

Multi-commodity allocation for dynamic demands using PageRank vectors

Fan Chung¹, Paul Horn², and Jacob Hughes¹

¹ University of California, San Diego

² Harvard University

Abstract. We consider a variant of the contact process concerning multi-commodity allocation on networks. In this process, the demands for several types of commodities are initially given at some specified vertices and then the demands spread interactively on a contact graph. To allocate supplies in such a dynamic setting, we use a modified version of PageRank vectors, called Kronecker PageRank, to identify vertices for shipping supplies. We analyze both the situation that the demand distribution evolves mostly in clusters around the initial vertices and the case that the demands spread to the whole network. We establish sharp upper bounds for the probability that the demands are satisfied as a function of PageRank vectors.

1 Introduction

Efficient allocation of resources to meet changing demands is a task arising in numerous applications. For example, institutions such as governments or corporations respond to the needs of a populace, and wish to meet the demands within allowed expenditure of resources. In some cases where demand spreads, one has to be able to act before demand becomes unmanageable. In the case of an epidemic, for instance, one desires to find a way to distribute medicine so that the disease will be contained. Such problems have been studied in several contexts using the contact process model [8], [6], [1], [10], [4]. In [4], it was demonstrated how to use PageRank vectors to both restrict the number of nodes inoculated and to provide certain containment guarantees.

In this paper, we study a variant of the classical contact process, a continuous time Markov process on a contact graph. In our scenario, vertices in the graph each have varying levels of demand for multiple commodities. Demand at a vertex propagates to its neighbors at a rate depending on the current demand. Our model allows for interactions between different commodities; demand for one commodity may influence demand for another. This fits many scenarios that arise, for instance demand for iPhones may accelerate the demand for iPads. As another example, demand at a node can be viewed as a measure of discontent with the current supply of a resource. It is natural for an unhappy node to create unrest in its neighbors. As the modified contact process continues, demands at a vertex are increased at a rate based on the demands at neighboring vertices. There are also decreased at a satisfaction rate, which can be thought of as a frequency of shipments. Demand spreads at rates which are a linear combination of demands from neighboring vertices. These rates are encapsulated in a spread matrix, B , roughly analogous to the infectivity parameter in the classical contact process. The goal of this paper is to find satisfaction rates, dependent on the spread matrix B and the geometry of the contact graph which ensure that eventually all vertices have no demand and the process dies out. Our process will be defined, in detail, in Section 2.

To satisfy the demands which evolve according to our model as defined in Section 2, the goal is to ship commodities and supply vertices with unsatisfied demands in an efficient way. The model here differs somewhat from typical resource allocation problems in the sense that we do not specify the location of the “warehouses” for the supply. We will not be concerned with either the sources of the supply or the detailed incremental costs of shipping supply. Instead, our goal is to identify *how often* to ship each commodity to a particular vertex, in order to contain and satisfy demands, given an initial seed set. The reader is referred to [3] for the usual resource allocation problem.

Contact graphs of interest take many forms: Cities and countries exert trade pressure on neighboring cities and countries. Communication on the internet can also spread demand for products, or discontent leading to a revolution. Instead of studying this problem on particular models for these contact graphs, we study the problem on arbitrary finite graphs. Two schemes of making shipments are considered. First is a global solution which involves “scheduling shipments” to all vertices in the graph, and ensures that all demand is satisfied in $O(\log n)$ time, with high probability and regardless of the initial demand. This is made precise in Theorem 1 once the model is formally defined. The next scheme is a local solution, in the sense that shipments are scheduled to only a subset of vertices which contain the initial demand. In particular, when the contact graph has some clustering structure we are interested in subsets so that the demand within the subset is satisfied quickly (in $O(\log n)$ time) and demands reach a vertex not receiving shipments with low probability. Precise results to this end are given in Section 5.

This latter scheme relies on understanding the geometry of the particular contact graph being studied. Our scheme uses PageRank to identify important vertices and to bound the probability that demand in our process leaves a set. We also introduce a variant of PageRank, which we call Kronecker PageRank and is introduced in Section 3, which provides sharper bounds by better utilizing the structure of the spread matrix B as well as the geometry of the graph in its estimates. Our analysis provides a tradeoff in the following sense: we may use PageRank estimates to *identify* a set of vertices containing the initial demand which are important to ship to or we may use our PageRank estimates to give a *guarantee* on the escape probability of leaving a particular set of our choosing. Precise results to this end are given in Theorems 3 and 4, using standard PageRank and Kronecker PageRank respectively.

2 Preliminaries and The Demand Model

We model demand spreading within an undirected simple graph, $G = (V, E)$. We write this $v \sim w$ when v and w are adjacent. For each vertex $v \in V$, let d_v be the degree of v , which is the number of neighbors of v . While not strictly necessary, we assume that there is a self loop at each vertex. In this case, d_v includes v in the count of neighbors and hence the loop counts as 1 towards the degree. We let $n = |V|$, the number of nodes of G . An exponential random variable with parameter λ has probability density function given by $f(x) = \lambda e^{-\lambda x}$ for $x \geq 0$, and 0 for $x < 0$. This distribution will be denoted $\text{Exp}(\lambda)$. One important property of exponential random variables is the memoryless property: if X is an exponential random variable then for any constants $a, b > 0$,

$$\mathbb{P}(X > a + b | X > a) = \mathbb{P}(X > b).$$

If X and Y are independent and $X \sim \text{Exp}(\lambda_1)$, $Y \sim \text{Exp}(\lambda_2)$ then $\min\{X, Y\} \sim \text{Exp}(\lambda_1 + \lambda_2)$. A Poisson point process at rate λ is a sequence of random variables $\{X_i\}_{i=1}^\infty$ so that X_1 and $X_i - X_{i-1}$, for $i \geq 2$, has distribution $\text{Exp}(\lambda)$.

