful for the inverse problem. The existence of a k-path between any two boundary vertices is a necessary and sufficient condition for this result, but the simpler condition that $k > 0$ everywhere would be sufficient.

Lemma 4.3. *Assume that for any two boundary vertices $p, q \in \partial V$, a k-path exists from p to q. Then all the entries of A are positive.*

Proof. Suppose that an entry $a_{ij} = 0$. This means that, with incoming flux boundary conditions of 1 at boundary node j and 0 elsewhere, the measured outgoing flux at node i is equal to zero. If rs is the outgoing boundary edge incident to node i, the transport equation (3) at rs gives

$$(-1 - \sigma_a(rs))f(rs) = f(s'r) + \sum_{\substack{t \in \mathcal{N}(r) \\ t \neq s'}} k(tr, rs)f(tr). \tag{10}$$

By assumption, there is a k-path from node j to node i. The last edge in this path is rs, and let qr be the next-to-last edge. All the direct predecessors of rs appear in (10), including qr. Since $f(rs) = 0$ and all the terms on the right hand side of (10) are nonnegative, each term is zero. If $q = \phi_r(s)$, then this implies $f(qr) = 0$. Otherwise, it implies $k(qr, rs)f(qr) = 0$. But $k(qr, rs) \neq 0$ by the defining property of a k-path, so $f(qr) = 0$ in this case as well. Proceeding in this way, we find that for each edge e in the k-path from boundary node j to boundary node i, $f(e) = 0$, and that the incoming flux at node j is zero, a contradiction. □

4.2 The Differential of S

In this section we derive sufficient conditions for the map S, which takes σ_a and k to A, to be locally one-to-one. We will show that our conditions imply the differential of S is injective. For notational simplicity, we drop the subscript a in $\sigma_a = \sigma$.

Denote the columns of A by $\mathbf{a_1}, \mathbf{a_2}, \ldots, \mathbf{a_n}$, and concatenate these columns into a vector $\mathbf{a}^* = (\mathbf{a_1}, \mathbf{a_2}, \ldots, \mathbf{a_n}) \in \mathbb{R}^{n^2}$. We will regard σ and k as vectors having $n+m$ and M components respectively. Thus, we are considering

$$\begin{aligned} S : \mathbb{R}^{n+m+M} &\rightarrow \mathbb{R}^{n^2} \\ (\sigma, k) &\mapsto \mathbf{a}^* \end{aligned}$$

Our goal is to find the Jacobian matrix

$$D_{(\sigma,k)}S = \begin{bmatrix} \frac{\partial \mathbf{a_1}}{\partial \sigma_1} & \cdots & \frac{\partial \mathbf{a_1}}{\partial \sigma_m} & \frac{\partial \mathbf{a_1}}{\partial k_1} & \cdots & \frac{\partial \mathbf{a_1}}{\partial k_M} \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{a_n}}{\partial \sigma_1} & \cdots & \frac{\partial \mathbf{a_n}}{\partial \sigma_m} & \frac{\partial \mathbf{a_n}}{\partial k_1} & \cdots & \frac{\partial \mathbf{a_n}}{\partial k_M} \end{bmatrix}$$

of S and compute its rank. As of yet, we are only able to find a lower bound for the rank of $D_{(\sigma,k)}S$ by finding subsets of the columns of $D_{(\sigma,k)}S$ that are provably linearly independent.

Letting ν denote a general component of (σ, k), we have from (8) that

$$\begin{aligned} \partial_\nu A = {} & (\partial_\nu Q_1^{-1})(J_1 - Q_2 Q_4^{-1} J_2) \\ & + \; Q_1^{-1}(\partial_\nu J_1 - (\partial_\nu Q_2) Q_4^{-1} J_2 - Q_2(\partial_\nu Q_4^{-1}) J_2 - Q_2 Q_4^{-1}(\partial_\nu J_2)). \end{aligned} \quad (11)$$

This motivates us to define the following matrices in terms of factors of A: let $B = Q_1^{-1} Q_2 Q_4^{-1}$, and let $C = Q_4^{-1} J_2$. It is clear that when Q_2 and J_2 are of full rank, then B and C will be as well (Lemma 3.4 gives conditions for this to be the case). By Lemma 3.3, Q_4^{-1} is nonpositive. Lemma 3.1(d) implies that Q_2 is nonnegative, and Lemma 3.1(a) and (b) imply that Q_1^{-1} is nonpositive. We therefore have that both B and C are nonnegative.

