

COMPRESSION AND EXPANSION IN GRAPHS USING OVERLAYS

BILAL KHAN * AND KIRAN R. BHUTANI †

Abstract. In this paper, we describe how to construct physical computer network topologies which can support the establishment of overlays that reduce or increase the distances between nodes. Reducing pairwise distances (i.e., *compression*) implies that the overlay enjoys significantly lower inter-node latencies compared to the ambient physical network; such an overlay can be used to implement a “high-performance mode” for disaster situations in which network responsiveness is of critical importance. On the other hand, increasing pairwise distances (i.e., *expansion*) implies that the overlay exhibits significantly higher inter-node latencies compared to the ambient physical network; such an overlay can be used to implement a brief “dilated state” in networks that have been infected by a malicious worm, where slowing down the infection spread allows greater time for antidote generation. We show that it is possible to design physical networks which support overlays whose logical link bandwidth is *equal* to the physical link bandwidth while providing arbitrarily high compression or expansion. We also show that it is possible to “grow” such networks over time in a scalable way, that is to say, it is possible to retain the compression/expansion properties while augmenting the network with new nodes, by making relatively small adjustments to the physical and overlay network structure.

Key words. graphs embeddings, diameter, expansion, compression, overlays

1. Introduction. In the last few years we have witnessed a renewed interest in applications of peer-to-peer overlay networks. Within an overlay network, each logical link is implemented as a multi-hop connection over a path in an underlying physical network. The overlay’s configuration and protocols ensure that data can travel along a logical link without undergoing expensive processing at intermediate physical nodes. Overlays thus enable network designers to superimpose alternate topologies (i.e., “virtual networks”) whose properties diverge considerably from those of the ambient physical network. The malleability thus obtained is very valuable, since modifications to the physical network topology are both time consuming and expensive, while alterations to the overlay topology are, by comparison, implementable inexpensively and quickly. The mathematical theory of overlays has been developed and applied in many areas, including the simulation of one parallel architecture on another [4, 8, 10, 11], the design of active networks [7], the layout of virtual paths in circuit switched networks [1, 6, 9], scalable distributed content-addressable storage [13, 18], and construction of efficient multicast trees [3]. Today, the design and implementation of network overlays is made commonplace by readily available tools such as Xbone, 6Bone, and Dynabone [14, 17].

*DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, CITY UNIVERSITY OF NEW YORK AND JOHN JAY COLLEGE OF CRIMINAL JUSTICE, NEW YORK, NY 10019.

†DEPARTMENT OF MATHEMATICS, THE CATHOLIC UNIVERSITY OF AMERICA, WASHINGTON D.C. 20064.

Overlays have a profound impact on the realities of physical network design. Network designers have historically sought physical networks that exhibit good performance metrics, yet struggled with the reality that these metrics are often in opposition with one another and face inherent tradeoffs. The widespread acceptance of overlay technologies makes it advantageous to seek physical networks which support a *wide range of overlays*, each with its own merits vis-a-vis performance metrics. In this paper we focus on the performance metric of inter-node latency, which is taken to be proportional to the shortest path distance between nodes in the appropriate (physical or overlay) graph.

Reducing pairwise distances (i.e., *compression*) lowers mean latencies, meaning that such networks can be made to support a “high-performance mode” that may be useful in mission-critical circumstances. The idea of using overlays to improve network performance characteristics is certainly not new. For example, Anderson et al. [2] considered the use of overlays to improve fault resilience in the RON project. There have also been attempts to use overlays to provide quality of service in wireless mobile ad-hoc networks [15] and sensor networks [12]. Researchers have considered mechanisms by which overlays may be used to meet bandwidth requirements of VOIP [5] applications. By contrast, in this paper, we investigate mechanisms for reducing pairwise and worst-case distances between nodes, for the purposes of lowering network latency.

Increasing pairwise distances (i.e., *expansion*) on the other hand, may be useful when the network is known to have been infected by a malicious agent (e.g., a worm or virus), since the use of such an overlay would slow down the rate at which the infection saturates the network’s hosts, allowing greater time for antidote generation. The idea of using overlays to mitigate network threats has been considered by previous researchers. Wang and Chien [16] explored the possibility of using dynamic overlays to defend and deflect denial of service attacks. By contrast, in this paper, we investigate mechanisms for increasing pairwise and worst-case distances between nodes, for the purpose of raising infection percolation times.

We note that increasing or reducing network diameter using overlays is trivial if the overlay’s logical links are permitted to be of lower capacity than the physical links in the ambient network. To see this, note that given a physical network on n nodes, it is always possible to use an overlay network that is a complete graph on n nodes—this reduces the diameter to its minimal value. However, if the physical network is not a complete graph, some edges of the overlay will necessarily traverse the same physical link, implying that the logical links are of *lower capacity* than the physical links in the ambient network. Similarly, it is always possible to use an overlay network that is a path of length $n - 1$; this has the effect of increasing the diameter to its maximal value. If, however, the physical network does not contain a Hamilton path, some edges of the overlay will necessarily traverse the same physical link, implying

once again that the logical links are of *lower capacity* than the physical links in the ambient network. In this paper we explore the extent to which it is possible to develop physical networks which support overlays having logical link capacity *equal* to the physical link bandwidth, and yet provide very high compression or expansion.