Before we describe our model, let us briefly recall the contact process on a graph G , which we denote $CP(T, \beta, \sigma, G)$. In the contact process (see for example [1] or [4]), a disease initially infects a set $T \subseteq V(G)$. The disease has an infectivity parameter, β , and each vertex has a certain amount of “medicine” σ_v . An infected vertex v infects its neighbor u at times given by a Poisson point process $\{X^{uv}\}$ at rate β , and each infected vertex is cured at times given by a Poisson point process at rate σ_v . In the most frequently studied case, σ is constant and the host graph is an infinite graph. The process ends when all vertices are cured, and the basic problem is to determine under which conditions on σ , and β the process ends almost surely. In the case of finite graphs, if $\sigma_v > 0$ for every vertex, it is easy to observe that the process ends a.s., so the problem becomes determining *how fast* the process ends.

The k -commodity dynamic demand model on a graph G is a variant of the contact process, $DD(\tau(0), B, \sigma, G, N)$. In this situation, B is a real valued $k \times k$ -matrix (not assumed to be symmetric, or even non-negative), which we call the spread matrix. The supply function is $\sigma : V \rightarrow \mathbb{R}^k$, and $\tau(0) : V \rightarrow \mathbb{N}^k$ is the initial demand. The state of the process at time t is given by $\tau(t) : V \rightarrow \mathbb{R}^k$, which gives the demands for each vertex at time t . We use $\tau_v(t) \in \mathbb{N}^k$ to denote the demand at vertex v at time t , and $\tau_v^j(t) \in \mathbb{N}$ to denote the demand for commodity j at time t . A node v is said to be *satisfied* at time t if $\tau_v(t) = \mathbf{0}$, and *unsatisfied* otherwise. $N \in \mathbb{N}$ serves as a uniform bound for the maximum demand for any resource at any point (for instance, it could be the population of a city, if a vertex is a city.) The existence of N is simply to ensure integrability of some random variables. In the case of the contact process, $N = 1$.

The spread matrix $B = [\beta_{ij}]$ describes how the demand for one commodity influences demands for other commodities. The i, j entry of B , β_{ij} , determines the spread rate of the demand for commodity j that is caused by demand for commodity i . In particular, we can describe the rate of spread events as follows. If v is a node that is unsatisfied at time t , and w an adjacent vertex, then there are spread events from v to w with rates $\max\{\tau_v(t)B, 0\}$. That is the rate at which τ_w^j increases due to the demand at v is given by $\max\{\sum_i \tau_v^i(t)\beta_{ij}, 0\}$. Here, when we say an event occurs with rate λ , we mean that the elapsed time until that event takes place is distributed as $\text{Exp}(\lambda)$. Because the minimum of exponential random variables is itself an exponential random variable, we can capture the total spreading rates in a condensed form. We define the rate function at time t , $\rho(t) : V \rightarrow \mathbb{R}^k$, by

$$\rho_v = \sum_{w \sim v} \tau_w(t)B = (\tau(t)(A \otimes B))_v,$$

where $\tau(t)$ is viewed as a vector with indices indexed by $V \times k$. $\rho_v^i(t)$ is the rate at which τ_v^i is increasing at time t . Any spread events that would raise τ_v^i above N are ignored.

Supply events occur with rates given by $\tau(t)\text{Diag}(\sigma)$, independently of any neighboring supply events. That is, the time until τ_v^i is decreased by 1 is distributed as $\text{Exp}(\sigma_v^i \tau_v^i)$.

We briefly give a construction of the process, to show it is well-defined. Let \vec{E} denote the set of ordered edges; that is, ordered pairs that are edges in the graph, so that uv and vu are

distinct. We run independent Poisson point processes $\{X_e^{j,\rho}\}_{e \in \vec{E}(G), j \in [k], \rho \in [N]^k}$ so that $X_e^{j,\rho}$ is at rate $\max\{0, [\tau_v B]_j\}$ and independent Poisson point processes $\{X_v^{i,n}\}_{v \in V(G), i \in [k], n \in [N]}$ so that $X_v^{i,n}$ is at rate $n\sigma_v^i$. Then these finitely many point processes can easily be seen to define the entire process; a spread event of type j from a vertex v to a vertex u which is currently in state ρ is controlled by the point process $X_{vu}^{j,\rho}$ with satisfaction events handled similarly.

Such a formulation is that it gives an easy coupling between processes that shows that if $B' \leq B$ pointwise, the stochastic process $DD(\tau(0), B, \sigma, G, N)$ stochastically dominates $DD(\tau(0), B', \sigma, G, N)$. That is, in the coupling the demands in the B process are always at least those in the B' process. This is accomplished by noting that the rates $\rho B \geq \rho B'$ pointwise for all $\rho \in \mathbb{N}^k$. We thus take point processes $Y_e^{j,\rho}$ at rate $[\rho B - \rho B']_i$. If the point processes $\{X_e^{j,\rho}\}$ and $\{X_v^{i,n}\}$ are used to determine $DD(\tau(0), B, \sigma, G, N)$, then the point processes $\{X_e^{j,\rho} \cup Y_e^{j,\rho}\}$ and $\{X_v^{i,n}\}$ are used to determine $DD(\tau(0), B', \sigma, G, N)$.

In particular, this allows us to replace B with B' , where $B'_{ij} = \max\{B_{ij}, 0\}$, and conclusions about the extinction of the B' process still hold for B . Furthermore, this turns out not to be entirely unreasonable. One hopes that the negative entries in B would afford better bounds on the extinction time, but in many cases with negative entries in B extinctions of some demand types mean that the process is eventually run in a non-negative case. In light of this, we will assume for the rest of this paper that B is *non-negative* for convenience.

