

Euclidean versus Hyperbolic Congestion in Idealized versus Experimental Networks

Edmond Jonckheere, Mingji Lou, Francis Bonahon, and Yuliy Baryshnikov

Abstract. This paper proposes a mathematical justification of the phenomenon of extreme congestion at a very limited number of nodes in very large networks. It is argued that this phenomenon occurs as a combination of the negative curvature property of the network together with minimum-length routing. More specifically, it is shown that in a large n -dimensional hyperbolic ball B of radius R viewed as a roughly similar model of a Gromov hyperbolic network, the proportion of traffic paths transiting through a small ball near the center is $\Theta(1)$, whereas in a Euclidean ball, the same proportion scales as $\Theta(1/R^{n-1})$. This discrepancy persists for the traffic load, which at the center of the hyperbolic ball scales as $\text{volume}^2(B)$, whereas the same traffic load scales as $\text{volume}^{1+1/n}(B)$ in the Euclidean ball. This provides a theoretical justification of the experimental exponent discrepancy observed by Narayan and Saniee between traffic loads in Gromov-hyperbolic networks from the Rocketfuel database and synthetic Euclidean lattice networks. It is further conjectured that for networks that do not enjoy the obvious symmetry of hyperbolic and Euclidean balls, the point of maximum traffic is near the center of mass of the network.

1. Introduction

One of the most important challenges in large and wide area networks is to overcome the traffic congestion problem. The queuing feature at routers can create a logic bottleneck between two users. Correspondingly, insufficient bandwidth on the physical links between routers is a contributor to congestion. The current congestion control technologies in communication networks are based on the feedback from the congested node to the source to slow down the packet flow rate, such as bidirectional congestion control and random early detection (RED). However, these technologies can be applied only after the congestion has happened to some degree, and it is based only on the *local* point of view of some queue overflow along the source-to-target path.

From the large-scale point of view, it has been experimentally observed that on the Internet and other networks, traffic seems to concentrate quite heavily on some very small subsets. The major result of this paper is to show that the deeper reason behind this congestion *in the large scale* is the combination of the least-cost routing together with the negative curvature of the network.

Roughly speaking, a network is negatively curved if its graph can be approximated by a negatively curved Riemannian manifold. In practice, however, it has become customary to check the Gromov *thin triangle condition* (TTC), meaning that the least-cost paths between three vertices “arch inside” the triangle, giving it a “thin” external appearance. The connection between negatively curved surfaces and the TTC can be understood on the basis that a triangle drawn on a negatively curved surface has the sum of its internal angles less than π , giving it a “thin” appearance. The formalization of this equivalence is, however, much harder and is known as the Bonk–Schramm theorem [Bonk and Schramm 00].

Over the past few years, there has been mounting evidence that many (wired and wireless) communication networks are negatively curved [Jonckheere et al. 08, Jonckheere and Lohsoonthorn 02, Jonckheere and Lohsoonthorn 04, Lohsoonthorn 03]. Hierarchical networks have been found to have “hidden hyperbolic structure” [Krioukov et al. 09]. In view of the above graph–manifold identification, we can certainly argue on a negatively curved n -dimensional Riemannian manifold. To simplify the exposition in this introduction, we invoke the nontrivial fact, proved in the appendix (Section 10), that congestion in very large hyperbolic ball models does not depend on the dimension. Therefore, it is legitimate to argue on a negatively curved surface, as we will do in the remainder of this introduction.

A well-known feature of least-length paths on a negatively curved surface is the fact that two geodesics starting from points within an arbitrarily small disk eventually diverge exponentially. In this paper, we somehow reverse that argument

and show that least-length paths departing at remote points of a convex subset X of the surface “converge” to a single point where the “density of geodesics” is the highest. The latter is the point of maximum congestion. We conjecture that this congestion point is near the *center of mass* of the convex subset, and prove that the maximum congestion point and the center of mass coincide for very large disks embedded in the surface.

Naturally, there might be some questions as to the relevance of this idealized analysis to real networks. In fact, some recent experimental results [Narayan and Saniee 09] have confirmed the validity of our analysis. The network congestion metric adopted in [Narayan and Saniee 09] is the betweenness centrality [Zhao et al. 05]. The continuous geometry traffic congestion metric adopted here is the average length of the geodesics in a small convex subset X . We call the latter *load measure*, denoted by $\Lambda_t(X)$, because if X is thought of as a subnetwork, then $\Lambda_t(X)$ is meant to be the number of packets in that subnetwork. Consistent with the major theme of this paper, we prove that if X is a disk of fixed radius, then $\Lambda_t(X)$ is maximal when the disk has its center at the center of mass. But more importantly, we show that in a large disk $B_R(0)$ of hyperbolic surface, $\max_{X \subseteq B_R(0), \text{area}(X)=1} \Lambda_t(X)$ scales as $(\text{area}(B_R(0)))^2$. On the other hand, it was shown in [Narayan and Saniee 09] that for networks for the Rocketfuel database [Spring et al. 04], the maximum betweenness centrality scales as N^2 , where N is the number of vertices. Since N can be interpreted as the area, this confirms the validity of our theoretical model.