In the matrix equation (7) for the forward problem, each component of (σ, k) appears exactly once. To see this, consider $k_i = k(p\,q, qr)$ for vertices p, q, and r. k_i will appear once in the equation associated with $f(qr)$, and nowhere else. Similarly, $\sigma_i = \sigma_a(p\,q)$ appears only in the equation associated with $f(p\,q)$. This simplifies the computation of $\partial_\nu A$ for each ν, and we can classify the columns of $D_{(\sigma,k)}S$, into six types based on the submatrix of Q or

J in which the corresponding component of (σ, k) is present. In computing the derivative of Q_4^{-1} with respect to σ and k, we will make use of the following well-known identity, whose proof we recall for completeness.

Lemma 4.4. *For a matrix valued function $M(x)$, if $M(x)$ is invertible at x, then*

$$\frac{d}{dx}M^{-1}(x) = -M^{-1}(x)\left(\frac{d}{dx}M(x)\right)M^{-1}(x). \tag{12}$$

Proof. Observe that

$$M^{-1}(x+h)-M^{-1}(x) = M^{-1}(x+h)M(x)M^{-1}(x)-M^{-1}(x+h)M(x+h)M^{-1}(x).$$

Dividing both sides by h and taking the limit as $h \to 0$, we obtain formula (12). □

Now we characterize the six types of columns of $D_{(\sigma,k)}S$.

Columns of $D_{(\sigma,k)}S$ We find $\partial_\nu A$, where ν is an element of (σ, k). The column vector $\partial_\nu \mathbf{a}^*$ is obtained by concatenating the columns of $\partial_\nu A$. It is clear from the discussion of Section 3 that an element of σ may appear in Q_1 or in Q_4, and that an element of k may appear in Q_2, Q_4, J_1, or J_2, giving us six possible forms for $\partial_\nu A$. We wish to keep track of these forms, and also the number of columns of $D_{(\sigma,k)}S$ that fall into each category. We index σ by i_1 and i_2 and index k by $j_1, \ldots j_4$, and classify the columns of $D_{(\sigma,k)}S$ according to the corresponding components of (σ, k):

1. $\sigma_{i_1} = \sigma(p\,q)$ for $p\,q \in \partial E_+$. In this case, σ_{i_1} appears in Q_1 only, so $\partial_{\sigma_{i_1}} A = (\partial_{\sigma_{i_1}} Q_1^{-1})(J_1 - Q_2 Q_4^{-1} J_2)$. If $p\,q$ is the rth boundary edge in the indexing of ∂E, then $\{Q_1^{-1}\}_{rr} = -\frac{1}{1+\sigma_{i_1}}$, and $\partial_{\sigma_{i_1}} Q_1^{-1}$ has a single nonzero entry: $\frac{1}{(1+\sigma_{i_1})^2}$ in position rr. Thus,

$$\partial_{\sigma_{i_1}} A = c_{i_1}\begin{pmatrix} -0- \\ -a_r- \\ -0- \end{pmatrix},$$

where $c_{i_1} = -\frac{1}{1+\sigma_{i_1}}$ and a_r is the rth row of A. The index i_1 runs from 1 to n.

2. $\sigma_{i_2} = \sigma(p\,q)$ for $p\,q \in \operatorname{Int} E$. Here, σ_{i_2} appears in Q_4 only, so using (12), we have $\partial_{\sigma_{i_2}} A = Q_1^{-1} Q_2 Q_4^{-1} (\partial_{\sigma_{i_2}} Q_4) Q_4^{-1} J_2$, or more compactly, $\partial_{\sigma_{i_2}} A = B(\partial_{\sigma_{i_2}} Q_4) C$. If $p\,q$ is the vth interior edge in the indexing of $\operatorname{Int} E$, then $\{Q_4\}_{vv} = -1 - \sigma_{i_2}$, and $\partial_{\sigma_{i_2}} Q_4$ has a single nonzero entry: -1 in position vv. Thus,