2. Mathematical Model. The **physical network** is represented by a graph G with vertices V and unweighted undirected edges $E \subset V \times V$. Each link in the overlay network is implemented at the physical level by a **walk** $p = (v_0, v_1, \dots, v_l)$ in G , where (v_i, v_{i+1}) is in E , for each $i = 0, \dots, l-1$. The **length** of p is denoted $|p| = l$, and the **boundary** of p is denoted $\partial p = \{v_0, v_l\}$. For each walk p , we define a corresponding **characteristic function** $\chi_p : E \rightarrow \mathbb{N}$ whose value at an edge e is defined to be the number of times e is traversed by p . For each $l \in \mathbb{Z}^+$, let $\mathcal{P}^l(G)$ be the set of all walks in G whose length does not exceed l , and then let $\mathcal{P}^*(G) = \bigcup_{l \in \mathbb{Z}^+} \mathcal{P}^l(G)$. An **overlay network** corresponding to a set of walks $Q \subset \mathcal{P}^*(G)$ is the (multi)graph $Q^\circ = (V, \{\partial p \mid p \in Q\})$.

DEFINITION 2.1. *Let G be a finite simple connected undirected graph, and let $Q \subset \mathcal{P}^*(G)$ be a set of walks in G . For each $e \in E$, the **congestion of Q on e** is defined to be the number of times e appears in Q :*

$$\Phi_Q(e) = \sum_{p \in Q} \chi_p(e).$$

The **congestion of Q on G** is taken as the largest congestion seen (over edges of G):

$$\tau_G(Q) = \max_{e \in E} \Phi_Q(e).$$

A set of walks $Q \subset \mathcal{P}^*(G)$ is called **sparse** if $\tau_G(Q) = 1$, and is called **spanning** if Q° is a single connected component on V . Let $\mathcal{S}(G) \subset 2^{\mathcal{P}^*(G)}$ be the collection of all sets of walks in G that are both sparse and spanning.

We remark that if $G = (V, E)$ is a physical network in which every (physical) link has capacity c , and Q is a set of walks in G , then the overlay graph Q° also enjoys the property that its (logical) links each have capacity no less than $c/\tau_G(Q)$. Hence if $Q \in \mathcal{S}(G)$ is a sparse spanning set of walks, then logical link capacities in Q° coincide with physical link capacities in G . Note that $\mathcal{S}(G) \neq \emptyset$ since $E \in \mathcal{S}(G)$. Toward this end, let $d_G(u, v)$ (resp. $d_{Q^\circ}(u, v)$) denote the geodesic distance between u and v in G (resp. Q°).

DEFINITION 2.2. *Let G be a finite simple connected undirected graph, and let $Q \subset \mathcal{P}^*(G)$ be a set of walks in G . The **pointwise expansion** (resp. **compression**)*

of Q on G are defined (respectively) as

$$\begin{aligned}\bar{\kappa}_0(G, Q) &= \max_{\substack{u, v \in V \\ u \neq v}} \frac{d_{Q^\circ}(u, v)}{d_G(u, v)}, \\ \underline{\kappa}_0(G, Q) &= \min_{\substack{u, v \in V \\ u \neq v}} \frac{d_{Q^\circ}(u, v)}{d_G(u, v)} = \max_{\substack{u, v \in V \\ u \neq v}} \frac{d_G(u, v)}{d_{Q^\circ}(u, v)}.\end{aligned}$$

In contrast with pointwise compression which is defined by considering distances between corresponding vertices in G and Q° , the next definition considers the aggregate property of a one-dimensional measure, namely diameter. Towards this end, let $\text{Diam}(G)$ (resp. $\text{Diam}(Q^\circ)$) denote the diameter of G (resp. Q°).

DEFINITION 2.3. *Let G be a finite simple undirected connected graph, and consider $Q \subset \mathcal{P}^*(G)$ a set of walks in G . Then the **diametric expansion** (and **compression**) of Q on G are both defined as*

$$\bar{\kappa}_1(G, Q) = \frac{\text{Diam}(Q^\circ)}{\text{Diam}(G)} = \underline{\kappa}_1(G, Q).$$

Building on Definitions 2.2 and 2.3 we introduce the following definitions:

DEFINITION 2.4. *For $i = 0, 1$ let:*

$$\begin{aligned}\bar{\kappa}_i(G) &= \max_{Q \in \mathcal{S}(G)} \bar{\kappa}_i(G, Q), \\ \underline{\kappa}_i(G) &= \min_{Q \in \mathcal{S}(G)} \underline{\kappa}_i(G, Q).\end{aligned}$$

We now extend our definitions of pointwise/diametric compression and expansion to sequences of graphs.