Given an initial demand $\tau(0)$ and spread matrix B , our goal is to find a supply function σ such that demand is satisfied. Ideally we would like to do this with small supply rates. Furthermore, the supply rates should only depend on the contact graph G , the spread matrix B , and the initial demand $\tau(0)$, but not on t or $\tau(t)$.

Because we take B to be an arbitrary positive $k \times k$ matrix, we will at times need to use various matrix norms in order to understand the process. For a square matrix B , there are many different matrix norms that can be used (see [7]). We will use the following notation for the following norms:

1. $\|B\|_1 = \sum_{i,j} |a_{ij}|$ is the ℓ_1 norm.
2. $\|B\|_1 = \max_j \sum_i |a_{ij}|$ is the *maximum column sum norm*.
3. $\|B\|_\infty = \max_i \sum_j |a_{ij}|$ is the *maximum row sum norm*.
4. $\|B\|_2 = \max\{\sqrt{\lambda} \mid \lambda \text{ is an eigenvalue of } A^*A\}$ is the *spectral norm*.

3 PageRank and Kronecker PageRank

The notion of PageRank was first introduced by Brin and Page [2] in 1998 for Google's search algorithms. Although PageRank was originally used for the Web graph, it is well defined on any finite graph G . The basis of PageRank is random walks on graphs. A walk is a sequence of vertices (v_0, v_1, \dots, v_k) where $v_i \sim v_{i+1}$. A simple random walk of length k is a sequence of random variables (x_0, \dots, x_k) where the starting vertex x_0 is chosen according to some distribution, and

$$\mathbb{P}(x_{i+1} = v \mid x_i) = \begin{cases} \frac{1}{d_{x_i}} & \text{if } x_i \sim v \\ 0 & \text{if } x_i \not\sim v \end{cases}.$$

Let D be the diagonal degree matrix with entries $D_{vv} = d_v$, and let A be the adjacency matrix with entries

$$A_{vw} = \begin{cases} 1 & \text{if } v \sim w \\ 0 & \text{if otherwise.} \end{cases}$$

Then the transition probability matrix for a random walk on G is given by $W = D^{-1}A$.

We will mainly use a modified version of the PageRank, called personalized PageRank. Personalized PageRank has two parameters, a jumping constant $\alpha \in [0, 1]$ and a seed \mathbf{s} which is some probability distribution on the vertex set V of G .

The personalized PageRank vector $\mathbf{pr}(\alpha, \mathbf{s})$ for jumping constant α and the seed distribution \mathbf{s} on V is given by

$$\mathbf{pr}(\alpha, \mathbf{s}) = \alpha \sum_{\ell=0}^{\infty} (1 - \alpha)^\ell \mathbf{s} W^\ell.$$

Note that here we view \mathbf{s} as a row vector, which is our convention for all vectors throughout this paper. We note that the PageRank vector is also the solution to the recurrence relation

$$\mathbf{pr}(\alpha, \mathbf{s}) = \alpha \mathbf{s} + (1 - \alpha) \mathbf{pr}(\alpha, \mathbf{s}) W.$$

The original definition of PageRank [2] is the special case where \mathbf{s} is the uniform distribution over all the vertices.

For a subset of vertices $H \subset V$, the volume of H is the sum of degrees of the vertices of H . The Cheeger Ratio of H , $h(H)$, measures the cut between H and \bar{H} via the relationship

$$h(H) = \frac{e(H, \bar{H})}{\min\{\text{vol}(H), \text{vol}(\bar{H})\}}.$$

The α -core of a subset H is the set of vertices

$$C_\alpha = \left\{ v \in H \mid \mathbf{pr}(\alpha, 1_v) 1_H \geq 1 - \frac{h}{\alpha} \right\}.$$

Personalized PageRank naturally gives bounds on the probability that demands leave a given set in the k -commodity dynamic demand model. These bounds lose some of the structure of the process given by B , however. In order to better understand the process and get tighter bounds, we will use generalization of the personalized PageRank vector, the Kronecker PageRank vector.

If B is an $k \times k$ matrix, and A an $n \times n$ matrix, then the Kronecker product $A \otimes B$ is the $nk \times nk$ block matrix

$$A \otimes B = \begin{pmatrix} a_{11}B & \cdots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{n1}B & \cdots & a_{nn}B \end{pmatrix}$$

With this, we can define Kronecker PageRank.

Definition 1 (Kronecker PageRank) Let B be a square $k \times k$ matrix with spectral radius strictly less than 1, and W be the transition matrix for a random walk on a graph G . Let \mathbf{s} be a non-negative vector in $\mathbb{R}^{k \times |V|}$. The Kronecker PageRank vector with parameters B and \mathbf{s} is defined as

$$\mathbf{Kpr}(B, \mathbf{s}) = \sum_{\ell=0}^{\infty} \mathbf{s}(W \otimes B)^\ell = \sum_{\ell=0}^{\infty} \mathbf{s}(W^\ell \otimes B^\ell)$$

Requiring the spectral radius of B less than 1 is necessary to ensure convergence of the infinite sum, as the spectrum of $W \otimes B$ is the product of the spectra of W and B . Since the eigenvalues of W have absolute value at most 1, the sum will converge.

We note that in the case where B is a 1×1 matrix $B = \beta < 1$ and \mathbf{s} is a probability distribution, then we have the relationship

$$\mathbf{Kpr}(B, \mathbf{s}) = \sum_{l=0}^{\infty} \mathbf{s}(W \otimes \beta)^l = \sum_{l=0}^{\infty} \mathbf{s}\beta^l W^l = \frac{1}{1-\beta} \mathbf{pr}(1-\beta, \mathbf{s}),$$

so the Kronecker PageRank is a natural extension of personalized PageRank. We will see in Theorem 4 that the Kronecker PageRank will arise naturally in our analysis in Section 3, and give better bounds than those that will be afforded by standard PageRank by incorporating the spread matrix.