$$\partial_{\sigma_{i_2}} A = B \begin{pmatrix} 0 & & \cdots & & 0 \\ \vdots & \ddots & & & \vdots \\ 0 & \cdots & -1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & & \cdots & & 0 \end{pmatrix} C,$$

or

$$\partial_{\sigma_{i_2}} A = - \begin{pmatrix} | \\ b_v \\ | \end{pmatrix} (-c_v-)\,,$$

where b_r is the vth column of B and c_v is the vth row of C. The index i_2 runs from $n+1$ to $n+m$.

3. $k_{j_1} = k(p\,q, qt)$ for $p\,q, qt \in \operatorname{Int} E$. In this situation, k_{j_1} appears in Q_4 only, so using (12) again, $\partial_{k_{j_1}} A = Q_1^{-1} Q_2 Q_4^{-1} (\partial_{k_{j_1}} Q_4) Q_4^{-1} J_2$, or $\partial_{k_{j_1}} A = B(\partial_{k_{j_1}} Q_4) C$. If $p\,q$ and qt are the vth and uth interior edges in the indexing of $\operatorname{Int} E$ respectively, then $\{Q_4\}_{uv} = k(p\,q, qt) = k_{j_1}$, and $\partial_{k_{j_1}} Q_4$ has a single nonzero entry: 1 in position uv. Thus,

$$\partial_{k_{j_1}} A = B \begin{pmatrix} 0 & & \cdots & & 0 \\ \vdots & \cdots & 1 & \cdots & \vdots \\ 0 & \cdots & \ddots & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & & \cdots & & 0 \end{pmatrix} C,$$

or

$$\partial_{k_{j_1}} A = \begin{pmatrix} | \\ b_u \\ | \end{pmatrix} (-c_v-)\,.$$

Unlike in (2), the column of B and the row of C do not have the same index. The index j_1 runs from 1 to α (we are now indexing the vector

k), where α is the number of components of k present in Q_4. More specific information about α is given by Lemma 4.5 below.

4. $k_{j_2} = k(p\,q, qt)$ for $p\,q \in \operatorname{Int} E, qt \in \partial E_+$. This time, k_{j_2} appears in Q_2 only, so $\partial_{k_{j_2}} A = -Q_1^{-1}(\partial_{k_{j_2}} Q_2) Q_4^{-1} J_2 = -Q_1^{-1}(\partial_{k_{j_2}} Q_2) C$. If $p\,q$ is the vth edge of $\operatorname{Int} E$ and qt is the rth edge in ∂E_+, then $\{Q_2\}_{rv} = k(p\,q, qt) = k_{j_2}$, and the only nonzero entry of $\partial_{k_{j_2}} Q_2$ is a 1 in position rv. Thus,

$$\partial_{k_{j_2}} A = -Q_1^{-1} \begin{pmatrix} 0 & & \cdots & & 0 \\ \vdots & \cdots & 1 & \cdots & \vdots \\ \vdots & & & \ddots & \vdots \\ 0 & & \cdots & & 0 \end{pmatrix} C,$$

or

$$\partial_{k_{j_2}} A = c_{j_2} \begin{pmatrix} -0- \\ -c_v- \\ -0- \end{pmatrix},$$

where the rth row of $\partial_{k_{j_2}} A$ is the vth row of C, and $c_{j_2} = \frac{1}{1+\sigma_r}$. The index j_2 runs from $\alpha+1$ to $\alpha+\beta$, where β is the number of components of k present in Q_2 (see Lemma 4.5 below).

