

An Axiomatic Approach to Routing

Omer Lev
Hebrew University and
Microsoft Research, Israel
omerl@cs.huji.ac.il

Moshe Tennenholtz
Technion
moshet@ie.technion.ac.il

Aviv Zohar
Hebrew University and
Microsoft Research, Israel
avivz@cs.huji.ac.il

ABSTRACT

Information delivery in a network of agents is a key issue for large, complex systems that need to do so in a predictable, efficient manner. The delivery of information in such multi-agent systems is typically implemented through routing protocols that determine how information flows through the network. Different routing protocols exist each with its own benefits, but it is generally unclear which properties can be successfully combined within a given algorithm. We approach this problem from the axiomatic point of view, i.e., we try to establish what are the properties we would seek to see in such a system, and examine the different properties which uniquely define common routing algorithms used today.

We examine several desirable properties, such as robustness, which ensures adding nodes and edges does not change the routing in a radical, unpredictable ways; and properties that depend on the operating environment, such as an “economic model”, where nodes choose their paths based on the cost they are charged to pass information to the next node. We proceed to fully characterize minimal spanning tree, shortest path, and weakest link routing algorithms, showing a tight set of axioms for each.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Computer-Communication Networks—*Network Architecture and Design*; C.2.2 [Computer Systems Organization]: Computer-Communication Networks—*Network Protocols*; C.2.6 [Computer Systems Organization]: Computer-Communication Networks—*Internetworking*; G.2.2 [Mathematics of Computing]: Discrete Mathematics—*Graph Theory*

General Terms

Design, Algorithms, Theory

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Tark 15 Pittsburgh, Pennsylvania USA
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

Keywords

Routing, Axiomatic Approach, Routing algorithms

1. INTRODUCTION

The proper way to distribute power, disseminate information, or establish hierarchies in organizations is an issue encountered whenever there is a large enough network of agents that needs to interact in an orderly manner. For example, when trying to establish efficient lines of communications between agents which all need to reach a central hub, there are various properties we may desire in our system. We might want the system to be able to handle small changes in connections without causing disruptions throughout the network; we may want it to be flexible when we change its parameters so that various routing options are possible, and more. Indeed, the search for the right communication structure has played a role in early work on the foundations of the area of multi-agent systems [7, 15, 5], based on classical work in organization theory [8, 16].

More concretely, examining networking, one of the most important aspects of the design of a communication network is the way it routes information through its physical links. Routing protocols, such as those used in packet switching networks, circuit switching, or ad-hoc networks are designed with many goals in mind. They must adapt to changing network conditions, withstand failures, and operate in a distributed fashion while constructing a “good” routing scheme. Nodes in the network are, in fact, autonomous agents that can control the flow of information through them and can choose to forward it according to their own considerations. Agents may be controlled by different economic entities (such as in the internet, where different internet service providers control some of the routers), and may route according to complex preferences that are derived from economic relations [9, 14]. Even in the cooperative local-network setting where all routers are controlled by a single network operator, different considerations such as bandwidth utilization, latency, and the risks of link failures come into play.

The multitude of previous treatments of the problem suggest a myriad of routing protocols, each with their own benefits and shortcomings. In contrast, this work examines the routing problem through the lens of the *axiomatic approach*, which seeks to formulate different elementary properties that are desirable in this context. One approach to an axiomatic treatment, which we take in this work, is that of characterization: a set of elementary properties is shown to uniquely determine some routing algorithm, and hence the routing

outcome on any specific graph. From the designer’s perspective, such a result implies a great deal – any additional property that is not already achieved by the protocol cannot be added to it without giving up on another basic property. The approach thus provably bounds the design space of algorithms and makes explicit the choices made when selecting one over the other.

As we are not aware of any previous axiomatic treatment of routing, we focus our attention on a domain that most closely resembles the internet as it is built today, and focus our efforts within this domain on what one may consider classic, or natural routing schemes. In particular, we assume that routing choices are independent of the congestion on links (such is the case in the internet, where routing protocols such as BGP first establish paths, and congestion control protocols such as the one embedded into TCP manage the load on each flow’s path and ensures that rates are throttled to match the bottleneck of the flow). Furthermore, as with internet routing where routers decide on the next hop of each packet using a routing table that maps its destination to the next hop, routing choices made to different destinations are done independently. Finally, packets addressed to the same destination are not split between different paths, and are routed in the same manner regardless of their source. These choices, which greatly restrict the power of any routing algorithm may seem arbitrary, but are in fact derived from real-world design considerations. For example, the need to quickly forward packets towards their destination at each router mandated that most routing be done in specialized hardware. No complex computation is performed (only a lookup into a routing table) and no deep inspection of the packet is performed. Keeping routing simple has made it fast and robust.