An outline of the paper follows: In Section 2, we set the basic traffic metrics in graph models. In Section 3, a simple concept of negatively curved planar graphs is introduced and the general facts about congestion in negatively curved graphs are proposed as conjectures. Immediately thereafter, in Section 4, we show that our conjectures hold in selected planar graph examples. In Section 5, we propose the more general concept of Gromov-hyperbolic graphs and we invoke the Bonk-Schramm theorem to argue that traffic in Gromov-hyperbolic graphs can be analyzed on negatively curved Riemannian manifold models. In Section 6, we formulate our major results dealing with traffic in subsets of \mathbb{H}^n . Probably the most significant results are Theorems 6.2 and 6.1, proving that traffic in a small subset X of large hyperbolic and Euclidean balls $B_R(0)$ scales as $\text{vol}(B_R(0))^2$ and $\text{vol}(B_R(0))^{1.5}$, respectively, where “ $\text{vol}(\cdot)$ ” denotes the *volume measure*, the n -dimensional generalization of the elementary 3-dimensional concept, also referred to as “area” when $n = 2$. (The precise differential-geometric definition of the volume and the proofs of the scaling results are relegated to the appendix.)

Theorems 6.3 and 6.4 formulate the basic minimum inertia versus maximum traffic issues, proving the conjectures for hyperbolic balls X . Finally, in Section 7, we show that no matter how theoretical our models are, they turn

out to be surprisingly accurate at predicting asymptotic traffic distribution in realistic networks.

2. Traffic Metrics

Let $G = (V, E)$ be a (connected) graph specified by its vertex set V and its edge set E and endowed with a (symmetric) distance function $d : G \times G \rightarrow \mathbb{R}^+$. A *path* $p(s, t)$ from s to t is a continuous map $[0, l] \rightarrow G$ such that $p(s, t)(0) = s$ and $p(s, t)(l) = t$. The *weight* of an edge $e = xy$ is defined as $w(e) = d(x, y)$. The *length* of the path is defined as $\ell(p(s, t)) = \sum_{e \subseteq p(s, t)} w(e)$. A *geodesic* $[s, t]$ is a path such that $\ell([s, t]) \leq \ell(p(s, t))$ for all paths $p(s, t)$ joining s to t .

The traffic on the graph is driven by a *demand measure* $\Lambda_d : V \times V \rightarrow \mathbb{R}^+$, where the demand $\Lambda_d(s, t)$ is the traffic rate (e.g., number of packets per second) to be transmitted from the source s to the destination target t . Assume that the routing protocol sends the packets from the source s to the target t along the geodesic $[s, t]$ with probability $\pi([s, t])$. It is indeed customary as a load-balancing strategy to randomize the Dijkstra algorithm so as to distribute the traffic more evenly [Cisco 05, Lou 08]. Under this scheme, the geodesic $[s, t]$ inherits a traffic rate measure $\tau([s, t]) = \Lambda_d(s, t)\pi([s, t])$. An edge e lying on the path $[s, t]$ inherits from that path a traffic $\tau([s, t])$. Aggregating this traffic over all source–target pairs and all geodesics traversing the edge e yields the traffic rate sustained by the edge e :

$$\tau(e) = \sum_{(s,t) \in V \times V} \sum_{[s,t] \supseteq e} \tau([s,t]) = \sum_{(s,t) \in V \times V} \sum_{[s,t] \supseteq e} \Lambda_d(s,t)\pi([s,t]).$$

Observe that $\tau(e)\ell(e)$ can be interpreted as the *traffic load*, that is, the number of packets in the edge e .

This paper is essentially concerned with existence of a subnetwork that has extremely high traffic load. Formally, given a connected subgraph $X \subseteq G$ containing some edges, we define its *traffic load* to be representative of the number of packets in it:

$$\begin{aligned} \Lambda_t(X) &= \sum_{e \in X} \ell(e)\tau(e) = \sum_{s,t \in V} \left(\sum_{e \in [s,t] \cap X} w(e) \right) \Lambda_d(s,t)\pi([s,t]) \\ &= \sum_{s,t \in V} \ell([s,t] \cap X) \Lambda_d(s,t)\pi([s,t]). \end{aligned} \quad (2.1)$$