5. $k_{j_3} = k(p\,q, qt)$ for $p\,q \in \partial E_+, qt \in \operatorname{Int} E$. In this case, k_{j_3} appears in J_2 only, so $\partial_{k_{j_3}} A = -Q_1^{-1} Q_2 Q_4^{-1}(\partial_{k_{j_3}} J_2) = -B(\partial_{k_{j_3}} J_2)$. If $p\,q$ is the rth edge of ∂E_+ and qt is the vth interior edge in $\operatorname{Int} E$, then $\{J_2\}_{vr} = -k(p\,q, qt) = -k_{j_3}$, and the only nonzero entry of $\partial_{k_{j_3}} J_2$ is a -1 in position vr. Thus,

$$\partial_{k_{j_3}} A = -B \begin{pmatrix} 0 & & \cdots & 0 \\ \vdots & \ddots & & \vdots \\ 0 & \vdots & \ddots & 0 \\ \vdots & -1 & \cdots & \vdots \\ 0 & & \cdots & 0 \end{pmatrix},$$

or

$$\partial_{k_{j_3}} A = \begin{pmatrix} | & | & | \\ 0 & b_r & 0 \\ | & | & | \end{pmatrix},$$

where the vth column of $\partial_{k_{j_3}} A$ is the rth column of B. The index j_3 runs from $\alpha+\beta+1$ to $\alpha+\beta+\gamma$, where γ is the number of components of k present in the matrix J_2.

6. $k_{j_4} = k(p\,q, qt)$ for $p\,q \in \partial E_-$, $qt \in \partial E_+$. In this case, k_{j_4} appears in J_1 only, so $\partial_{k_{j_4}} A = -Q_1^{-1}(\partial_{k_{j_4}} J_1)$. If $p\,q$ is the rth edge of ∂E_- and qt is the sth edge in ∂E_+, then $\{J_1\}_{sr} = -k(p\,q, qt) = -k_{j_4}$, and the only nonzero entry of $\partial_{k_{j_4}} J_1$ is a -1 in position sr. Note that here, s and r may be equal. We have

$$\partial_{k_{j_4}} A = -Q_1^{-1} \begin{pmatrix} 0 & & \cdots & & 0 \\ \vdots & \ddots & & & \vdots \\ 0 & \cdots & -1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & & \cdots & & 0 \end{pmatrix},$$

or

$$\partial_{k_{j_4}} A = \begin{pmatrix} 0 & & \cdots & & 0 \\ \vdots & \ddots & & & \vdots \\ 0 & \cdots & -\frac{1}{1+\sigma_s} & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & & \cdots & & 0 \end{pmatrix},$$

where entry sr of $\partial_{k_{j_4}} A$ is equal to $-\frac{1}{1+\sigma_s}$. The index j_4 runs from $\alpha+\beta+\gamma+1$ to $\alpha+\beta+\gamma+\delta$, where δ is the number of components of k present in the matrix J_1.

For $h = 1, \ldots, 6$, let L_h denote the collection of columns of $D_{(\sigma,k)}S$ of the form h as numbered in the above list. To completely characterize the columns of the differential of S, it remains to determine how many vectors are in each L_h. We already know that $|L_1| = n$ and $|L_2| = m$, and we have assigned $\alpha = |L_3|$, $\beta = |L_4|$, $\gamma = |L_5|$, and $\delta = |L_6|$. It is clear that since $|L_3| + \cdots + |L_6|$ is equal to the length of the vector k, we must have $\alpha+\beta+\gamma+\delta = M$. The next lemma relates the sizes of $|L_h|$ to the geometry of the graph G. Recall that $\mathcal{N}_\partial(p)$ is the set of boundary neighbors of p and $\mathcal{N}_{\text{Int}}(p)$ is the set of interior neighbors of p.

Lemma 4.5. *Suppose that G has n boundary edges and g corners, and that for l incoming corner edges rs, $\phi_s(r) \in \partial V$ (i.e. the preferred edge at rs is*

in ∂E_+). Then,
(a) $\alpha = |L_3| = \sum_{p\,q \in Int\,E} (|\mathcal{N}_{Int}(q)| - 1) + n - l.$
(b) $\beta = |L_4| = \sum_{sr \in \partial E_+} (|\mathcal{N}_{Int}(p)| - 1) + l$
(c) $\gamma = |L_5| = \beta.$
(d) $\delta = |L_6| = n - l + \sum_{i=1}^{g} |\mathcal{N}_\partial(s_i) - 1|$*, the sum extending over all corners* s_i*.*

Proof. (a) α, the number of occurrences of components of k in the matrix Q_4, is equal to the number of ways a particle can scatter from an interior edge to another interior edge. If $p\,q$ is an interior edge, a particle can continue along $p' = \phi_q(p)$ or it can scatter, so if $p' \notin \partial V$, there are $|\mathcal{N}_{\mathrm{Int}}(q)| - 1$ ways to scatter from pq. If $p' \in \partial V$, then there are $|\mathcal{N}_{\mathrm{Int}}(q)|$ ways to scatter from $p\,q$ to an interior edge. Summing over $\mathrm{Int}\, E$, we observe that $p' \in \partial V$ exactly once for each boundary edge rs unless $\phi_s(r) \in \partial V$, or $n - l$ times.