More advanced routing schemes that have been proposed in the literature may split traffic, allow routing choices to depend on the source of the packet or its previous hops, or may even change the routes in response to link congestion. These are notoriously difficult to coordinate and to implement. We leave treatment of these more advanced schemes to future work.

Our set of axioms or “desirable properties” are also motivated by similar considerations. For example, one of the fundamental features we desire in our algorithms is one of **robustness**, which is the ability of a system to endure changes in the network without creating disruption in parts of the network that have not undergone changes.

A different feature, which might be desirable only in certain cases, is “**first hop**”, which is particularly relevant for diffuse networks with independent nodes. It means, broadly, that network nodes care only about their immediate surroundings, or the “next step” in the network data transfer. Such a property might be relevant when nodes pursue an “economic model”, paying for transferring information, and hence only caring about the cost they need to pay to move their information to the next node, and following that, they have no preference on the route the information should pass en route to its destination. Other properties, desirable only in some cases include an indifference between two parallel paths, as long as they change their weights by the same amount concurrently.

Ultimately, after devising our axioms we successfully fully characterized 3 natural routing algorithms:

- **Minimum spanning tree:** A tree with the smallest

overall weight is a result, among others, of the “first hop” axiom (the “economic model”).

- **Shortest path:** A tree where each node has the shortest possible path to its destination is a result, among other axioms, of viewing as immaterial to the routing decision any parallel paths which change their weight by the same amount.
- **Weakest link tree:** A tree where each node takes the path with the maximal “lightest” weight available to it. This results from considering higher edge-weights as beneficial (e.g., representing bandwidth which one wishes to increase in contrast to delay that one wishes to decrease), and from considering designers that choose between parallel paths in a slightly different manner.

We proceed to review relevant previous research and then continue to define our model and expand on the axioms, which are motivated with a brief explanation and presented formally. Following that we show (and prove) our characterization of the minimal spanning tree, the shortest path tree, and the weakest link tree.

2. RELATED WORK

In the past decade, as routers became more flexible, research on routing (particularly inter-domain) and its techniques has been rekindled and extended beyond the technical issues dealt with in the past. The harbinger for much of this research was [10], which was further expanded by several researchers (see updating report here: <http://www.cl.cam.ac.uk/~tgg22/metarouting/>). However, this line of research, while introducing many interesting mathematical and theoretical concepts to the field of routing, has refrained from phrasing its models as requirements by users, to be filled by various routing algorithms.

The axiomatic approach, which does approach problems with this outlook, has been first introduced in CS contexts as extensions to the classical theory of choice [4], and has been applied to ranking systems [1, 2] and trust systems [3], as well as to other multi-agent setups such as multi-level marketing [6].

In relation to networking, usage of the axiomatic approach has generally been concentrated in two main areas: applying to general graph theory (e.g., [18]) or in more technical approaches to networks: papers such as [13] which deal with particular wireless models and implementations, and, somewhat closer to our line of work, [12], whose basic axioms are basic enough to be covered through our models, while the routing related axioms involve various assumptions on how routers work (tables, etc.), which we refrain from approaching in our more abstract considerations.

Further work connecting networking and the axiomatic approach has focused on particular instances of problems: [11] try to use the axiomatic approach to extract the costs of multicast routing and decide who is to pay them. Trust networks and social networks (e.g., recommendation systems) have been analysed many times using the axiomatic approach to understand their desirable features and better understand desirable algorithms in these cases [17, 3]. However, none of these papers deal with the basic routing mechanism by which messages and information arrive at each node.

3. SETUP

Before introducing our axioms, we begin by setting up our routing model. It is, naturally, only a simplification of routing as it is done in large, complex networks such as the internet, but we believe it is robust enough to display many networking characteristics.

Our world will be a weighted graph $G(V, E, W)$ and a destination d , where V is a set of nodes, E is a set of edges, and W is a function assigning weights to edges, and $d \in V$. A routing solution is a tree T over that graph, as defined below (we do not concern ourselves with non-tree routing, as passing through the same node several times does not serve any purpose).