The above definition does not allow X to be reduced to a vertex, but we can identify the smallest Λ_t -measurable neighbor of a vertex. Define the star of a vertex, $\text{star}(x)$, to be the smallest subgraph of G containing x ; let $\ell(\text{star}(x))$ be

the sum of the lengths of the set of edges abutting x ; the latter vertex set is denoted by $E_x = \{xy : y \in V, xy \in E\}$. Then consider

$$\frac{\Lambda_t(\text{star}(x))}{\ell(\text{star}(x))} = \frac{\sum_{e \in E_x} \Lambda_t(e)}{\sum_{e \in E_x} \ell(e)} \leq \sum_{e \in E_x} \frac{\Lambda_t(e)}{\ell(e)} = \sum_{e \in E_x} \tau(e) =: \tau(x).$$

The inequality is a well-known fact, and equality is achieved for the number-of-hops metric, that is, $w(e) = 1, \forall e \in E$. Since the traffic in any edge connected to x must be “serviced” by the “hub” x , the interpretation of $\tau(x)$ is the *traffic rate* sustained by x . Therefore, the fundamental question addressed by this paper can be reformulated, from a very local point of view, as that of existence of vertices with very high traffic rate. The above string indicates that such vertices with very high traffic rates can be sought via the load measure Λ_t . The latter aspect will play a crucial role in Section 6.

The traffic $\tau(x)$ can be computed by counting the number of geodesics $[s, t]$ having x as a vertex. One should consider two classes of such geodesics, though: those that traverse x , that is, $x \neq s, x \neq t$, and those that either start at x , $[x, t]$, or terminate at x , $[s, x]$. The first ones have to be serviced twice, once in the input queue, once in the output queue; the others have to be serviced just once. There are $2(N - 1)$ geodesics starting or terminating at x . The *betweenness centrality* $\beta_c(x)$ of a vertex x is defined, for a uniquely geodesic graph, as the number of shortest paths that have x as a vertex (this includes those geodesics starting at x or terminating at x). For uniformly defined demand, that is, $\Lambda_d(x, y) = 1, \forall x \neq y$, the connection between the traffic rate and the betweenness is easily seen to be

$$\tau(x) = 2(\beta_c(x) - 2(N - 1)) + 2(N - 1) = 2\beta_c(x) - 2(N - 1).$$

3. Basic Conjectures

For the sake of simplicity, we introduce a network curvature concept restricted to planar communication graphs and based on Alexandrov angles [Bridson and Haefliger 99].

Let $(ab_1 = ab_{\text{deg}(a)+1}, ab_2, \dots, ab_{\text{deg}(a)})$ be a cyclic ordering of the set of edges attached to the vertex a . A geodesic *triangle* is defined as

$$\Delta abc = [a, b] \cup [b, c] \cup [c, a].$$

The *Alexandrov angle* α_k at the vertex a of the geodesic triangle $\Delta ab_k b_{k+1}$ is

$$\alpha_k = \cos^{-1} \frac{d(a, b_k)^2 + d(a, b_{k+1})^2 - d(b_k, b_{k+1})^2}{2d(a, b_k)d(a, b_{k+1})}.$$

Then the (Gaussian) *curvature* at the vertex a is defined as

$$\kappa(a) = \frac{2\pi - \sum_{i=1}^{\deg(a)} \alpha_k}{\sum_{k=1}^{\deg(a)} \text{area}(\Delta ab_k b_{k+1})},$$

where $\text{area}(\Delta ab_k b_{k+1})$ denotes the area of the geodesic triangle $\Delta ab_k b_{k+1}$, easily computable from the distances via Heron's formula.

It is easily seen that for the number-of-hops metric ($w(e) = 1$), we have $\alpha_k = \pi/6$; therefore, $\kappa(a) < 0$, $\kappa(a) = 0$, or $\kappa(a) > 0$ depending on whether $\deg(a) > 6$, $\deg(a) = 6$, or $\deg(a) < 6$, respectively.

This simple definition is introduced to construct some easily understood illustrative examples. In the main body of the paper, though, we will use the Gromov definition of negative curvature for graphs (see Section 5).

Definition 3.1. [Jost 97] The *moment of inertia* of a connected weighted graph G with respect to a vertex v is defined as $\phi_G^{(2)}(v) = \sum_{v^i \in V} d^2(v, v^i)$.

Since the edges are treated as massless, this concept refers more to the underlying metric space structure (V, d) than to the weighted graph structure (V, E, d) . Observe that this inertia may be infinite.

Definition 3.2. [Jost 97] A *center of mass* or *centroid* of the weighted graph G is defined as a vertex that achieves the infimum of the inertia: $\phi_G^{(2)}(\text{cm}(G)) = \inf_{v \in V} \phi_G^{(2)}(v) < \infty$.

If the graph is infinite ($|V| = \infty$), this definition requires existence of a vertex v such that $\phi_G^{(2)}(v) < \infty$. If we relax the minimum to be anywhere on the graph, it will in general be achieved on an edge, but restricting it to vertices makes it easier to relate it to the traffic.

We can now pose our conjectures.