(b) β is the number of components of k occuring in Q_2, or equivalently, the number of ways for a particle to scatter from an interior edge to an outgoing boundary edge. For an outgoing boundary edge sr, if $\phi_s(r) \in \mathrm{Int}\, V$, then there are $|\mathcal{N}_{\mathrm{Int}}(s)| - 1$ ways to scatter from an interior edge to sr. If $\phi_s(r) \in \partial V$ (this happens a total of l times), there are $|\mathcal{N}_{\mathrm{Int}}(s)|$ ways to scatter from an interior edge to sr.

(c) γ, the number of occurrences of components of k in the matrix J_2, is equal to the number of ways a particle can scatter from an incoming boundary edge to an interior edge. Since $\phi_p^2 = 1$ for all p, if a particle can scatter from $p\,q$ to qr, then it can scatter from rq to qp. This gives a bijection between incoming boundary edge-interior edge pairs $(p\,q, qr)$ where $k(pq, qr) \neq 0$ and interior edge-outgoing boundary edge pairs (rq, qp) where $k(rq, qp) \neq 0$.

(d) δ gives the number of components of k in the matrix J_1, or the number of ways to scatter from outgoing to incoming boundary. If $rs \in \partial E_-$ is not a corner edge, there is only one way to scatter from rs to ∂E_+, namely to sr. If rs is an incoming corner edge, it may scatter to sr or scatter to ∂E_- in $|\mathcal{N}_\partial(s)| - 1$ additional ways, excepting the l cases where $\phi_s(r) \in \partial V$. □

The Jacobian matrix $D_{(\sigma,k)}$ is $n^2 \times (m + n + M)$, so for injectivity, we must have $n^2 - n \geq m + M$. It is not yet known whether this condition is sufficient. If one could show that the all the columns of $D_{(\sigma,k)}S$ found above are independent, then injectivity would follow. At present, we are able to find linearly independent subsets of columns; namely, $|L_1| \cup |L_6|$ forms a basis for a subspace of dimension $n + \delta$, adopting the notation of Lemma 4.5. This result relies upon Lemma 4.3, meaning that we need a k-path to exist from any boundary node to any other boundary node.

Theorem 4.6. *Let g and l be as in Lemma 4.5, and assume that the conditions of Lemma 4.3 hold. If $n - l + \sum_{i=1}^{g} |\mathcal{N}_\partial(s_i) - 1| \geq m + M$, then $D_{(\sigma,k)}S$ is one-to-one.*

Proof. The collection $|L_1|$ is comprised of n matrices of the following form: all rows are zero except the ith row, which is a scaled row of A. Regarding these matrices as vectors in $\mathbb{R}^{n^2}$, they are clearly linearly independent (note that we do not have to use the invertibility of A). The collection $|L_6|$ is comprised of δ matrices, each with one nonzero entry. No two of the matrices in $|L_6|$ have their nonzero entries in the same position, so $|L_6|$ spans a δ-dimensional subspace of $\mathbb{R}^{n^2}$.

Linear combinations of elements of $|L_6|$ can only have nonzero entries in positions where components of k appear in J_1. Since every row of J_1 contains at least one 0 or 1 (i.e. an entry not depending on k), every linear combination of elements of $|L_6|$ has at least one 0 in each row. But by Lemma 4.3, each entry of A is positive, so linear combinations of elements of $|L_1|$ have the property that each row either contains no zeroes or is all zeroes. Therefore, no linear combination of vectors in $|L_6|$ can be equal to a linear combination of vectors in $|L_1|$, and $|L_1| \cup |L_6|$ spans an $n + \delta$-dimensional subspace of $\mathbb{R}^{n^2}$. The Jacobian matrix $D_{(\sigma,k)}S$ has $m + n + M$ columns , so if $n + \delta \geq m + n + M$, i.e. $\delta \geq m + M$, the matrix will have full column rank. By Lemma 4.5, $\delta = n - l + \sum_{i=1}^{g} |\mathcal{N}_\partial(s_i) - 1|$, completing the proof. □

Now the Inverse Function Theorem gives us the desired result.