DEFINITION 1. A routing function $f_d : \mathcal{G} \rightarrow \mathcal{T}$ is a function from connected weighted graph $G(V, E, W) \in \mathcal{G}$ in which $d \in V$, to a tree $T(V, E, W) \in \mathcal{T}$ such that $T \subseteq G$.

We can look at the graph as one with directed edges if we consider each edge's direction to be the one pointing at the vertex from which there is a path to d (without going through the same edge again).

We discuss 3 different routing options:

- *Minimum spanning tree (MST)*: a tree connecting all nodes in the graph with the minimal weight, i.e., for every tree $T' \subseteq G$ that encompasses all of G 's nodes, $\sum_{e \in f_d(G)} W(e) \leq \sum_{e \in T'} W(e)$.
- *Shortest path*: each node is connected to d using a shortest length path in the graph. For every node $v \in V$, let (e_1, \dots, e_s) be a path without cycles from v to d such that $e_i \in T$, and let (e'_1, \dots, e'_k) a different path from v to d , then $\sum_{i=1}^s W(e_i) \leq \sum_{j=1}^k W(e'_j)$.
- *Weakest link*: looking at each potential path from each node to d , we give each path the value of its smallest valued edge. The routing tree will contain, for each node the path to d with the maximal value. So for every node $v \in V$, let (e_1, \dots, e_s) be a path without cycles from v to d such that $e_i \in T$, and let (e'_1, \dots, e'_k) a different path from v to d , then $\min_{1 \leq i \leq s} W(e_i) \geq \min_{1 \leq j \leq k} W(e'_j)$.

Notice that while for the minimal spanning tree and shortest path routing options weights are interpreted as costs (e.g. payments, delays), so these algorithms seek to minimize them, the weakest link views weights as measure for capability such as bandwidth, so seeks to maximize the weight.

4. AXIOMS

Having introduced our framework, we introduce our axioms, which are, basically, desirable properties of the function f_d (in the axioms below we use f , as these are properties which do not depend on a specific d destination).

Robustness indicates the routing being quite unsusceptible to changes – only if a path in the routing is destroyed, will it require any change. As indicated in Figure 1, the path from node a changes, but not from node b .

AXIOM 1 (ROBUSTNESS). f is robust if removing an edge $e \in E$ from $G(V, E, W)$, yielding G' , then for every vertex $v \in V$: if the cycle-less path from v to d in $f_d(G)$ did not contain e , then this is still the selected path according to $f_d(G')$ (see example Figure 1).




Figure 1: An edge is removed, but only a , whose path used that edge changes its path (the left side is the graph, the right side is the routing algorithm's output)

The following axioms deal with global changes to the graph weights, additive or multiplicative:

AXIOM 2 (SCALE INVARIANCE). f is scale invariant if for a graph $G(V, E, W)$, for any positive scalar $\alpha \in \mathbb{R}_+$, defining $G'(V, E, \alpha W)$, for every $d \in V$, $f_d(G) = f_d(G')$.

AXIOM 3 (SHIFT INVARIANCE). f is shift-invariant if for a graph $G(V, E, W)$, for any $\alpha \in \mathbb{R}$, defining $G'(V, E, \alpha + W)$, for every $d \in V$, $f_d(G) = f_d(G')$.

The monotonicity axiom below seeks to establish that if an edge does not have to be in every tree, if its weight increases enough, it will not be a part of the routing tree:

AXIOM 4 (MONOTONICITY). f is monotone if for a graph $G(V, E, W)$ and $d \in V$, for $e' \in E$, if $e' \notin f_d(G)$, then for every $G'(V, E, W')$, there is a value $M_{W'}$ such that for W'' such that $W''(e) = W'(e)$ for all $e \in E \setminus \{e'\}$ and $W''(e') \geq M_{W'}$, $e' \notin f_d(G''(V, E, W''))$. Similarly, we can define the opposite direction, an edge in $f_d(G)$ will not be in the routing tree if it has a small enough value; we will refer to it as inverse monotonicity.

While the phrasing of the following axiom is somewhat technical, the *first hop* axiom below simply means that if a vertex has several potential edges to connect to a path to d , the routing only depends on the weights of the edges connecting it to these potential paths, and unrelated to weights of other edges in the graph.