Conjecture 3.3. (Negative curvature.) Consider a large but finite ($|V| < \infty$) negatively curved graph G , subject to uniformly distributed demand. Then:

1. There are a very few nodes v that have very high traffic rate $\tau(v)$ as measured by $\beta_c(v)$; furthermore, the vertices of highest traffic rate are in a small neighborhood of the vertices of minimum inertia.
2. If the graph has a symmetry group that fixes some point 0, this point achieves the unique minimum of the inertia and the maximum of the traffic rate.

The first part of this conjecture is illustrated on planar graphs in Section 4. As a model of a *large* negatively curved graph, it is tempting to take the full hyperbolic space \mathbb{H}^n . The latter is (globally) affine symmetric [Helgason 00, Chapter IV], [Jost 98, Section 6.2], [Kobayashi and Nomizu 96, Chapter XI, Example 10.2], and the transitivity of the symmetry group [Kobayashi and Nomizu 96, XI, Theorem 1.4] would make the traffic uniformly distributed if it weren't for the lack of convergence of the traffic function for uniformly distributed demand on infinite space.

Restricting the network model to finite subsets of \mathbb{H}^n breaks enough of the symmetry to create traffic spikes, even asymptotically as the size increases to infinity. An estimate of how close the center of mass and the point of maximum traffic are seems to be beyond our reach for general network models, at least at this stage. However, for hyperbolic balls, we prove that the two coincide and, as our major result, that the sharpest traffic spike scales as $\text{vol}(G)^2$.

Conjecture 3.4. (Nonnegative curvature.) *Consider a large but finite ($|V| < \infty$) nonnegatively curved graph G subject to uniformly distributed demand. Then:*

1. *Both the traffic and inertia functions τ and β_c are more evenly distributed than in the case of a negatively curved graph.*
2. *If the graph has a vertex-transitive symmetry group, then both the traffic and the inertia are uniform.*

Even though this conjecture asserts that for both zero-curvature graphs (e.g., Euclidean lattices) and positively curved graphs, the traffic is more smoothly distributed than in the case of negatively curved graphs, there is a significant difference between the two cases. A Euclidean lattice graph could be infinite, while positively curved (cubic) graphs are *finite* by Higuchi's theorem [DeVos and Mohar 07, Higuchi 01, Réti et al. 05, Sun and Yu 04]. Positively curved graphs need not be truncated and hence enjoy more symmetry than truncated Euclidean or hyperbolic graphs. This point is easy to illustrate on the 1-skeletons of the boundaries of the Platonic solids, all of which are positively curved. (For those Platonic solids that have triangular faces, the above curvature formula applies; for the other Platonic solids, the more general Higuchi–Mohar–DeVos formula [DeVos and Mohar 07] should be applied.) The transitivity of the symmetry groups of the Platonic solids, together with the uniform distribution of the demand, implies the uniform distribution of the traffic on the boundary graphs. For more general positively curved graphs, this conjecture is proved using a Riemannian manifold model with its curvature bounded

as $0 < k_1^2 \leq \kappa(x) \leq k_2^2$. Here the symmetry is broken by the nonisotropic curvature.

Truncating \mathbb{E}^n to secure convergence results in symmetry-breaking, creating traffic “bumps,” not as sharp as those of the negatively curved case, but sharper than those in the positively curved spaces. This statement is quantified by showing that the maximum traffic scales as $\text{vol}(G)^{1+1/n}$.

4. Some Examples

We consider two simple examples: The first one (symmetric graph) highlights the worsening of the congestion as the curvature decreases. The second one (asymmetric graph) highlights the relationship between the curvature and the inertia.

4.1. Almost Symmetric Graph: Congestion versus Curvature

We construct a graph from a single vertex (#1) followed by the addition of seven neighbors (#2–#8) in a counterclockwise sense, as shown in Figure 1, top. We then proceed from vertex #2 and add vertices #9, #10, #11, #12, so that vertex #2 has valence 7. We then add neighbors to vertex #3, and proceed recursively, until we obtain the graph shown in Figure 1, top. With seven neighbors for each vertex, except the boundary ones, the graph is negatively curved, because

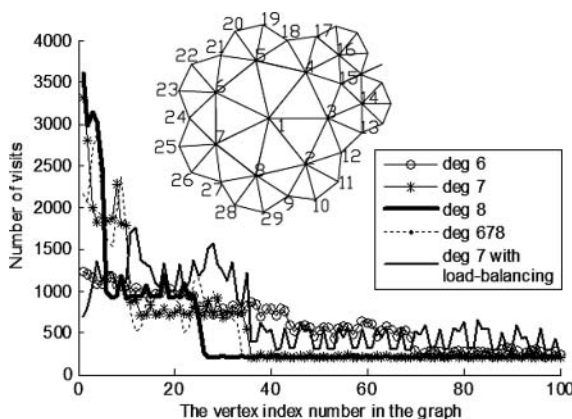


Figure 1. Top: Almost symmetric negatively curved network of valence 7. Bottom: Traffic function $2\beta_c$, “number of visits,” versus vertex number; clearly, congestion at the centroid increases as the curvature becomes more negative. Also shown is the traffic function after curvature-based load balancing (see [Lou 08] for details).