4 Global analysis: Supplying every vertex

Here we show that if supply rates are above a certain threshold, then with probability approaching 1 demands will be satisfied. Recall that B is a non-negative real valued $k \times k$ matrix, and $\boldsymbol{\sigma}$ is the vector of supply rates for the process $DD(\boldsymbol{\tau}(0), B, \boldsymbol{\sigma}, G, N)$.

Theorem 1 Consider the k -commodity demand model on a graph G with n vertices parameterized by spread matrix $B = [\beta_{ij}]$. Let $X(t) = \|\boldsymbol{\tau}(t)\|_1$, the total amount of demand at time t . If the supply rates to each vertex v are $\sigma_v^i > d_v \left(\sum_j \frac{\beta_{ij} + \beta_{ji}}{2} \right) + \delta$ for $\delta > 0$ then with probability $1 - \epsilon$ all vertices are satisfied at time t for all

$$t > \frac{1}{\delta} \left(\frac{1}{2} \log(nk) + \log(X(0)) + \log\left(\frac{1}{\epsilon}\right) \right).$$

Proof.

We begin by considering the quantity $\frac{\partial}{\partial t} \mathbb{E}[\boldsymbol{\tau}(t)]$. From the discussion in Section 2, we know that demand is increasing with rates given by $\boldsymbol{\rho}(t) = \boldsymbol{\tau}(t)(A \otimes B)$, but also demand decreases proportionally to the supply rates and current demand. Indeed, let $S = \text{diag}(\boldsymbol{\sigma})$, the diagonal $nk \times nk$ matrix with entries given by the supply vector. Then demand decreases at each vertex according to rates given by the supply rate vector $\boldsymbol{\tau}(t)S$.

It is not difficult to encapsulate all of the above information in the simple expression

$$\frac{\partial}{\partial t} \mathbb{E}[\boldsymbol{\tau}(t)] = \mathbb{E}[\boldsymbol{\rho}(t) - \boldsymbol{\tau}(t)S] = \mathbb{E}[\boldsymbol{\tau}(t)](A \otimes B - S). \quad (1)$$

A detailed proof of (1) is left to the appendix for space reasons.

Solving the matrix differential equation with initial condition $\mathbb{E}[\boldsymbol{\tau}(0)] = \boldsymbol{\tau}(0)$ yields

$$\mathbb{E}[\boldsymbol{\tau}(t)] = \boldsymbol{\tau}(0)e^{t(A \otimes B - S)}. \quad (2)$$

Let $Q = A \otimes B - S$. Then by [5], $\|e^{tQ}\|_2 \leq e^{t\nu}$, where ν is the largest eigenvalue of $\frac{Q+Q^*}{2}$. We note that $\frac{Q+Q^*}{2} = A \otimes (\frac{B+B^*}{2}) - S$, which has diagonal terms $\beta_{ii} - \sigma_v^i$, ranging over all values of v and i . By the Gershgorin Circle Theorem, the eigenvalues of $\frac{Q+Q^*}{2}$ are contained in the intervals

$$\left[-(d_v - 2)\beta_{ii} - d_v \left(\sum_{j \neq i} \frac{\beta_{ij} + \beta_{ji}}{2} \right) - \sigma_v^i, \quad d_v \left(\sum_j \frac{\beta_{ij} + \beta_{ji}}{2} \right) - \sigma_v^i \right].$$

Since $\sigma_v^i > d_v \left(\sum_j \frac{\beta_{ij} + \beta_{ji}}{2} \right) + \delta$ all the eigenvalues of $\frac{Q+Q^*}{2}$ are less than $-\delta$. Therefore

$$\begin{aligned} \mathbb{E}[X(t)] &= \|\tau(0)e^{t(A \otimes B - S)}\|_1 \leq \sqrt{nk} \|\tau(0)e^{t(A \otimes B - S)}\|_2 \\ &\leq \sqrt{nk} \|\tau(0)\|_2 \|e^{t(A \otimes B - S)}\|_2 \leq \sqrt{nk} \|\tau(0)\|_1 e^{t\nu} \\ &\leq \sqrt{nk} X(0) e^{-t\delta} \end{aligned}$$

Thus Markov's inequality gives that $\mathbb{P}(X(t) > 0) < \epsilon$ if $t > \frac{1}{\delta} \left(\frac{1}{2} \log(nk) + \log(X(0)) + \log\left(\frac{1}{\epsilon}\right) \right)$.

□

We note that this approach works for all initial distributions $\boldsymbol{\tau}(0)$, and on any graph G . In particular, it is agnostic to the shape of the graph. This indicates that in many situations this approach may be overkill and that we could have used smaller supply rates. In the next section, we analyze the process more carefully and give conditions that depend on the initial distribution of demand and the geometry of the underlying contact graph.

5 Local analysis: Supplying a small subset

For the remainder of the discussion, it is convenient to introduce reformulation of the model that takes advantage of the fact that demands take on integer values. Rather than view demands as a function $\tau : V \rightarrow \mathbb{N}^k$, we view demands as discrete objects sitting on each node. Borrowing language from chip-firing games on graphs (see, for example, [9]) we view units of the demand as chips located on vertices of the graph. For example, if $k = 7$ then for a vertex v with $\boldsymbol{\tau}_v(t) = (0, 1, 2, 0, 2, 0, 3)$ then we would say that at time t there was 1 2-chip, 2 3-chips, 2 5-chips, and 3 7-chips at vertex v , corresponding to 1 “unit of demand” for commodity 2, etc. Unlike in classical chip-firing games, the number of chips is not static, and the process is continuous time. We restate the possible transitions in terms of demand chips. For an i -chip at vertex v , there are two types of transition events:

- For each vertex $w \sim v$ and each $j = 1, \dots, \ell$, a j -chip is added at w at rate β_{ij} . When this occurs we say that the new j -chip is created by the i -chip.
- The i -chip itself is removed with rate σ_v^i .