AXIOM 5 (FIRST HOP). Let $G(V, E, W)$ be a weighted graph and let $v, d \in V$ and $d \neq v$. Suppose $C = \{c_1, \dots, c_s\}$ are the vertices such that $(v, c_i) \in E$ and there is a path from c_i to d in $f_d(G)$ which does not pass through v . W.l.o.g., let (v, c_1) be the first step in the path from v to d in $f_d(G)$. We say that f satisfies first hop if for any W' such that $W'(v, c_i) = W(v, c_i)$ and if for all $c_i \in C$ $f_d(G'(V, E, W'))$ contains paths to d from c_i that do not pass through v , and




Figure 2: Selected path does not change when each path from the top node is added 2.

there is no $c' \notin C$ such that $(v, c') \in E$ and there is a path from c' to d in $f_d(G')$, then the cycle-less path from v to d in $f_d(G'(V, E, W'))$ starts with (v, c_1) .

The rationale for the *first hop* axiom is to capture a common economic model, in which edge weights indicate the cost of passing information. In distributed networks, such as the internet, each agent only minds the amount it needs to pay to transfer its data to the next node, not caring about the path the data will take from there.

Path cardinal/ordinal invariance intends to see the planner's considerations when multiple paths exist. As there might be many potential behaviours, we only limit ourselves to examining the narrow case of what the planner considers important when there is only one cycle in the graph (i.e., the axiom does not strongly enforce a general behaviour on the planner). Cardinal invariance deals with adding the same weight to potential paths, and how it does not effect the routing. Ordinal invariance similarly does not change the routing if all that has changed are the weights of the competing paths, as long as edges in each path maintain their relative position.

AXIOM 6 (PATH CARDINAL INVARIANCE). Let $G(V, E, W)$ be a graph which contains a single cycle, $d \in V$, and let $d \neq v \in V$ be a part of this cycle. Hence there are two alternative paths from v to d – $p_1 \subset E$ and $p_2 \subset E$ (one of them is actually a part of $f_d(G)$). f is path cardinal invariant if it treats those paths as such: Choosing an edge $e' \in p_1$ and $e'' \in p_2$, for any $\alpha \in \mathbb{R}$, we define W' as $W(e) = W'(e)$ for $e \in E \setminus \{e' \cup e''\}$ and $W'(e') = W(e') + \alpha$ and $W'(e'') = W(e'') + \alpha$, the path from v to d will not change in $f_d(G(V, E, W'))$ (see example Figure 2).

AXIOM 7 (PATH ORDINAL INVARIANCE). Let $G(V, E, W)$ be a graph which contains a single cycle, $d \in V$, and let $d \neq v \in V$ be a part of this cycle. Hence there are two alternative paths from v to d – $p_1 \subset E$ and $p_2 \subset E$ (one of them is actually a part of $f_d(G)$). f is path ordinal invariant if it treats those paths as such: Taking an edge $e' \in p_i$ ($i \in \{1, 2\}$) that is not maximal or minimal in $p_1 \cup p_2$, we define W' as $W(e) = W'(e)$ for $e \in E \setminus \{e'\}$ and allow $W'(e')$ to be any value it chooses as long as for every $e'' \in p_i$ if $W(e') \geq W(e'')$ then $W'(e') \geq W'(e'') = W(e'')$, and the path from v to d will not change in $f_d(G(V, E, W'))$ (see example Figure 3).




Figure 3: Selected path does not change when the bottom right edge is slightly increased.

5. MINIMAL SPANNING TREE

THEOREM 1. A robust, scale invariant, shift invariant, monotone, first-hop (axioms 1-5) routing function f , for any graph $G(V, E, W)$ and $d \in V$, $f_d(G)$ will always be a minimal spanning tree of G .

REMINDER 1. As our minimal spanning tree proof relies on the Kruskal algorithm, we will briefly describe it:

1. Order edges according to weights
2. Define a set S , initialized to the empty set.
3. Going over edges from lightest to heaviest, if the set $S \cup \{e\}$ has no cycles, $S = S \cup \{e\}$.

PROOF OF THEOREM 1. We shall prove the theorem using complete induction on the number of non-cycle lightest edges in the tree $f_d(G)$. Hence, we shall begin by proving that the lightest edge in the graph G is in the routing tree $T = f_d(G)$. Assuming we are mistaken, let us consider the lightest edge in $G - e = (u, v) \in E$ – and assume $e \notin T$. We create $G'(V, E', W) = f_d(G) \cup \{e\}$, and thanks to the robustness axiom, we know $f_d(G) = f_d(G')$.