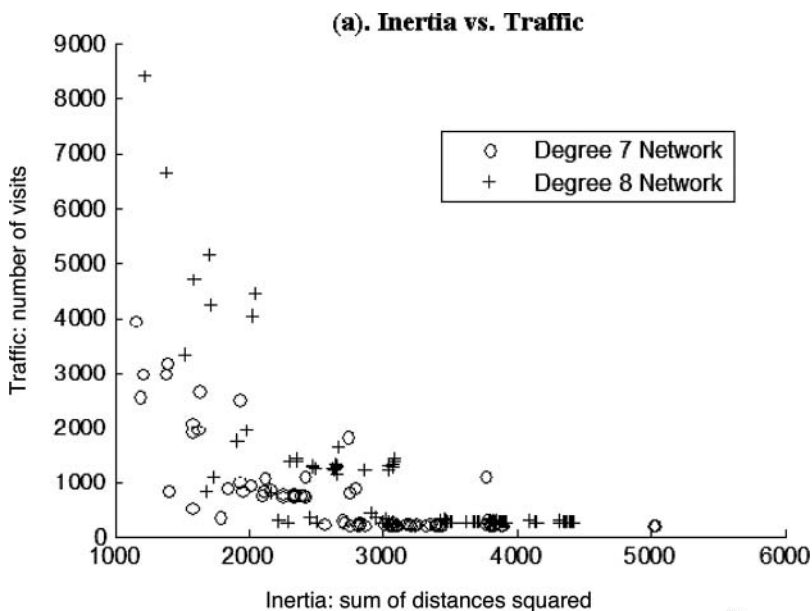


Figure 2. Scatter plot of traffic, measured as number of visits $2\beta_c$, and inertia $\phi^{(2)}$ on planar graphs of valence 7 and 8. The valence-7 graph is shown in Figure 3; the valence-8 graph is a valence-8 version of the same graph. Observe that the $(\phi^{(2)}, \beta_c)$ data seem to cluster around a relationship of the form $\phi^{(2)} \times \beta_c = \text{constant}$. From the scatter plot, it is quite obvious that the maximum traffic occurs at the minimum inertia, and on the other hand, the minimum traffic occurs at the maximum inertia.

$\sum_{i=1}^7 \alpha_i = 7 \times \frac{\pi}{3} > 2\pi$. For the traffic simulation, the construction is iterated until 100 vertices are obtained. The same construction can be used to generate graphs of valence 6 (vanishing curvature) and 8 (negative curvature). The plots of Figure 1, bottom, clearly demonstrate that the maximum congestion worsens as the curvature becomes more negative.

4.2. Asymmetric Graph: Congestion versus Inertia

Here we consider a graph of valence 7, constructed the same way as in the preceding case, except that some “appendices” have been deliberately added to make the graph asymmetric. The results are shown in Figures 2 and 3. Because the graph is negatively curved, and as conjectured, the traffic has its maximum precisely at the point where the inertia is minimal, we observe that the traffic and the inertia are inversely proportional.

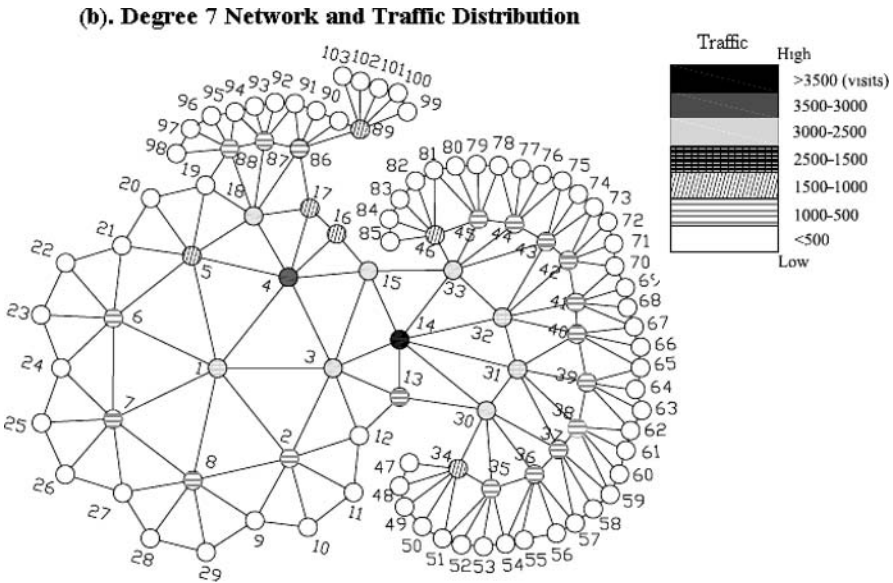


Figure 3. Spatial distribution of traffic on a planar graph of valence 7; the graph has been made asymmetric so that the centroid does not coincide with the obvious symmetry center. The centroid occurs at vertex #14. Also observe the heavy traffic at vertex #4, because it is the gateway between the top appendix and the right side of the network.