Corollary 4.7. *If $n - l + \sum_{i=1}^{g} |\mathcal{N}_\partial(s_i) - 1| \geq m + M$ and Lemma 4.3 holds, the map S is locally one-to-one.*

Remark 3. We have not actually used Lemma 3.4 or Lemma 4.2, which gave us conditions for Q_2,J_1, J_2, and A to have full rank. These results may be useful in improving the lower bound on the rank of $D_{(\sigma,k)}S$ (see also the next section on future work).

5 Open Problems

5.1 Weaker Conditions for Local Uniqueness

The linearized problem is to determine whether the matrix $D_{(\sigma,k)}S$ is injective. That matrix has dimensions $n^2 \times (n+m+M)$, so clearly the condition $n^2 - n \geq m + M$ is necessary. We conjecture that it is also sufficient.

Conjecture 5.1. *If $n^2 - n \geq m + M$ and the conditions of Lemmas 3.4, 4.2, and 4.3 hold, then the differential of S is injective.*

If this statement is true, it should be provable through study of the collections $|L_h|$ of columns of $D_{(\sigma,k)}S$. If it is false, determining how many of these columns are indeed independent may give a condition for injectivity of $D_{(\sigma,k)}S$ which is more restrictive than that of Conjecture 5.1 but less restrictive than Theorem 4.6. The nonnegativity of B and C and positivity of A (Lemma 4.3), as well as the fact that A, B, and C are of full rank (Lemma 4.2, Lemma 3.4), may help to show linear independence of some of the columns of $D_{(\sigma,k)}S$.

5.2 Global Uniqueness

We have no results about global injectivity of the map S. The most natural way to prove such a result would be to find an algorithm for reconstructing σ_a and k from A. To this end, we could impose incoming flux of 1 at boundary node p_1 and 0 at every other boundary node and use the transport equation (3) at each edge in $\operatorname{Int} E \cup \partial E_+$ to build a system of $m+n$ equations for $m+m+n+M$ unknowns (m unknown values of f, $m+n$ values of σ_a, and M values of k). Here, we have access to the measured outgoing flux values, which are given by the first column of A. In general, this system will be underdetermined, so we repeat the process n times, placing boundary conditions of 1 at boundary node p_i and 0 elsewhere and measuring the outgoing flux values (column i of the matrix A). Each iteration gives us m additional unknown values of f, while σ_a and k remain the same at each edge, for a total of $n(m+n)$ equations and $nm+m+n+M$ unknowns. Thus, if $n^2 - n \geq m + M$ (the same condition that is necessary for local uniqueness), the system is overdetermined. Even in this case, however, there is no simple way to decide whether this nonlinear system is uniquely solvable.

For a general transport network, there are no obvious shortcuts to this "brute force" method, and even for small graphs, the number of equations involved is daunting. It is straightforward to check that on the b-star for b even (see Figure 1, left), A determines σ_a and k. In such a case, where there are no interior edges, every value of f can be found directly from A, and the recovery process is trivial. But on a 2×2 lattice graph (see Figure 1, right), where $n = 8$, $m = 8$, and $M = 48$, we have 128 equations and 128 unknowns. Analysis of this network, probably with the help of a computer algebra system, may produce an example of a nontrivial graph on which A determines σ_a and k, or, if uniqueness is found to be false, it may provide a counterexample. To this point, neither a counterexample nor a nontrivial example (i.e. a graph with interior edges on which S is one-to-one) has been found.