If v 's path to d in $f_d(G')$ goes through u , we shall switch the nodes' names, so that v 's path to d does not pass through u . As e is not in $f_d(G')$, there is an edge $e' = (u, s)$ that is the first step from u towards d . We now define $x = W(e)$ and $y = W(e')$, and due to our minimality assumption, we know $x < y$.

Using the monotonicity axiom, we change graph weights to W' that is identical to W except that e' weight is large enough so that we create a tree T' in which there is a path from s to d that does not pass through u (e.g., the same path that is in $f_d(G')$), and v passes through u towards d (i.e., $e \in T'$). We define $y' = W''(e')$.

Using scale invariance we now multiply all edges by $\frac{y-x}{y'-x}$, and using shift invariance, we add to all edges $y - \frac{y-x}{y'-x}y'$. This means the weight of edge e is now

$$x \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = (x-y') \frac{y-x}{y'-x'} + y = x$$

While the weight of edge e' is now

$$y' \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = y$$

However, the routing tree contains e and not e' , and a path from both v and s to d , contradicting the “first hop” axiom, which should have caused e' to be chosen over e , as the edge weights for e and e' have not changed.

We now turn to the induction step – we assume all bottom weighted $k-1$ edges that do not create a cycle are included in

the tree $T = f_d(G)$, and we now seek to include the k -lightest edge that does not create a cycle. We pursue a similar path as we did as previously, and we shall mark the edge as $e = (u, v)$, and assume it is not included in $T = f_d(G)$, and instead $e' = (u, s)$ is included, and there is a path to d from v and s . Again, we create $G'(V, E', W) = f_d(G) \cup \{e\}$, and thanks to the robustness axiom, we know $f_d(G) = f_d(G')$. Using monotonicity we create weights W' that just increase e' weight, so that $G'' = (V, E', W')$ has $T' = f_d(G'')$ which include the same bottom k which do not create cycles (from the induction hypothesis), and u reaches d via the edge e . Recall that we know the bottom $k - 1$ edges will definitely be in $f_d(G'')$, and we wish to ensure that there will still be a path from v to d and from s to d . The same arguments used in the initial step of the induction ensure that, as well as returning the weights of e and e' to their values in G , while routing u through e and not e' in the routing tree, reaching a contradiction with our initial assumption due to the “first hop” axiom.

What is left is to show MST indeed follows our axioms:

Robustness (axiom 1) Trivial thanks to the Kruskal algorithm – if the removed edge (e') was not in the routing tree, it means it was not selected in the first place, and hence the same routing tree will be chosen. If it was, then any edge added after its removal (e'') closed a cycle with it, and hence, if affecting the edges in any path that did not include e' , it means e'' closes a cycle with them, hence e' would have closed a cycle as well.

Scale invariance (axiom 2) Multiplying all edges by a fixed amount does not change their order in relation to others, hence Kruskal will choose the same routing tree.

Shift invariance (axiom 3) Adding a fixed amount to all edges does not change their order in relation to others, hence Kruskal will choose the same routing tree.

Monotonicity (axiom 4) Giving an edge the maximal possible edge value ensures it will only be selected if no other edge can replace it – and if there exists a tree without some edge, we know it will be chosen before.

“First hop” (axiom 5) Kruskal ensures that if there are the same possible options of connecting a node to the tree, only the lightest edge will be chosen.

□

6. SHORTEST PATH

THEOREM 2. *A robust, scale invariant, monotone, and path cardinal invariant (axioms 1-2, 4, 6) routing function f , for any graph $G(V, E, W)$ and $d \in V$, $f_d(G)$ will always be a shortest path graph to d of G .*

PROOF. Suppose $T = f_d(G)$ is not a shortest path routing tree. Let u be the closest node to d that is not connected to d with a shortest path. Hence, there is an edge $e = (u, v)$ which will make u 's path a shortest path one (v , being closer to d , is already connected to d with a shortest path), but $e \notin T$, and instead $e' = (u, s)$ is included in T . Using robustness, we create $G'(V, E', W) = T \cup e$. G' contains two alternate paths from u to d , and $f_d(G) = f_d(G')$.