DEFINITION 2.5. *Let $\mathcal{G} = (G_j \mid j \in \mathbb{N})$ be a sequence of graphs with $\mathcal{Q} = (Q_j \mid j \in \mathbb{N}, Q_j \in \mathcal{S}(G_j))$ a sequence of sparse spanning sets of walks. We define, for $i \in \{0, 1\}$:*

$$\begin{aligned}\bar{\kappa}_i(\mathcal{G}, \mathcal{Q}) &= \lim_{j \in \mathbb{N}} \bar{\kappa}_i(G_j, Q_j), \\ \underline{\kappa}_i(\mathcal{G}, \mathcal{Q}) &= \lim_{j \in \mathbb{N}} \underline{\kappa}_i(G_j, Q_j),\end{aligned}$$

where the limits are taken over the affinely extended reals $\mathbb{R} \cup \{\pm\infty\}$. A pair consisting of a sequence of graphs \mathcal{G} and a sequence of sets of paths \mathcal{Q} is designated

- **pointwise ultracompressed** if $\underline{\kappa}_0(\mathcal{G}, \mathcal{Q}) = 0$;
- **diametrically ultracompressed** if $\underline{\kappa}_1(\mathcal{G}, \mathcal{Q}) = 0$;
- **pointwise ultraexpanded** if $\bar{\kappa}_0(\mathcal{G}, \mathcal{Q}) = \infty$; and
- **diametrically ultraexpanded** if $\bar{\kappa}_1(\mathcal{G}, \mathcal{Q}) = \infty$.

A sequence of graphs $\mathcal{G} = (G_j \mid j \in \mathbb{N})$ is called **pointwise ultracompressible** if for some $\mathcal{Q} = (Q_j \mid j \in \mathbb{N}, Q_j \in \mathcal{S}(G_j))$ the pair $(\mathcal{G}, \mathcal{Q})$ is ultracompressed. We define the terms **diametrically ultracompressible**, **pointwise ultraexpandible** and **diametrically ultraexpandible** analogously.

3. Compressibility and Expandibility. The next theorem shows a connection between sparse spanning sets of paths and the phenomena of pointwise compression and expansion. Specifically, it demonstrates that a sparse spanning set of walks which exhibits non-trivial pointwise compression ($\underline{\kappa}_0(G, Q) < 1$) necessarily exhibits non-trivial pointwise expansion ($\overline{\kappa}_0(G, Q) > 1$).

THEOREM 3.1 (No Free Lunch). *Let $G = (V, E)$ be a non-trivial ($|V| > 1$) finite simple connected undirected graph, and let Q be a sparse spanning set of walks for which $\underline{\kappa}_0(G, Q) < 1$. Then $\overline{\kappa}_0(G, Q) > 1$.*

Proof. Suppose, to the contrary, that $\overline{\kappa}_0(G, Q) \leq 1$. Then for any u', v' in V ,

$$\frac{d_{Q^\circ}(u', v')}{d_G(u', v')} \leq \max_{\substack{u, v \in V \\ u \neq v}} \frac{d_{Q^\circ}(u, v)}{d_G(u, v)} = \overline{\kappa}_0(G, Q) \leq 1.$$

It follows that $d_{Q^\circ}(u', v') \leq d_G(u', v')$ for all u', v' in V . Since G is non-trivial and connected, let us take u' and v' to be any distinct adjacent vertices in G ; then $d_{Q^\circ}(u', v') \leq 1$. Now the simplicity of G mandates that every path in Q must have distinct endpoints, and hence $d_{Q^\circ}(u', v') = 1$. We have shown that every edge (u', v') in G appears as a length one path in Q , and hence $Q \supseteq E$. Since by hypothesis Q is sparse, it must be that $|Q| = |E|$, and thus $Q = E$. This, however, implies that $\underline{\kappa}_0(G, Q) = 1$, a contradiction. \square

The next two theorems show that it is possible to construct sequences of graphs that are pointwise ultracompressible and diametrically ultracompressible, respectively.

THEOREM 3.2. *There is a pointwise ultracompressible sequence of graphs.*

Proof. We define, for any $m \in \mathbb{N}$, the bridge graph B_m as shown in Figure 3.1. In each graph B_m we select paths Q_m to form the logical edges of the overlay network, as shown in Figure 3.2.

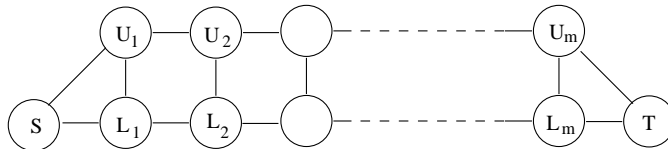
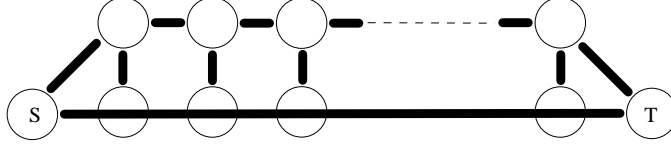


FIG. 3.1. The physical network B_m .

Consider the sequence of graphs $\mathcal{G}_1 = (B_{2^j} \mid j \in \mathbb{N})$, and take the sequence of sets of paths $\mathcal{Q}_1 = (Q_{2^j} \mid j \in \mathbb{N})$. To see that \mathcal{G}_1 is pointwise ultracompressible, note that

FIG. 3.2. The overlay network Q_m on B_m .