Due to the properties of exponential random variables, the rates add linearly, and the model is equivalent to the original description discussed in Section 2. The main advantage of this reformulation is the ability to trace back the history of a chip. If there is a chip c at vertex v at time t , then either c existed at time $t = 0$, or there is a sequence of ℓ chips $(c_0, \dots, c_\ell = c)$ located at vertices along a walk $\pi = (v_0, v_1, \dots, v_\ell = v)$ with the following properties:

1. c_0 existed at $t = 0$
2. c_r is created by c_{r-1} for $r = 1, \dots, \ell$.

We allow π to have repeated vertices to allow for the case where demand created more demand at the same vertex. If a chip c exists at time 0, we refer to it as an *initial chip*.

For a path $\pi = (v_0, v_1, \dots, v_\ell)$ and a chip c_0 located at v_0 , we define the event \mathcal{S}_{π, c_0} to be the event there is a sequence of m chips (c_0, \dots, c_ℓ) located respectively at $(v_0, v_1, \dots, v_\ell)$ and c_r is created by c_{r-1} for $r = 1, \dots, \ell$.

It is important to note that \mathcal{S}_{π, c_0} occurring does not imply that there is any demand at v_ℓ at time t because it could be satisfied sometime before t . However, if there is a demand at v_ℓ at time t , then $\mathcal{S}_{\pi, c}$ must have occurred for some initial chip c at vertex v_0 and some walk π from v_0 to v_ℓ .

We begin by relating $\mathbb{P}(\mathcal{S}_{\pi, c_0})$ to the length of the walk π . Inspired by Theorem 1 we make the assumption that supply rates are proportional to the degree of the vertices. That is, we assume that $\sigma_v^i > \mu_i(d_v)$ for all v for constants $\mu_i > 0$.

Lemma 2 *Let $M = \text{diag}(\mu_1, \dots, \mu_k)$, $\hat{B} = M^{-1}B$ and $\zeta = \min\{\|\hat{B}\|_1, \|\hat{B}\|_\infty\}$*

Then for any chip c_0 located at v_0 and any walk $\pi = (v_0, \dots, v_\ell)$ of length ℓ ,

$$\mathbb{P}(\mathcal{S}_{\pi, c_0}) \leq k\zeta^\ell \prod_{j=0}^{\ell} \frac{1}{d_{v_j}}$$

Proof.

Let S_r denote the event that a chip c_r at v_r creates a chip at v_{r+1} . If c_r is an i -chip, then for it to create any chip at v_{r+1} a spread event must occur before c_r is removed. The time until c_r creates a j -chip at v_{r+1} is an exponential random variable with rate β_{ij} . Since the time until c_r is removed is given by $\text{Exp}(\sigma_v^i)$, the probability of c_r creating a j -chip is $\frac{\beta_{ij}}{\beta_{ij} + \sigma_v^i} \leq \frac{\beta_{ij}}{\sigma_v^i} < \frac{\beta_{ij}}{\mu_i d_{v_r}}$. Thus $\mathbb{P}(S_r) < \sum_{i,j} \frac{\beta_{ij}}{\mu_i d_{v_r}} = \frac{1}{d_{v_r}} \mathbf{1}\hat{B}\mathbf{1}^*$.

For a walk π of length ℓ , we want to consider the intermediate steps more carefully. Since there are ℓ transitions that occur, we can use the same reasoning as above to obtain the bound

$$\mathbb{P}(\mathcal{S}_{\pi, c}) < \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \mathbf{1}\hat{B}^\ell \mathbf{1}^* = \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \|\hat{B}^\ell\|_1 \leq k \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \|\hat{B}^\ell\|_1 \leq k \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \|\hat{B}\|_1^\ell.$$

The factor of k that appears in the final lines above is just a consequence of switching from the vector 1-norm $\|\hat{B}^\ell\|_1$ to maximum column sum norm $\|\hat{B}^\ell\|_1$ (see [7]).

We could have just as easily switched to the maximum row sum norm and obtained the term $k\|\hat{B}\|_\infty^\ell$, and so it follows that

$$\mathbb{P}(S_{\pi,c}) < \min\left\{k \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \|\hat{B}\|_1^\ell, k \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \|\hat{B}\|_\infty^\ell\right\} = k\zeta^\ell \prod_{j=0}^{\ell} \frac{1}{d_{v_j}}.$$

□

We note that the use of $\zeta = \min\{\|\hat{B}\|_1, \|\hat{B}\|_\infty\}$ in Lemma 2 reflects the difficulty in working with arbitrary spread matrices B . For certain classes of spread matrices (e.g. if B is symmetric or diagonalizable) it is possible to obtain tighter bounds. Lemma 2 will allow us to obtain a bound using PageRank, but note that our use of matrix norms ignores some of the structure of the spread matrix B . Following the proof of Theorem 3 below, a more careful analysis fully using the structure of the spread matrix B will lead naturally to use of Kronecker PageRank, which we explore in Theorem 4.