Using path cardinal invariant, we “move” all the value of the edges on each path to its “source”, i.e., to (u, v) or (u, s) (we do this by adding to the weight of (u, v) and (u, s) the value of $\sum_{e \in (p_1 \cup p_2) \setminus (p_1 \cap p_2)} W(e) - W((u, v)) - W((u, s))$, and reduce from $W((u, s))$ the weight of all edges of the path from u to d through (u, v) and vice versa). We shall refer to $W(e) = x$ and $W(e') = y$. We now use monotonicity to create a new tree, with e but without e' , with the graph's weight now W' (identical to W except for increase in e' weight). Once again, we transfer all value of the paths from u to d to e and e' respectively, with everything else being 0. Now, using monotonicity, we increase the weight of e' above that of e , with the weight of (u, v) being x (its path weights have not changed) and (u, s) being y' .

Finally, we multiply all edges by $\frac{y'-x}{y'-x}$ (using scale invariance), and using path cardinal invariance, we add to e and e' the amount $y - \frac{y'-x}{y'-x}y'$. The weight of e is now:

$$x \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = (x-y') \frac{y-x}{y'-x} + y = x$$

While the weight of edge e' is now

$$y' \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = y$$

As all edges are the same weight as before, therefore we reached a contradiction regarding the inclusion of e' instead of e (whose weights are the same as well).

We will now show shortest path follows our axioms:

Robustness (axiom 1) Removing an edge, at most, eliminates a potential path from a node to the destination d . If the path was not on the shortest path, the previous shortest path remains so.

Scale invariance (axiom 2) Multiplying by a fixed amount all edges means the value of each path is multiplied by the same amount, maintaining their relative ordering, hence what was shortest remains so.

Monotonicity (axiom 4) Giving an edge the value of the sum of all other edges ensures it will only be selected if no other path can replace it — and if there exists a tree without some edge, we know there is such a path.

Path cardinal invariance (axiom 6) Having multiple paths from a node, adding the same amount to each path doesn't change the ordering of the paths (i.e., which path is “shorter” than another), hence selection of shortest path will be identical.

□

7. WEAKEST LINK

THEOREM 3. *A robust, scale invariant, shift invariant, inverse monotone, and path ordinal invariant (axioms 1-4, 7) routing function f , for any graph $G(V, E, W)$ and $d \in V$, $f_d(G)$ will always be a weakest link graph to d of G .*

PROOF. Suppose $T = f_d(G)$ is not a weakest link routing tree. Let u be a node that requires just one edge missing from T that is not connected to d with a weakest link¹, and

¹such a node exists as there is a node not connected by weakest link in T , hence adding the necessary path for that node, taking the node just before the final edge that we add to T (i.e., closest to d), answers our criterion.

we mark this edge as $e = (u, v)$. Since $e \notin T$, there is an edge instead $e' = (u, s)$ that is included in T . Using robustness, we create $G'(V, E', W) = T \cup e$. G' contains two alternate paths from u to d , and $f_d(G) = f_d(G')$.

Using path ordinal invariant, we change the value of all edges on each alternate path from u to d to its weakest link value (we do this by taking the 2nd smallest edge in the path and changing its value to that of the weakest link, which by the axiom does not change the path chosen, and we proceed doing so to all edges on the path). We shall refer to $W(e) = x$ and $W(e') = y$ (from assuming u is not in a weakest link path we know $x > y$). Using inverse monotonicity, we create W' identical to W except for e' weight, that is low enough that it is not included in $f_d(G''(V, E', W'))$. Once again, we change the values of the paths from u to d to their weakest link value (this is only relevant for the path through e' , as the other path has not changed). We term the the new value for $e' - y'$, and we know $x > y'$.

Finally, we multiply all edges by $\frac{y-x}{y'-x}$ (using scale invariance), and using shift invariance, we add to all edges $y - \frac{y-x}{y'-x}y'$. Edge e now has the weight:

$$x \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = (x-y') \frac{y-x}{y'-x} + y = x$$

While the weight of edge e' is now

$$y' \frac{y-x}{y'-x} + y - \frac{y-x}{y'-x}y' = y$$

As all edges are the same weight as before, hence we reached a contradiction regarding the inclusion of e' instead of e .

We shall now show weakest link also follows our axioms:

Robustness (axiom 1) Removing an edge, at most, eliminates a potential path from a node to the destination d . If the path was not a weakest link, the previous weakest link remains so.