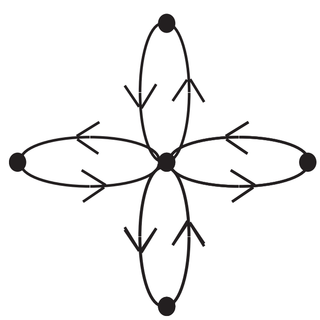

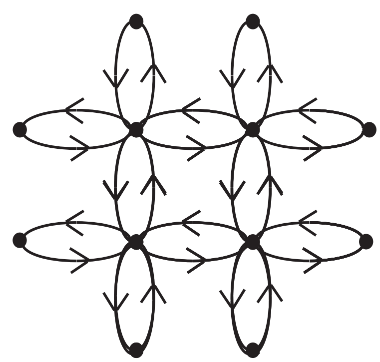

Figure 1: Left: A directed 4-star. Right: A 2×2 directed rectangular lattice.

There are many simplifications that may make the global inverse problem more manageable. As noted above and in [1] and [4], in the continuous inverse problem one must assume that $\sigma_a(x, v) = \sigma_a(x)$ for uniqueness. The analogous assumption for our problem would be to assume σ_a depends only on the vertex where a directed edge begins, i.e. $\sigma_a(p\,q) = \sigma_a(p)$. To see why this condition may be useful to us, consider a $d \times d$ rectangular lattice graph (see Figure 2). Vertices are located at each point in $\{0, d+1\} \times \{0, d+1\}$ except the corners $(0, 0)$, $(0, d+1)$, $(d+1, 0)$, and $(d+1, d+1)$. Vertices of the form $(0, t)$, $(s, d+1)$, $(d+1, t)$, or $(s, 0)$ for $1 \leq s, t \leq d$ are designated boundary vertices. When working with lattice graphs, we usually define ϕ_p in the obvious way, so that for each p, it takes the "left-hand" neighbor of p

to the "right-hand" neighbor of p, and so on.

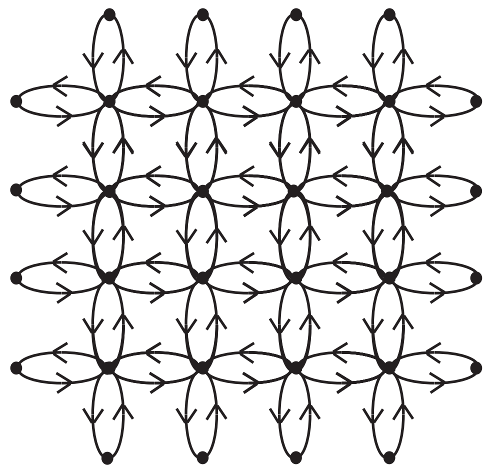

Figure 2: A 4×4 directed rectangular lattice graph.

In this situation, $n = 4d$, $m = 4d^2 - 4d$, and $M = 12d^2$. When $d = 2$ as above, $m = n$ and the system is precisely determined, but for $d \geq 3$, the inequality $n^2 - n \geq m + M$ is not satisfied, so the system is underdetermined. If, on the other hand, we assume $\sigma_a(p\,q) = \sigma_a(p)$, then the number of values of σ_a becomes $m' = d^2$. For the system to be overdetermined, we need $n^2 - n \geq m' + M$, or $16d^2 - 4d \geq d^2 + 12d^2$, which is always satisfied for $d \geq 2$. As far as we know, this still does not imply recoverability of σ_a and k.

Assumptions may also be made on the form of the scattering kernel k. For the continuous transport equation, the inverse problem was solved with $k(x, v', v) = k(x, v' \cdot v)$ (see the references in [4]) before it was solved in the general case. For a precise notion of the "dot product" of two edges, the graph G must have some regular structure; for instance, G could be a rectangular lattice. In that case, we can let $k(p\,q, qr) = k(q, p \cdot r)$ where $p \cdot r$ at q takes values in $\{left, right, back\}$ and is defined as the direction a particle must turn to scatter from $p\,q$ to qr. For a square lattice graph, we then have $M' = 3d^2$ unknown values of k rather than $M = 12d^2$. Other possible simplifications include $k(p\,q, qr) = k(q)$ (isotropic scattering) for any graph and $k(p\,q, qr) = k(\theta(p\,q), \theta(qr))$ for a lattice graph, where $\theta(p\,q)$ takes values in $\{\uparrow, \leftarrow, \downarrow, \rightarrow\}$ and is defined as the direction a particle travels as it moves along edge $p\,q$. If G is a $d \times d$ lattice graph, these assumptions reduce the number of unknown values of k to d^2 and 12 respectively.