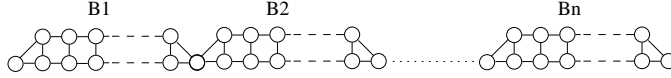
$d_{Q_{2^j}^\circ}(S, T) = 1$ while $d_{B_{2^j}}(S, T) = 2^j + 1$, and hence

$$\underline{\kappa}_0(\mathcal{G}_1, \mathcal{Q}_1) = \lim_{j \in \mathbb{N}} \underline{\kappa}_0(B_{2^j}, Q_{2^j}) \leq \lim_{j \in \mathbb{N}} \frac{1}{2^j + 1} = 0.$$

It follows by Definition 2.5 that $(\mathcal{G}_1, \mathcal{Q}_1)$ is pointwise ultracompressed, and hence \mathcal{G}_1 is pointwise ultracompressible. \square

THEOREM 3.3. *There is a diametrically ultracompressible sequence of graphs.*

Proof. Define the chain of bridges $C_{n,m}$ as seen in Figure 3.3, consisting of a sequence of n copies of B_m . Since $C_{n,m}$ consists of n copies of B_m we extend the

FIG. 3.3. The chain of bridges graph, $C_{n,m}$.

vertex naming conventions of B_m to $C_{n,m}$ in the following natural way: We denote S^k , T^k , U_i^k and L_i^k as the vertices S , T , U_i and L_i in the k th copy of B_m which appears in $C_{n,m}$ (indexing copies sequentially as k varies from 1 to n). Note that $S^1 = S$, $S^{k+1} = T^k$ and $T^n = T$. We use an overlay network $Q_{n,m}$ which consists of n copies of Q_m (introduced in the proof of Theorem 3.2 and depicted in Figure 3.2).

Consider the sequence of graphs $\mathcal{G}_2 = (C_{2^j, 6 \cdot 2^j - 1} \mid j \in \mathbb{N})$, and take the sequence of sets of paths $\mathcal{Q}_2 = (Q_{2^j, 6 \cdot 2^j - 1} \mid j \in \mathbb{N})$. Then

$$\text{Diam}(C_{2^j, 6 \cdot 2^j - 1}) = 2^j(6 \cdot 2^j) = 6 \cdot 2^{2j}$$

for all $j \in \mathbb{N}$. On the other hand,

$$\text{Diam}(Q_{2^j, 6 \cdot 2^j - 1}^\circ) = 7 \cdot 2^j$$

whenever $j > 0$ (and equals 4 for the case when $j = 0$). Thus

$$\underline{\kappa}_1(\mathcal{G}_2, \mathcal{Q}_2) = \lim_{j \in \mathbb{N}} \underline{\kappa}_1(C_{2^j, 6 \cdot 2^j - 1}, Q_{2^j, 6 \cdot 2^j - 1}) = \lim_{j \in \mathbb{N}} \frac{7 \cdot 2^j}{6 \cdot 2^{2j}} = 0.$$

It follows by Definition 2.5 that $(\mathcal{G}_2, \mathcal{Q}_2)$ is diametrically ultracompressed, and hence \mathcal{G}_2 is diametrically ultracompressible. \square

The next two theorems show that it is possible to construct sequences of graphs that are diametrically ultraexpandible and pointwise ultraexpandible, respectively.

THEOREM 3.4. *There is a diametrically ultraexpandible sequence of graphs.*

Proof. Consider the sequence of graphs $\mathcal{G}_3 = (G_j \mid j \in \mathbb{N})$ in which the G_j are defined inductively as follows: The initial element $G_0 = (O, \emptyset)$ is a trivial graph, consisting of a single isolated vertex O . Each subsequent graph G_j is obtained from G_{j-1} by taking a disjoint copy of the complete graph on 2^j vertices in which one vertex is distinguished as u_j , and adding the edge (O, u_j) , thereby connecting a copy of K_{2^j} to G_{j-1} through O using this new edge. The construction process is depicted in Figure 3.4.

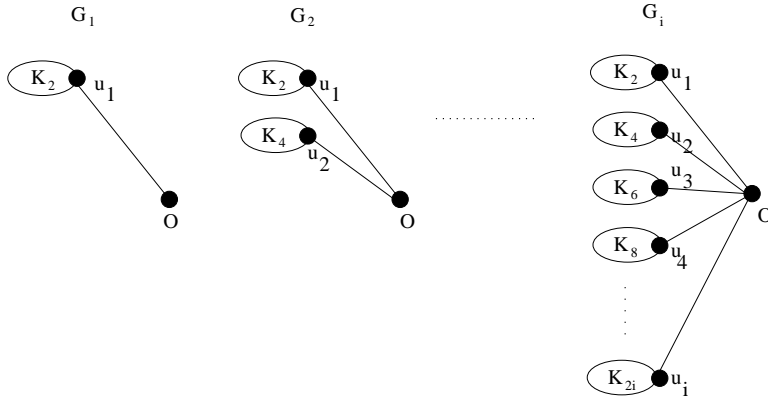


FIG. 3.4. An ultraexpandible sequence of graphs G_i .

We take the sequence of sets of paths $\mathcal{Q}_3 = (Q_j \mid j \in \mathbb{N})$ where Q_j is defined inductively as follows. The initial element $Q_0 = \emptyset$. Each subsequent set of paths Q_j is obtained by augmenting Q_{j-1} by adding 2^j paths, each of length 1: the $2^j - 1$ edges of any Hamilton path in K_{2^j} starting at vertex u_j , as well as the path consisting of just the edge (O, u_j) . For reference, we denote the endpoint vertices of the chosen Hamilton path in K_{2^j} as u_j and w_j .