Scale invariance (axiom 2) Multiplying by a fixed amount all edges means the value of each path (its smallest edge) is multiplied by the same amount, maintaining their relative ordering, hence what was weakest link remains so.

Shift invariance (axiom 3) Adding a fixed amount all edges means the value of each path (its smallest edge) is added the same amount, maintaining their relative ordering, hence what was weakest link remains so.

Monotonicity (axiom 4) Giving an edge the value of the minimum of all other edges ensures it will only be selected if no other path can replace it — and if there exists a tree without some edge, we know there is such a path.

Path ordinal invariance (axiom 7) Having multiple paths from a node, the weakest link edge (the one with smallest value) of the selected path can't become lower than the weakest link of the non-selected path, hence weakest link choice does not change.

□




Figure 4: Lack of robustness results in a minimal spanning tree/shortest path routing (above) ending up in a routing tree that is weakest link but not MST or shortest path (below).

8. TIGHTNESS OF AXIOMS

We will now show that the above characterizations are tight, and that without each axiom, other routing algorithms become possible.

THEOREM 4. *All MST axioms (1-5) are necessary, and without even one of them, other routing algorithms are possible.*

PROOF. Going over all MST axioms, we detail potential algorithms which work with all axioms except that one, and are not MST. We will refer below to each relaxed axiom, and to the new/additional system which can be obtained by that relaxation:

Robustness See example in Figure 4. Apply MST to any other graph that is not a linear transformation of the bottom one.

Scale invariance See example in Figure 5. On all graphs except those which contain as a subgraph a linear transformation of the bottom one, apply MST.

Shift invariance See example in Figure 6. On all graphs except those which contain as a subgraph a linear transformation of the bottom one, apply MST.

Monotonicity A maximal spanning tree implements all axioms but monotonicity.

First Hop Weakest link implements all of the other axioms.

□

THEOREM 5. *All shortest path axioms (1-2, 4, 6) are necessary, and without even one of them, other routing algorithms are possible.*

PROOF. Going over all shortest path axioms, we detail potential algorithms which work with all axioms except one, and are not shortest path. We will refer below to the each relaxed axiom, and to the new/additional system which can be obtained by that relaxation:




Figure 5: Eliminating scale invariance results in a minimal spanning tree/shortest path/weakest link routing (above) ending up in neither (below).




Figure 6: Eliminating shift invariance results in a minimal spanning tree/weakest link routing (above) ending up in neither (below).

Robustness See example in Figure 4. Apply shortest path to any other graph that isn't a scale of the structure of the bottom one. Any edge in that structure that is 100 times all the others is removed in the tree.

Scale invariance See example in Figure 5. Taking the bottom example and for the group that includes all graphs for which it is a subgraph and those that can be formed by path cardinal invariance, and only for them do not apply shortest path but rather the example (it won't trample on the top example, as if the upper example adds y to lower-right edge, and y to the rest, and the bottom example adds x , it would require $2 + x = 4 + x$, reaching an impossibility).

Inverse Monotonicity A longest path tree implements all axioms but monotonicity.

Path cardinal invariant Minimal spanning tree implements all other axioms.

□

THEOREM 6. All weakest link axioms (1-4, 7) are necessary, and without even one of them, other routing algorithms are possible.

PROOF. Going over all weakest link axioms, we detail potential algorithms which work with all axioms except one, and are not weakest link. We will refer below to the each relaxed axiom, and to the new/additional system which can be obtained by that relaxation:

Robustness See example in Figure 4. Apply weakest link to any other graph that isn't of the structure of the bottom one, Any edge that in that structure that is 100 times less than all the others' weight is removed.

Scale invariance See example in Figure 5. Taking the bottom example and for the group that includes all graphs for which it is a subgraph and those that can be formed by shift invariance and only for them do not apply weakest link but rather the example (it won't trample on the top example, as it can't be reached by shift invariance, and as the edge weights are all minimal/maximal, they change change by ordinal invariance).

Shift invariance See example in Figure 6. Taking the bottom example and for the group that includes all graphs for which it is a subgraph and those that can be formed by scale invariance and only for them do not apply weakest link but rather the example (it won't trample on the top example, as it can't be reached by shift invariance, and as the edge weights are all minimal/maximal, they change change by ordinal invariance).

Monotonicity A strongest link tree implements all axioms but monotonicity.

Path ordinal invariant Minimal spanning tree implements all other axioms.