Theorem 3 *Suppose that initial demand is contained in $S \subset H \subset V$ with and each vertex $v \in H$ has supply rates $\sigma_v^i > \mu_i d_v$, and $\sigma_w^i = 0$ for $w \in \bar{H}$. Let $M = \text{diag}(\mu_1, \dots, \mu_k)$, $\hat{B} = M^{-1}B$ and $\zeta = \min\{\|\hat{B}\|_1, \|\hat{B}\|_\infty\}$. Let $\mathbf{x}(t)$ be defined by $x_v(t) = \sum_i \tau_v^i(t)$, and $X(t) = \|\boldsymbol{\tau}(t)\|_1$. Let E_H denote the event that demands spread outside the set H . Then*

1. $\mathbb{P}(E_H) \leq \frac{X(0)}{\zeta} \text{pr}\left(1 - \zeta, \frac{\tau(0)}{X(0)}\right) 1_{\bar{H}}^*$
2. If S in the $(1 - \zeta)$ core of H , then $\mathbb{P}(E_H) \leq \frac{2X(0)h(H)}{\zeta(1-\zeta)}$, where $h(H)$ is the Cheeger ratio of H .

Proof. Let P_ℓ denote the set of all paths of length ℓ from an initial chip in S to \bar{H} such that the first $\ell - 1$ steps are in H . Let $P = \bigcup_{\ell=1}^{\infty} P_\ell$. The key observation is that if $w \in \bar{H}$ ever has demand, then $S_{\pi,c}$ must have occurred for some initial chip c and path π from the location of c to w . Thus we can use the union bound to get that

$$\begin{aligned} \sum_{\pi \in P} \mathbb{P}(S_{\pi,c}) &\leq \sum_{\ell} \sum_{(\pi,c) \in P_\ell} \mathbb{P}(S_{\pi,c}) \\ &\leq \sum_{\ell} \sum_{v_0 \in S} \sum_{c \text{ at } v_0} \sum_{v_\ell \in \bar{H}} \sum_{\pi=(v_0, \dots, v_\ell) \in P_\ell} \mathbb{P}(S_{\pi,c}) \\ &\leq \sum_{\ell} \sum_{v_0 \in S} \sum_{c \text{ at } v_0} \sum_{v_\ell \in \bar{H}} \sum_{\pi=(v_0, \dots, v_\ell) \in P_\ell} \zeta^\ell \prod_{r=0}^{\ell} \frac{1}{d_{v_r}} \\ &= \sum_{\ell} \mathbf{x}(0) \zeta^\ell (D^{-1}A)^\ell 1_{\bar{H}}^* \\ &= \sum_{\ell} \mathbf{x}(0) \zeta^\ell W^\ell 1_{\bar{H}}^* = \frac{X(0)}{\zeta} \text{pr}\left(1 - \zeta, \frac{\mathbf{x}}{X(0)}\right) 1_{\bar{H}}^*. \end{aligned}$$

proving the first statement. The second statement follows the same proof as Theorem 3.2 of [4]. \square

Finally we show how Kronecker PageRank arises in a natural way as the bound of the escape probability of this process.

Theorem 4 *Suppose that the initial demand is contained in $S \subset H \subset V$ with and each vertex $v \in H$ has supply rates $\sigma_v^i \geq \mu_i d_v$. Let $M = \text{diag}(\mu_1, \dots, \mu_k)$, $\hat{B} = M^{-1}B$ and $\zeta = \|\hat{B}\|_1$. Let $X(t) = \|\tau(t)\|_1$, the total amount of demands at time t . Let \mathcal{E}_H denote the event that demands spread outside the set H . Then \mathcal{E}_H can be bounded above using the Kronecker PageRank vector via the relationship:*

$$\mathbb{P}(\mathcal{E}_H) \leq X(0) \mathbf{Kpr} \left(\hat{B}, \frac{\tau(0)}{X(0)} \right) \mathbf{1}_{\bar{H}}$$

Proof. Let f be a vector indicator function of commodity type on chips, that is $f(c) = \mathbf{e}_i$ if c is an i -chip, where \mathbf{e}_i denotes the i th standard basis vector for \mathbb{R}^k . Let C_0 denote the set of initial chips. By the same methods that were used in the proof of Lemma 2, we can bound the probability that demand originating from c ever spreads along a path $\pi = (v_0, v_1, \dots, v_\ell)$ by the sum

$$\mathbb{P}(\mathcal{S}_{\pi,c}) \leq f(c) \hat{B}^\ell \mathbf{1}^* \prod_{r=0}^{\ell-1} \frac{1}{d_{v_r}}$$

Therefore using the same technique as in the proof of Theorem 3 we obtain the bound

$$\begin{aligned} \sum_{\pi \in P} \mathbb{P}(\mathcal{S}_{\pi,c}) &\leq \sum_{\ell} \sum_{u \in S} \sum_{\pi \in B_\ell} \mathbb{P}(\mathcal{S}_{\pi,u}) \leq \sum_{\ell} \sum_{c \in C_0} \sum_{v_\ell \in \bar{H}} \sum_{\pi=(v_0, \dots, v_\ell) \in P_\ell} \mathbb{P}(\mathcal{S}_{\pi,u}) \\ &\leq \sum_{\ell} \sum_{c \in C_0} \sum_{v_\ell \in \bar{H}} \sum_{\pi=(v_0, \dots, v_\ell) \in P_\ell} f(c) \hat{B}^\ell \mathbf{1}^* \prod_{r=0}^{\ell-1} \frac{1}{d_{v_r}} = \sum_{\ell} \tau(0) (D^{-1}A \otimes \hat{B})^\ell \mathbf{1}_{\bar{H}} \\ &= \sum_{\ell} \tau(0) (W \otimes \hat{B})^\ell \mathbf{1}_{\bar{H}} = X(0) \mathbf{Kpr} \left(\hat{B}, \frac{\tau(0)}{X(0)} \right) \mathbf{1}_{\bar{H}} \end{aligned}$$