Note that the diameter of G_j is no more than 4. On the other hand, the diameter of Q_j° (i.e., the longest shortest path in Q_j°) consists of the Hamilton path in K_{2^j} followed by edges (u_j, O) and (O, u_{j-1}) , followed by the Hamilton path in $K_{2^{j-1}}$. Thus $\text{Diam}(Q_j^\circ) = 4j - 2$. Since

$$\bar{\kappa}_1(\mathcal{G}_3, \mathcal{Q}_3) = \lim_{j \in \mathbb{N}} \bar{\kappa}_1(G_j, Q_j) = \lim_{j \in \mathbb{N}} \frac{4j - 2}{4} = \infty,$$

it follows by Definition 2.5 that $(\mathcal{G}_3, \mathcal{Q}_3)$ is diametrically ultraexpanded and hence that \mathcal{G}_3 is diametrically ultraexpandible. \square

THEOREM 3.5. *There is a pointwise ultraexpandible sequence of graphs.*

Proof.

Consider the sequence of graphs \mathcal{G}_3 and the sequence of sets of paths \mathcal{Q}_3 defined in the proof of Theorem 3.4. Note that

$$\begin{aligned} \bar{\kappa}_0(\mathcal{G}_3, \mathcal{Q}_3) &= \lim_{j \in \mathbb{N}} \bar{\kappa}_0(G_j, Q_j) = \lim_{j \in \mathbb{N}} \max_{\substack{u, v \in V[G_j] \\ u \neq v}} \frac{d_{Q_j^\circ}(u, v)}{d_{G_j}(u, v)} \geq \\ &= \lim_{j > 0} \frac{d_{Q_j^\circ}(w_j, w_{j-1})}{d_{G_j}(w_j, w_{j-1})} = \lim_{j > 0} \frac{4j - 2}{4} = \infty. \end{aligned}$$

It follows by Definition 2.5 that $(\mathcal{G}_3, \mathcal{Q}_3)$ is pointwise ultraexpanded and hence that \mathcal{G}_3 is pointwise ultraexpandible. \square

4. Scalable Growth Sequences. Theorems 3.2-3.5 of the previous section show that it is possible to construct graphs which are expandible and compressible to arbitrary extents, in both the pointwise and diametric sense. They do not, however, necessarily provide a way by which to efficiently “grow” the network over time so that it evolves towards ever greater expandibility or compressibility. The next three definitions, which culminate in the notion of *scalable growth*, make this objective precise. We begin by exploring the implications of “evolving” the network from a state where the physical topology is given by the graph $G = (V, E)$ and the logical topology by a sparse spanning set of paths $Q \subseteq \mathcal{S}^*(G)$, to a new state where the new physical topology is given by the graph $G' = (V', E')$ and the new logical topology by a sparse spanning set of paths $Q' \subseteq \mathcal{S}^*(G')$.

DEFINITION 4.1. *Given two graphs $G = (V, E)$ and $G' = (V', E')$, and an injective map $f : V \hookrightarrow V'$, we define*

$$f(E) = \{(f(u), f(v)) \mid (u, v) \in E \text{ and } (f(u), f(v)) \in E'\}.$$

Informally, $f(E)$ are the edges of G which survive the evolution f from G to G' . Thus $|(f(V) \times V') \cap E'| - |f(E)|$ is the number of new edges which are incident to at least one vertex of G , while $|E| - |f(E)|$ is the number of existing edges which were deleted from $f(V)$ during the transition.

DEFINITION 4.2. *Given graphs $G = (V, E)$, $G' = (V', E')$, an injective map $f : V \hookrightarrow V'$, and a walk $q = (v_1, v_2, \dots, v_k) \in \mathcal{P}^*(G)$ we say that q is **projected** by f if $(f(v_j), f(v_{j+1})) \in f(E)$ for all $j = 1, \dots, k - 1$. In this case, we define*

$$f(q) = (f(v_1), f(v_2), \dots, f(v_k)).$$

If $Q \subseteq \mathcal{P}^*(G)$ is a set of walks in G , we define

$$f(Q) = \{f(q) \mid q \in Q \text{ and } q \text{ is projected by } f\} \subseteq \mathcal{P}^*(G').$$

Given a walk $q' = (w'_1, w'_2, \dots, w'_m) \in \mathcal{P}^*(G')$ we say that q' **collides** with f if $w'_j \in f(V)$ for at least one value of $j \in \{1, \dots, m\}$. If $Q' \subseteq \mathcal{P}^*(G')$ is a set of walks in G' , we define

$$X_f(Q') = \{q' \in Q' \setminus f(Q) \mid q' \text{ collides with } f\} \subseteq Q'.$$

Informally, $f(Q)$ are the paths of Q which survive the evolution from G to G' , while $X_f(Q')$ are new paths in Q' which involve at least one existing vertex of $f(V)$. If Q is a set of paths in G , and Q' is a set of paths in G' , then $|Q| - |f(Q)|$ is the number of existing paths which were deleted from Q in the transition, while $|X_f(Q')|$ is the number of new paths that involved at least one existing vertex of G .