Remark 4. One could also study the linearized problem under simplifying assumptions on σ_a and k, but if Conjecture 5.1 is true, such study would be unneccessary, as $D_{(\sigma,k)}S$ would remain injective under any assumption that decreases the number of unknowns.

5.3 Asymmetric Directed Networks

Finally, we pose the inverse problem for simple directed graphs, that is, for graphs where any two vertices have at most one directed edge joining them. It may also be worthwhile to look at graphs where vertex pairs may have zero, one, or two (one in each direction) directed edges joining them, but here we look at the simpler case.

Let $G = (V, \partial V, E)$ be a simple directed graph with boundary. Assume G is finite and connected, and that each boundary vertex has degree one. Define the incoming and outgoing boundary ∂E_- and ∂E_+ exactly as in Section 2; the difference is that here, a given boundary vertex will be adjacent to ∂E_- or ∂E_+, not to both. Also assume that at each interior vertex p, we have a bijection ϕ_p from the set of edges ending at p ($in(p)$) to the set edges beginning at p ($out(p)$). In particular, this implies $|in(p)| = |out(p)|$ for all $p \in \mathrm{Int}\,(p)$. Interpreting ϕ_p as in Section 2, we can consider particles travelling along the directed edges of G. Letting $q' = \phi_p^{-1}(q)$, the equation governing this transport is

$$-[f(p\,q) - f(q'p)] - \sigma_a(p\,q)f(p\,q) + \sum_{\substack{r \in in(p) \\ r \neq q'}} k(rp, p\,q)f(rp) = 0, \quad p\,q \in E \setminus \partial E_-. \tag{13}$$

As before, we prescribe incoming flux,

$$f(rs) = f_-(rs), \quad rs \in \partial E_-, \tag{14}$$

and measure outgoing flux f_+ on ∂E_+.

The forward problem (13)-(14) is solved exactly as in the symmetric case (Theorem 3.2). The condition for solvability is now

$$\sum_{\substack{r \in out(p) \\ r \neq \phi_p(q)}} k(p\,q, qr) \leq \sigma_a(p\,q) \quad \text{for all } p\,q \in E.$$

This condition ensures the matrix Q in the equation $Qf = Jf_-$ is diagonally dominant.

Since the forward problem is solvable, we define $A : f_- \mapsto f_+$ exactly as in Section 4. Here, if G has n boundary vertices, A is an $(n/2) \times (n/2)$ matrix. The inverse problem is once again to recover σ_a and k from A. It is hoped that this problem will be simpler than the case where G is a symmetric directed graph because the number of unknown values of σ_a and k is lower.

For an edge $pq \in E$, define $in(pq) = in(p)$ and $out(pq) = out(q)$. One interesting feature of this problem is that if the directed edges of G do not form any loops, then the interior edges $e_1, \ldots, e_m$ can be indexed so that $e_j \notin in(e_i)$ for any $i < j$. That this indexing is possible follows from the fact that any partial order can be extended to a linear order. Since (13) for an edge e involves only edges in $in(e)$, this means that with the appropriate indexing of $\operatorname{Int} E$, the matrix Q_4 is upper triangular. This simplifies the calculation of Q_4^{-1}, which may lead to explicit formulas for the entries of A in terms of σ_a and k.

References

[1] M. Choulli and P. Stefanov, *An inverse boundary value problem for the stationary transport equation*, Osaka J. Math. 36:1 (1999), 87-104.

[2] A. Berman and R. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Academic Press, New York, 1979.

[3] J. Giansiracusa, *The Discrete Transport Equation*, University of Washington REU 2003 Student Report.

[4] P. Stefanov and G. Uhlmann, *Optical Tomography in two dimensions*, Methods Appl. Anal., 10:1 (2003), 1-10.

[5] F. Zhang, Ed. The Schur Complement and its Applications, Springer, New York, 2005.