\square

On the event of non-escape, we would like to guarantee that all demand is satisfied quickly. To make this precise, let \mathcal{S}_t denote the event that all of the vertices are satisfied at time t . In order to complete the analysis of the local case, we would like to bound $\mathbb{P}(\mathcal{S}_t | \bar{\mathcal{E}}_H)$, where \mathcal{E}_H is as in Theorems 3 and 4. Such a bound is not immediately given by Theorem 1. To derive a bound on $\mathbb{P}(\mathcal{S}_t | \bar{\mathcal{E}}_H)$, consider running a modified ‘Dirichlet’ process which is identical to the standard process with the same supply rates, except demand leaving H is ignored. Let \mathcal{S}'_t denote the event that in Dirichlet process, all of the events are satisfied at time t then $\mathbb{P}(\mathcal{S}'_t)$ can be bounded directly by Theorem 1 as this Dirichlet process restricted to vertices in H is the standard process on H . Furthermore $\mathbb{P}(\mathcal{S}_t \cap \mathcal{E}_H) \leq \mathbb{P}(\mathcal{S}'_t)$. Therefore

$$\mathbb{P}(\mathcal{S}_t | \bar{\mathcal{E}}_H) = \frac{\mathbb{P}(\mathcal{S}_t \cap \bar{\mathcal{E}}_H)}{\mathbb{P}(\bar{\mathcal{E}}_H)} \leq \frac{\mathbb{P}(\mathcal{S}'_t)}{\mathbb{P}(\bar{\mathcal{E}}_H)}.$$

Combining this observation along with Theorems 3 and 4, yields that the probability of escape from H is bounded and if the process does not escape from H it dies quickly.

Theorems 3 and 4 can be used in two different ways. As stated, they provide a way to bound the probability demands escape from a given subset. However, they can be also used to construct such a bounding subset. For example, given initial demand $\tau(0)$ contained in an initial set of vertices $S \subset V$, we can algorithmically construct H such that demand stays in H with probability $1 - \epsilon$ by iteratively selecting vertices with the highest (Kronecker) PageRank.

6 An Example

An immediate question is whether anything is actually gained by the introduction of Kronecker PageRank. Suppose we have a spread matrix B , and some initial demand on a graph G . We wish to identify a subset of vertices $H \subset G$ to make shipments to so that the escape probability is at most ϵ . We may use either Theorem 3 or Theorem 4 to identify such a set. The bound afforded by Theorem 4 is clearly sharper than the bound in Theorem 3 as the structure of the spread matrix B is taken into account, but it is not guaranteed that the identified set is actually smaller. In many cases it actually is, though depending on B it may not be significantly smaller.


To illustrate, we give a simple example calculation on synthetic data. Our graph G is an instance of the following random process, which is designed to create a graph which contains tighter clusters that are slightly more sparsely connected to neighboring clusters: We begin with a cycle on 20 vertices. Each vertex is then replaced by an instance of the random graph $G_{20,.3}$, that is a graph 20 vertices and each edge existing independently with probability .3. Inter-cluster edges are then created between vertices in neighboring clusters with probability .05.

We consider the case where $k = 4$, and

$$B = \begin{pmatrix} .8 & .4 & .3 & .2 \\ .2 & .7 & .2 & .1 \\ .1 & .2 & .9 & .3 \\ .4 & .2 & .3 & .6 \end{pmatrix}.$$

The initial demand is given by $\tau(0) = \{2, 1, 2, 0, 0, 1, 0, 0, \dots, 0\}$. In addition we assume $\mu_i = 2$, and $\zeta = .85$.

We demonstrate the difference between Theorems 3 and 4 in the following way. The figure below shows the graph G . The demands start in the large outlined vertices and spread outward from there. Theorem 4 states that with 95% probability, demands stay in the circular vertices. Theorem 3 states that with 95% probability, demands stay in the diamond and circular vertices. This small example illustrates how the Kronecker PageRank can be used to obtain improved results.



References

1. C. BORGS, J. CHAYES, A. GANESH, AND A. SABERI, *How to distribute antidote to control epidemics*, Random Structures & Algorithms, 37 (2010), pp. 204–222.
2. S. BRIN AND L. PAGE, *The anatomy of a large-scale hypertextual web search engine*, Computer Networks and ISDN Systems, 30 (1998), pp. 107–117.
3. Y. CHEVALEYRE, P. DUNNE, U. ENDRISS, J. LANG, M. LEMAITRE, N. MAUDET, J. PADGET, S. PHELPS, J. RODRIGUEZ-AGUILAR, AND P. SOUSA, *Issues in multiagent resource allocation*, Informatica, (2006).
4. F. CHUNG, P. HORN, AND A. TSIATAS, *Distributing antidote using PageRank vectors*, Internet Mathematics, 6 (2009), pp. 237–254.
5. G. DAHLQUIST, *Stability and error bounds in the numerical integration of ordinary differential equations*, Kungl. Tekn. Hgsk. Handl. Stockholm. No., 130 (1959), p. 87.
6. A. GANESH, L. MASSOULIE, AND D. TOWSLEY, *The effect of network topology on the spread of epidemics*, vol. 2, IEEE, 2005, pp. 1455–1466.
7. R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, 1990.
8. I. Z. KISS, D. M. GREEN, AND R. R. KAO, *Infectious disease control using contact tracing in random and scale-free networks*, Journal of the Royal Society, Interface / the Royal Society, 3 (2006), pp. 55–62. PMID: 16849217.
9. C. MERINO, *The chip-firing game*, Discrete Mathematics, 302 (2005), pp. 188–210.
10. M. NEWMAN, *Spread of epidemic disease on networks*, Physical Review E, 66 (2002).