Summarizing the discussions following the previous two definitions, the evolution f of a graph G (with overlay paths Q) to a graph G' (with overlay paths Q') requires a disruption of the existing physical network structure which we quantify with the expression:

$$|(f(V) \times V') \cap E'| - |f(E)| + |E| - |f(E)|,$$

and a disruption of the existing logical structure which we quantify as:

$$|X_f(Q')| + |Q| - |f(Q)|.$$

DEFINITION 4.3. Let $\mathcal{G} = (G_j = (V_j, E_j) \mid j \in \mathbb{N})$ be a sequence of graphs and take $\mathcal{Q} = (Q_j \mid j \in \mathbb{N}, Q_j \in \mathcal{S}(G_j))$ to be a sequence of sparse spanning sets of walks in the corresponding graphs of \mathcal{G} . The pair $(\mathcal{G}, \mathcal{Q})$ is said to be a **scalable growth sequence** if the following three conditions hold:

1. The sequence of graphs grows monotonically and unboundedly:

$$j > j' \Rightarrow |V_j| > |V_{j'}|. \tag{4.1}$$

2. Relatively few physical edges are added/removed from the existing network at each stage; i.e., there is a sequence of injections of vertices $f_j : V_j \hookrightarrow V_{j+1}$ (for each $j \in \mathbb{N}$) such that

$$\lim_{j \rightarrow \infty} \frac{|(f(V_j) \times V_{j+1}) \cap E_{j+1}| + |E_j| - 2|f(E_j)|}{|E_j|} = 0. \tag{4.2}$$

3. Relatively few logical edges are added/removed from the existing network overlay at each stage:

$$\lim_{j \rightarrow \infty} \frac{|X_{f_j}(Q_{j+1})| + |Q_j| - |f_j(Q_j)|}{|Q_j|} = 0. \tag{4.3}$$

In this manner, constraints (4.1)-(4.3) express the requirement that over time (index j) the network can be (1) made to grow, in a manner that (2) requires relatively few changes to be made to both the physical network (3) and logical network at each stage. The next four theorems extend Theorems 3.2-3.5 by showing that there exist a pair $(\mathcal{G}, \mathcal{Q})$ of sequences of graphs and sets of paths which are pointwise/diametrically ultraexpanded/ultracompressed and which, additionally, exhibit scalable growth.

THEOREM 4.4. *There is a pointwise ultracompressed scalable growth sequence.*

Proof. It suffices to show that the sequence of graphs \mathcal{G}_1 and the sequence of sets of paths \mathcal{Q}_1 (defined in the proof of Theorem 3.2) exhibit scalable growth. Towards this end we define $f_j : V[G_j] \hookrightarrow V[G_{j+1}]$ as

$$f_j : \begin{cases} S & \mapsto S \\ U_i & \mapsto U_i \quad (i = 1, \dots, 2^j) \\ L_i & \mapsto L_i \quad (i = 1, \dots, 2^j) \\ T & \mapsto L_{2^{j+1}} \end{cases}$$

where the function f_j is described as a map between the labels of vertices of G_j and the labels of vertices of G_{j+1} . It is easy to see that f_j is injective. Since $|V_{j+1}| > |V_j|$, constraint (4.1) holds. Since $|(f(V_j) \times V_{j+1}) \cap E_{j+1}| = 3 \cdot 2^j + 4$ and $|E_j| = 3 \cdot 2^j + 2$ and $|f_j(E_j)| = 3 \cdot 2^j + 1$, it follows that

$$\frac{|(f(V_j) \times V_{j+1}) \cap E_{j+1}| + |E_j| - 2|f_j(E_j)|}{|E_j|} = \frac{(3 \cdot 2^j + 4) + (3 \cdot 2^j + 2) - 2(3 \cdot 2^j + 1)}{3 \cdot 2^j + 2} = \frac{4}{3 \cdot 2^j + 2},$$

which is an expression which tends to 0 in the limit (as j tends to infinity) thereby showing that constraint (4.2) holds. Since $|X_{f_j}(Q_{j+1})| = 3$, $|Q_j| = 2(2^j + 1)$ and $|f_j(Q_j)| = 2^{j+1}$, it follows that

$$\frac{|X_{f_j}(Q_{j+1})| + |Q_j| - |f_j(Q_j)|}{|Q_j|} = \frac{3 + (2^{j+1} + 2) - 2^{j+1}}{2(2^j + 1)} = \frac{5}{2(2^j + 1)},$$

which is an expression which tends to 0 in the limit (as j tends to infinity) thereby showing that constraint (4.3) holds. We have shown that $(\mathcal{G}_1, \mathcal{Q}_1)$ is a pointwise ultracompressed sequence which exhibits scalable growth. \square

THEOREM 4.5. *There is a diametrically ultracompressed scalable growth sequence.*