5. From Graphs to Riemannian Manifolds

As stated in the introduction, it has become standard routine in complex networks to check their Gromov negatively curved property as a way to validate arguments based on such models as the Poincaré disk. For the sake of mathematical accuracy, we formulate the basic definition and result.

Definition 5.1. A geodesic metric space (G, d_G) , for example a graph, is said to be *Gromov hyperbolic* if there exists a $\delta < \infty$ such that every geodesic triangle $\triangle abc$ has an inscribed triangle $\triangle xyz$, $x \in [b, c], y \in [a, c], z \in [a, b]$, of perimeter not exceeding δ , that is, $d(x, y) + d(y, z) + d(z, x) \leq \delta$.

This definition makes sense only for infinite graphs. Even though there exists a Gromov concept for finite graphs [Jonckheere et al. 08], we really do not need the latter here, since our congestion analysis is asymptotic for very large graphs (or manifolds).

Theorem 5.2. [Bonk and Schramm 00] *Let (G, d_G) be a Gromov hyperbolic metric geodesic space with bounded growth at some scale. Then there exist an integer n , a convex subset $D \subseteq \mathbb{H}^n$, constants λ, k , and a map $f : G \rightarrow D$ such that*

$$|\lambda d_G(u, v) - d_D(f(u), f(v))| \leq k, \quad \forall u, v \in G,$$

and $\sup_{x \in D} d_D(x, f(G)) \leq k$.

This theorem makes precise the somewhat loose statement made in the introduction regarding graph–manifold identification. Clearly, G , after a scaling λ , can be identified with D via f and subject to a bounded error k . The condition of bounded growth at some scale is satisfied, for example, when G is a finite-valence graph.

Since this graph–manifold identification entails a bounded error, large-scale problems on graphs can be mapped to more manageable continuous geometry problems on manifolds.

6. Differential-Geometric Proof of Conjectures

To justify our numerical results related to congestion on planar graphs of uniform valence 7, 8, 9, we could develop a Poincaré disk model (see Figure 4, right panel). The latter is a faithful model in the sense that the graphs of valence 7, 8, 9 are quasi-isometric to the Poincaré disk. Recall that the Poincaré disk

$$\mathbb{D} = \{z = x + jy \in \mathbb{C} : |z| < 1\}$$




Figure 4. Traffic load in Euclidean space and Poincaré disk, the latter taken as a model of \mathbb{H}^2 . The radii of the balls in the complex plane model are written \bar{r}, \bar{R} , and the corresponding hyperbolic measurements are $r = \tanh^{-1}(\bar{r})$, $R = \tanh^{-1}(\bar{R})$.

inherits its hyperbolic structure through the metric

$$ds^2 = \frac{4 dz d\bar{z}}{(1 - |z|^2)^2}.$$

The latter leads to the area element

$$d \text{ area} = \frac{4 dx dy}{(1 - |z|^2)^2}.$$

In the spirit of generalization, we develop models in the hyperbolic space \mathbb{H}^n , even in a Riemannian manifold M with arbitrary Riemannian metric ds^2 and volume form $d \text{ vol}$, provided its sectional curvature is bounded as $-k_2^2 \leq \kappa(x) \leq -k_1^2 < 0$. To emphasize the role of the curvature, we also look at the Euclidean space \mathbb{E}^2 as a model of graphs of valence 6 and further generalize the results to \mathbb{E}^n .

6.1. Hyperbolic versus Euclidean Traffic Load

Consider, in a Riemannian manifold M , a large ball $B_R(0)$ of radius R with its center at the origin together with a convex subset $X \subset B_R(0)$. The *traffic load* in X is defined as (compare with (2.1))

$$\Lambda_t(X) = \iint_{(\mathbf{s}, \mathbf{t}) \in B_R(0) \times B_R(0)} \ell(X \cap [\mathbf{s}, \mathbf{t}]) d\Lambda_d(\mathbf{s}, \mathbf{t}).$$

As for the graph model, in this Riemannian context, the demand is *uniformly distributed* in the sense that $d\Lambda_d(\mathbf{s}, \mathbf{t})$ is the product volume $d \text{ vol}(\mathbf{s}) d \text{ vol}(\mathbf{t})$. The *normalized* traffic load in X is defined as follows:

$$\begin{aligned} \lambda_t(X) &:= \frac{\Lambda_t(X)}{\text{vol}(B_R(0))^2} \\ &= \frac{1}{\text{vol}(B_R(0))^2} \iint_{(\mathbf{s}, \mathbf{t}) \in B_R(0) \times B_R(0)} \ell(X \cap [\mathbf{s}, \mathbf{t}]) d\Lambda_d(\mathbf{s}, \mathbf{t}). \end{aligned} \tag{6.1}$$

We take the hyperbolic plane \mathbb{H}^2 as a roughly similar model of an infinite negatively curved planar graph [Bonk and Schramm 00] and carry over the analysis to n dimensions. In order to secure convergence of the traffic load, we restrict ourselves to the finite domain $B_R(0) \subset \mathbb{H}^n$, representative of a large but finite negatively curved graph G , and then we do the asymptotic analysis as $R \rightarrow \infty$.