7 Appendix

In this section we establish the differential equation used in the proof of Theorem 1, namely:

Lemma 5 *If $\tau(t)$ denotes the demand vector at time t , A is the adjacency matrix of G , B denotes the spread matrix and $S = \text{diag}(\sigma)$ denotes the supply matrix, then*

$$\frac{\partial}{\partial t} \mathbb{E}[\tau(t)] = \mathbb{E}[\tau(t)](A \otimes B - S).$$

The key of Lemma 5 are the following two well known and simple facts concerning exponentially distributed random variables. We use the notation $f(h) = O_{h \rightarrow 0}(g(h))$ to indicate that $f(h) \leq Cg(h)$ for h sufficiently small.

Lemma 6 *Suppose X is an exponentially distributed waiting time with rate λ , then*

$$\mathbb{P}(X < h) = \lambda h + O_{h \rightarrow 0}(h^2).$$

An immediate corollary is

Lemma 7 *Suppose X, Y are independent exponentially distributed waiting times with rates λ_1, λ_2 . Then*

$$\mathbb{P}(X, Y < h) = O_{h \rightarrow 0}(h^2).$$

Proof of Lemma 5.

Fix a vertex v and commodity i . We will show that

$$\frac{\partial}{\partial t} \mathbb{E}[\tau_v^i(t)] = [\mathbb{E}[\tau(t)](A \otimes B - S)]_v^i,$$

Since this holds for all v , and i the result will follow.

To do this, we compute the derivative by the definition, that is we compute

$$\lim_{h \rightarrow 0} \frac{\mathbb{E}[\tau_v^i(t) - \tau_v^i(t+h)]}{h}.$$

To do this, consider the conditional expectation, $\mathbb{E}[\tau_v^i(t) - \tau_v^i(t+h) | \tau(t)]$. Note that by Lemma 7, then probability that two independent events (either two spread events, or two satisfy events or a spread and a satisfy event) occur is $O_{h \rightarrow 0}(h^2)$. On the other hand, given a neighbor u of v , and a commodity j , the probability of a spread event originating from this neighbor and commodity in time $(t, t+h)$ is exactly $B_{ji}\tau_u^j(t)h + O_{h \rightarrow 0}(h^2)$. Likewise, the probability of a satisfaction event in this time is $\tau_v^i(t)\sigma_v^i h + O_{h \rightarrow 0}(h^2)$. Linearity of expectation yields

$$\begin{aligned} \mathbb{E}[\tau_v^i(t) - \tau_v^i(t+h) | \tau(t)] &= \mathbb{E}[\tau(t)(A \otimes B - S)h + O_{h \rightarrow 0}(h^2) | \tau(t)] \\ &= \tau(t)(A \otimes B - S)h + O_{h \rightarrow 0}(h^2) | \tau(t). \end{aligned}$$

In particular, note that the $O_{h \rightarrow 0}(h^2)$ term means that there is a (large) constant $C = C(\tau(t), A, B)$, so that $O_{h \rightarrow 0}(h^2) \leq C \cdot h^2$ for $h \leq 1$. Note that due to our conditioning this constant depends on $\tau(t)$, but critically not on h . Note that this constant is $\sigma(\tau(t))$ -measurable.

By the tower property of conditional expectation,

$$\begin{aligned} \lim_{h \rightarrow 0} \frac{\mathbb{E}[\tau_v^i(t) - \tau_v^i(t+h)]}{h} &= \lim_{h \rightarrow 0} \frac{\mathbb{E}[\mathbb{E}[\tau_v^i(t) - \tau_v^i(t+h) | \tau(0)]]}{h} \\ &= \lim_{h \rightarrow 0} \frac{\mathbb{E}[\tau(t)(A \otimes B - S)h] + \mathbb{E}[\mathbb{E}[O_{h \rightarrow 0}(h^2) | \tau(t)]]}{h} \\ &= \tau(t)(A \otimes B - S) + \lim_{h \rightarrow 0} \mathbb{E}[O_{h \rightarrow 0}(h)]. \end{aligned}$$

It suffices to show that

$$\lim_{h \rightarrow 0} \left| \mathbb{E}[O_{h \rightarrow 0}(h)] \right| \leq \lim_{h \rightarrow 0} \mathbb{E}[|O_{h \rightarrow 0}(h)|] = 0.$$

But recall that the $O_{h \rightarrow 0}(h)$ term is bounded by $C(\boldsymbol{\tau}(t), A, B) \cdot h$ for $h \leq 1$. Thus it is enough to show that

$$\lim_{h \rightarrow 0} \mathbb{E}[|C(\boldsymbol{\tau}(t), A, B)h|] = 0.$$

This follows from the monotone convergence theorem, so long as

$$\lim_{h \rightarrow 0} \mathbb{E}[|C(\boldsymbol{\tau}(t), A, B)|] < \infty.$$

To complete the proof we note that we can give an upper bound on $C(\boldsymbol{\tau}(t), A, B)$ in terms of $\|\boldsymbol{\tau}(t)\|_1$, n and $\max\{b_{i,j}\}$. Indeed, the rates of the active point processes are at most $\|\boldsymbol{\tau}(t)\|_1 \max\{b_{i,j}\}$; and thus the probability that any pair of point processes both have events in the period $(t, t + h)$ is bounded by

$$C(\boldsymbol{\tau}(t), A, B)h \leq C\|\boldsymbol{\tau}(t)\|_1^2 \max\{b_{i,j}\}^2 n^2 h.$$

But $\|\boldsymbol{\tau}(t)\|_1 \ll n \cdot N \cdot k$, where $n = |G|$, k is the number of demands and N is our uniform upper bound for the demand at a point (indeed, this is the precisely the motivation for such a bound.) \square