Proof. It suffices to show that the sequence of graphs \mathcal{G}_2 and the sequence of sets of paths \mathcal{Q}_2 (defined in the proof of Theorem 3.3) exhibit scalable growth. Towards

this end we define $f_j : V[G_j] \hookrightarrow V[G_{j+1}]$ as

$$f_j : \begin{cases} S^k \mapsto S^{(k+1)/2} & (1 \leq k \leq 2^j, k \text{ odd.}) \\ S^k \mapsto L_{6 \cdot 2^j}^{k/2} & (1 \leq k \leq 2^j, k \text{ even.}) \\ U_i^k \mapsto U_i^{(k+1)/2} & (1 \leq k \leq 2^j, k \text{ odd, } 1 \leq i \leq 6 \cdot 2^j - 1) \\ U_i^k \mapsto U_{6 \cdot 2^j + i}^{k/2} & (1 \leq k \leq 2^j, k \text{ even, } 1 \leq i \leq 6 \cdot 2^j - 1) \\ L_i^k \mapsto L_i^{(k+1)/2} & (1 \leq k \leq 2^j, k \text{ odd, } 1 \leq i \leq 6 \cdot 2^j - 1) \\ L_i^k \mapsto L_{6 \cdot 2^j + i}^{k/2} & (1 \leq k \leq 2^j, k \text{ even, } 1 \leq i \leq 6 \cdot 2^j - 1) \\ T^{2^j} \mapsto T^{2^j - 1} & (k = 1, \dots, m) \end{cases}$$

where the function f_j is described as a map between the labels of vertices of G_j and the labels of vertices of G_{j+1} . It is easy to see that f_j is injective. Since $|V_{j+1}| > |V_j|$, constraint (4.1) holds. Since $|(f(V_j) \times V_{j+1}) \cap E_{j+1}| - |f_j(E_j)| = 3 \cdot 2^{j-1} + 2$ and $|E_j| - |f_j(E_j)| = 2^j$,

$$\frac{|(f(V_j) \times V_{j+1}) \cap E_{j+1}| + |E_j| - 2|f_j(E_j)|}{|E_j|} = \frac{5 \cdot 2^{j-1} + 2}{2^j(9 \cdot 2^{j+1} - 1)},$$

an expression which tends to 0 in the limit (as j tends to infinity) thereby showing that constraint (4.2) holds. Since $|X_{f_j}(Q_{j+1})| = 2^{j+1} + 2$, $|Q_j| = 3 \cdot 2^{2j+2}$ and $|f_j(Q_j)| = 3 \cdot 2^{2j+2} - 2^{j+1}$, it follows that

$$\frac{|X_{f_j}(Q_{j+1})| + |Q_j| - |f_j(Q_j)|}{|Q_j|} = \frac{(2^{j+1} + 2) + 3 \cdot 2^{2j+2} - (3 \cdot 2^{2j+2} - 2^{j+1})}{3 \cdot 2^{2j+2}} = \frac{2^{j+2} + 2}{3 \cdot 2^{2j+2}},$$

an expression which tends to 0 in the limit (as j tends to infinity) thereby showing that constraint (4.3) holds. We have shown that $(\mathcal{G}_2, \mathcal{Q}_2)$ is a diametrically ultracompressed sequence which exhibits scalable growth. \square

THEOREM 4.6. *There is a pointwise ultraexpanded scalable growth sequence.*

Proof. It suffices to show that the sequence of graphs \mathcal{G}_3 and the sequence of sets of paths \mathcal{Q}_3 (defined in the proof of Theorem 3.5) exhibit scalable growth. The construction which defines the manner by which G_{j+1} is derived from G_j implicitly defines a graph embedding $f_j : G_j \hookrightarrow G_{j+1}$. First, since $|V_{j+1}| > |V_j|$, constraint (4.1) holds. Moreover, since

$$\begin{aligned} |E_j| - |f_j(E_j)| &= 0 \\ |(f(V_j) \times V_{j+1}) \cap E_{j+1}| - |f_j(E_j)| &= 1, \end{aligned}$$

and since $|E_j|$ goes to infinity as j goes to infinity,

$$\lim_{j \rightarrow \infty} \frac{|(f(V_j) \times V_{j+1}) \cap E_{j+1}| + |E_j| - 2|f(E_j)|}{|E_j|} = 0,$$

hence constraint (4.2) holds. It is easy to verify that $|X_{f_j}(Q_{j+1})| = 1$ and $|Q_j| - |f_j(Q_j)| = 0$, while $|Q_j| = j(j+1)$. Thus it follows that

$$\frac{|X_{f_j}(Q_{j+1})| + |Q_j| - |f_j(Q_j)|}{|Q_j|} = \frac{1}{j(j+1)},$$

an expression which tends to 0 in the limit (as j tends to infinity) thereby showing that constraint (4.3) holds. We have shown that $(\mathcal{G}_3, \mathcal{Q}_3)$ is a pointwise ultraexpanded sequence which exhibits scalable growth. \square

THEOREM 4.7. *There is a diametrically ultraexpanded scalable growth sequence.*

Proof. It suffices to show that the sequence of graphs \mathcal{G}_3 and the sequence of sets of paths \mathcal{Q}_3 (referenced in the proof of Theorem 3.4) exhibit scalable growth; this was shown as part of the proof of Theorem 4.6. \square

5. Conclusions and Future Work. We have seen that there are ways to evolve a physical network topology over time, so as to support the establishment of overlays that yield ever higher compression and expansion. Compression allows us to use the overlay to significantly lower mean latencies in the ambient physical network, and thereby implement a “high-performance mode” that could be enabled in mission-critical settings. Expansion, on the other hand, increases latencies, thereby increasing infection propagation times and helping to curtail the spread of a malicious worm or virus and providing greater time for signature and antivirus development.