□

9. DISCUSSION

In this paper we explore the basic issue of routing – how should information flow through a network and what properties might this process have. In the process of considering this issue we developed several properties we believe might be desirable by system planners. For example, *robustness*, or the ability of a routing protocol to keep small changes from disrupting the whole routing process, is a property especially required in fast, changing networks.

Naturally, creating a structure from possible interactions between agents defined by a connections' graph is not limited just to information routing in networks such as the internet. Looking at organizations, where workers are connected according to their ability to work with other workers, and instead of routing messages between them we seek to construct an organizational hierarchy, we face a similar challenge. Again, robustness is a desirable property, as it means that if some workers have a worsening relationship with others, if they're not very senior in the organization, it has little effect on many others. In this case, we may consider the "economic model" ("first hop" axiom) appropriate as well – if workers only interact with their boss, we only care about the edge from each worker to his/her boss, and each worker does not care what happens further up in the hierarchy².

Beyond setting up the axioms, we also examined common routing algorithms – minimal spanning tree, shortest path and weakest link, and fully characterized them. Obviously, this is only the beginning of the road for this line of research – further steps will entail developing more axioms and using them to characterize more algorithms, with the aim of giving a set of tools for system designers, allowing them to choose desirable properties which would dictate appropriate routing protocols.

²Similarly, in a highly centralized organization, a path cardinal invariance is probably a sensible axiom.

Acknowledgments

The authors thank Michael Schapira for his insightful discussions on this matter. Moshe Tennenholtz carried out this work while at Microsoft Research, Israel. Aviv Zohar is supported in part by the Israel Science Foundation (Grants 616/13 and 1773/13), and by the Israel Smart Grid (ISG) Consortium.

10. REFERENCES

- [1] A. Altman and M. Tennenholtz. Ranking systems: The pagerank axioms. In *Proceedings of the 6th ACM conference on Electronic commerce (EC)*, pages 1–8, Vancouver, Canada, June 2005.
- [2] A. Altman and M. Tennenholtz. An axiomatic approach to personalized ranking systems. *Journal of the ACM*, 57(4):1–35, 2010.
- [3] R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: an axiomatic approach. In *Proceedings of the 17th international conference on World Wide Web (WWW)*, pages 199–208, Beijing, China, April 2008.
- [4] K. J. Arrow. *Social Choice and Individual Values*. Yale University Press, 1951.
- [5] E. Durfee, V. R. Lesser, and D. D. Corkill. Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers*, 36(2):1275–1291, 1987.
- [6] Y. Emek, R. Karidi, M. Tennenholtz, and A. Zohar. Mechanisms for multi-level marketing. In *Proceedings of the 12th ACM conference on Electronic commerce (EC)*, pages 209–218, San Jose, California, June 2011.
- [7] M. S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man and Cybernetics*, 11(1):70–80, 1981.
- [8] J. R. Galbraith. *Designing Complex Organizations*. Addison-Wesley, 1973.
- [9] L. Gao and J. Rexford. Stable internet routing without global coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, December 2001.
- [10] T. G. Griffin and J. L. Sobrinho. Metarouting. In *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 1–12, Philadelphia, Pennsylvania, August 2005.
- [11] S. Herzog, S. Shenker, and D. Estrin. Sharing the “cost” of multicast trees: An axiomatic analysis. *IEEE/ACM Transactions on Networking*, 5(6):847–860, December 1997.
- [12] M. Karsten, S. Keshav, S. Prasad, and M. Beg. An axiomatic basis for communication. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pages 217–228, Kyoto, Japan, August 2007.
- [13] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems (MSWiM)*, pages 78–82, Venice, Italy, October 2004.
- [14] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *Proceedings of the 40th annual ACM symposium on Theory of computing (STOC)*, pages 57–66, Victoria, Canada, May 2008.
- [15] T. W. Malone. Informational efficiency in networks and hierarchies. Sloan School of Management Working Paper 1849, MIT, 1986.
- [16] J. G. March and H. A. Simon. *Organizations*. John Wiley and Sons, 1958.
- [17] Y. L. Sun, W. Yu, Z. Han, and K. J. Liu. Information theoretic framework of trust modeling and evaluation for ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(2):305–317, September 2006.
- [18] T. S. Verma and J. Pearl. Causal networks: Semantics and expressiveness. In *Proceedings of the 4th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 69–78, 1988.