We now state our two major results. The proofs are in the appendix (Section 10):

Theorem 6.1. *If $B_R(0)$ and $X = B_r(0)$, $r \ll R$, are concentric balls in \mathbb{H}^n and if $d\Lambda_d(\mathbf{s}, \mathbf{t})$ is the product volume measure, we have*

$$\lambda_t(B_r(0)) \asymp c_1(n)r^n,$$

where $c_1(n) > 0$. Furthermore, in a Riemannian manifold of curvature bounded as $-k_2^2 \leq \kappa(x) \leq -k_1^2 < 0 \forall x \in M$, the asymptotic traffic load is bounded as $c_{k_1}(n)r^n \leq \lambda_t(X) \leq c_{k_2}(n)r^n$.

Theorem 6.2. *If $B_R(0)$ and $X = B_r(0)$, $r \ll R$, are n -dimensional Euclidean balls (vanishing curvature) and if $d\Lambda_d(\mathbf{s}, \mathbf{t})$ is the product volume measure, we have*

$$\lambda_t(B_r(0)) \asymp c_0(n) \frac{r^n}{R^{n-1}}$$

for some constant $c_0(n) > 0$, independent of r, R .

Probably the most important conclusion to be drawn from the preceding two theorems is that the normalized traffic load in the small hyperbolic ball $B_r(0)$ remains bounded from below as $R \rightarrow \infty$, whereas the same normalized traffic but in the small Euclidean ball goes to zero as $R \rightarrow \infty$. Figure 4 provides an intuitive explanation as to why this discrepancy happens. Because the hyperbolic geodesics are “arched” toward the center where the small ball lies, their average length in the small hyperbolic ball is much larger than in the small Euclidean ball.

Figure 4 gives a clue about elementary proofs of Theorems 6.1 and 6.2 in two dimensions. Clearly, the natural parameterization of the $\mathbf{s}(x, y)$, $\mathbf{t}(x', y')$ points in the traffic load integral is via polar coordinates; the difficulty is to compute the Jacobian from $dx dy dx' dy'$, (respectively,

$$\frac{dx dy dx' dy'}{(1 - (x^2 + y^2))^2 (1 - (x'^2 + y'^2))^2})$$

to the area-squared element in polar coordinates of Euclidean (respectively hyperbolic) space. The details are available in [Lou 08]; this elementary but more explicit proof yields a specific value for $c_0(2) = 1/\pi$. The proof in the appendix is much more conceptual.

Observe that besides its definition as the *traffic load*, λ_t could be interpreted in another way. Instead of (6.1), consider the following:

$$\frac{1}{\text{vol}(B_R(0))^2} \iint_{B_R(0) \times B_R(0)} I_{X \cap [s, t]}(\mathbf{s}, \mathbf{t}) d\Lambda_d(\mathbf{s}, \mathbf{t}),$$

where $I_{X \cap [s, t]}(\mathbf{s}, \mathbf{t}) = 1$ if $X \cap [s, t] \neq \emptyset$ and 0 otherwise. The above is clearly the proportion of communication paths transiting through X . If X is between two

balls, that is, $B_{r_1}(0) \subseteq X \subseteq B_{r_2}(0)$, the above is of order $\lambda_t(X)/r_i$. Thus $\lambda_t(X)$ can be interpreted as the percentage of traffic passing through X .

6.2. Minimum Inertia

Let $D \subset M$ be a convex domain of finite volume in a Riemannian manifold with distance $d(\cdot, \cdot)$. For $x \in D$, define (compare with Definition 3.1)

$$\phi_D^{(p)}(\mathbf{x}) = \int_D d(\mathbf{x}, \mathbf{x}')^p d \text{vol}(\mathbf{x}').$$

For $p = 2$, the above is called the *moment of inertia* of D relative to x . Next, assume that there exists a point $\text{cm}(D) \in D$ such that (compare with Definition 3.2)

$$\phi_D^{(p)}(\text{cm}(D)) = \inf_{\mathbf{x} \in D} \phi_D^{(p)}(\mathbf{x}) < \infty.$$

For $p = 2$, $\text{cm}(D)$ is called a *center of mass* or *centroid* of D (see [Jost 97, Definition 3.2.1]). This concept was apparently introduced by Élie Cartan [Berger 00, p. 47].