We ask whether it is possible to construct a sequence of graphs $\mathcal{G} = (G_j \mid j \in \mathbb{N})$ that is both pointwise ultraexpandable and pointwise ultracompressible, i.e., which admits *two distinct sequences* of sparse spanning sets of paths

$$\begin{aligned} \mathcal{Q}^* &= (Q_j \mid j \in \mathbb{N}, Q_j \in \mathcal{S}(G_j)) \\ \mathcal{Q}_* &= (R_j \mid j \in \mathbb{N}, R_j \in \mathcal{S}(G_j)) \end{aligned}$$

such that $(\mathcal{G}, \mathcal{Q}_*)$ is pointwise ultracompressed, while $(\mathcal{G}, \mathcal{Q}^*)$ is a pointwise ultraexpanded. If in addition $(\mathcal{G}, \mathcal{Q}_*)$ and $(\mathcal{G}, \mathcal{Q}^*)$ could be taken to be scalable growth sequences, then \mathcal{G} would provide a recipe by which to “grow” a network in a manner that admits both radical increases or decreases in diameter—merely by instrumenting the appropriate overlay \mathcal{Q}^* or \mathcal{Q}_* , as the specific circumstances mandate. The authors hope to report on these types of questions, and their diametric analogues, in the future.

The authors would like to thank Daniel Tunkelang, Amitabha Bagchi, Joel Sandin,

Delaram Kahroabei and Amitabh Chaudhury for many spirited discussions about this problem during the early stages of our work.

REFERENCES

- [1] S. Ahn, R. P. Tsang, S.-R. Tong, and D. H.-C. Du, Virtual path layout design on ATM networks, Proc of IEEE INFOCOM, Toronto, Ontario, Canada, 1994, pp. 192–200.
- [2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, Resilient overlay networks, Proc of SOSP, Chateau Lake Louise, Banff, Canada, 2001, pp. 131–145.
- [3] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, and S. Khuller, Construction of an efficient overlay multicast infrastructure for real-time applications, Proc of IEEE INFOCOM, San Francisco, CA, 2003, pp. 1521–1531.
- [4] S. L. Bezrukov, B. Monien, W. Unger, and G. Wechsung, Embedding ladders and caterpillars into the hypercube, Discrete Applied Mathematics 83 (1998), 21–29.
- [5] C. Chuah, A scalable framework for IP-network resource provisioning through aggregation and hierarchical control, Ph.D. thesis, University of California, Berkeley, 2001.
- [6] T. Eilam, M. Flammini, and S. Zaks, A complete characterization of the path layout construction problem for ATM networks with given hop count and load, Proc of International Colloquium on Automata, Languages and Programming (ICALP), Bologna, Italy, 1997, pp. 527–537.
- [7] V. Ferreria, A. Rudenko, K. Eustice, R. Guy, V. Ramakrishna, and P. Reiher, Panda: Middleware to provide the benefits of active networks to legacy applications, Proc of DARPA Active Networks Conference & Expo, Los Alamitos, CA, 2002. pp. 319.
- [8] L. S. Heath, A. L. Rosenberg, and B. T. Smith, The physical mapping problem for parallel architectures, J ACM 35 (1988), 603–634.
- [9] J. D. Kock and A. Krzesinski, Computing an optimal virtual path connection network by simulated annealing, Proc of the 1st Annual South African Telecommunications Network and Applications Conference (SATNAC), 1998, pp. 611–617.
- [10] F. T. Leighton, Introduction to parallel algorithms and architectures: Arrays, trees & hypercubes, Morgan Kaufmann, San Francisco, CA, 1991.
- [11] B. Monien and H. Sudborough, Embedding one interconnection network in another, Computational Graph Theory, Computing Supplement 7, Springer-Verlag, 1990, 257–282.
- [12] J. Newsome and D. Song, Gem: Graph embedding for routing and data-centric storage in sensor networks without geographic information, Proc of 1st International Conference on Embedded Networked Sensor Systems, ACM Press, Los Angeles, CA, 2003, pp. 76–88.
- [13] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan, Chord: A scalable peer-to-peer lookup service for Internet applications, Proc of ACM SIGCOMM Conference, San Diego, CA, 2001, pp. 149–160.
- [14] J. Touch, Dynamic Internet overlay deployment and management using the X-Bone, Computer Networks 36 (2001), 117–135.
- [15] S. Waharte, J. Xiao, and R. Boutaba, Overlay wireless sensor networks for application-adaptive scheduling in WLAN, Proc of HSNMC, Toulouse, France, 2004, pp. 676–684.
- [16] J. Wang and A. Chien, Using overlay networks to resist denial-of-service attacks, Technical Report, CSE Department, University of California, San Diego, 2003.
- [17] Y. Wang and J. Touch, Application deployment in virtual networks using the X-Bone, Proc of DARPA Active Networks Conference & Expo, Los Alamitos, CA, 2002, pp. 484–493.
- [18] B. Y. Zhao, J. D. Kubiatowicz, and A. D. Joseph, Tapestry: An infrastructure for fault-tolerant wide-area location and routing, Technical Report UCB/CSD-01-1141, University of California, Berkeley, 2001.