Arguments related to existence and uniqueness of the center of mass rely on strict convexity of $\mathbf{x} \mapsto d(\mathbf{x}, \mathbf{y})^p$ and $\mathbf{x} \mapsto \phi^{(p)}(\mathbf{x})$. In negatively curved spaces, $p \geq 1$ suffices, whereas in nonpositively curved spaces (e.g., Euclidean spaces) the stricter condition $p > 1$ is required for strict convexity. If $\phi^{(p)}$ is strictly convex, uniqueness is guaranteed [Jost 97, Lemma 3.1.1].

Theorem 6.3. *Let M be a complete Riemannian manifold with its curvature bounded as $-k_2^2 \leq \kappa(\mathbf{x}) \leq -k_1^2 < 0$. Then the inertia of $B_R(0) \subset M$ relative to the point \mathbf{x} ,*

$$\phi_{B_R(0)}^{(2)}(\mathbf{x}) = \int_{B_R(0)} d(\mathbf{x}, \mathbf{x}')^2 d \text{vol}(\mathbf{x}'),$$

has a unique minimum; furthermore, for $k_1 = k_2$, this minimum is $\mathbf{x} = 0$.

Proof. Obviously, M is a complete Riemannian manifold of nonpositive curvature, and hence it is a Busemann nonpositively curved (NPC) space [Jost 97, p. 45]. Define the measure $\mu(\cdot) = \text{vol}(\cdot) I_{B_R(0)}(\cdot)$, where $I_{B_R(0)}$ denotes the indicator of $B_R(0) \subset M$. Obviously, the measure μ has finite support and $\mu(B_R(0)) < \infty$. Therefore, by [Jost 97, Theorem 3.2.1],

$$\int_M d(\mathbf{x}, \mathbf{x}')^2 \mu(d\mathbf{x}') = \int_{B_R(0)} d(\mathbf{x}, \mathbf{x}')^2 d \text{vol}(\mathbf{x}')$$

is finite and has a unique infimum. Furthermore, if $k_1 = k_2 = 1$, then M has orientation-preserving isometry group $\text{SO}^+(n, 1)$. Under the subgroup that fixes 0, $B_R(0)$ is invariant, and therefore $\phi_{B_R(0)}^{(2)}$ has the same symmetries. Therefore, the only way to secure uniqueness of the infimum is $\mathbf{x}^* = 0$. \square

Obviously, whenever D has a nontrivial symmetry group that fixes some point 0, the minimum inertia is unique and occurs at 0.

6.3. Maximum Load

Theorem 6.4. *Restricted to convex subsets X of $B_R(0) \subset \mathbb{H}^n$ with the same hyperbolic volume, $\lambda_t(X)$ reaches its maximum for a ball centered at the origin of $B_R(0)$.*

Proof. The proof is a corollary of the proof of the appendix. First, observe that $\mu_k(\mathbf{x})$ is spherically invariant, that is, $\mu_k(\mathbf{x}) = \mu_k(S\mathbf{x})$, for all S in the subgroup of $\text{SO}^+(n, 1)$ that fixes 0. Set $\mu_k(\|\mathbf{x}\|)$ as the common value of $\mu(\mathbf{x})$ for $\mathbf{x} \in \partial B_{\|\mathbf{x}\|}(0)$.

Next, we want to prove that

$$\int_{S_{\mathbf{x}}^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \text{sh}_k^{n-1}(x+y) dx dy d\vec{v}$$

subject to the restriction that $e(\vec{v}) + e(-\vec{v}) = E$ is maximal for $e(\vec{v}) = e(-\vec{v})$. This is proved using the augmented functional

$$\int_{S_{\mathbf{x}}^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \text{sh}_k^{n-1}(x+y) dx dy + L(e(\vec{v}) + e(-\vec{v})),$$

where L is the Lagrange multiplier. Setting the partial derivatives relative to $e(\vec{v})$ and $e(-\vec{v})$ to zero yields

$$\begin{aligned} \int_{S_{\mathbf{x}}^{n-1}} \int_0^{e(-\vec{v})} \text{sh}_k^{n-1}(e(\vec{v}) + y) dy + L &= 0, \\ \int_{S_{\mathbf{x}}^{n-1}} \int_0^{e(\vec{v})} \text{sh}_k^{n-1}(x + e(-\vec{v})) dx + L &= 0. \end{aligned}$$

Clearly, the above yields $e(\vec{v}) = e(-\vec{v})$, and the maximum is monotone increasing with E . It follows that

$$\mu_k(\|\mathbf{x}\|) \leq \frac{1}{\text{vol}(B_R(0))^2} \int_{S_{\mathbf{x}}^{n-1}} \int_0^R \int_0^R \text{sh}_k^{n-1}(x+y) dx dy d\vec{v}.$$

But

$$\mu_k(\mathbf{0}) = \frac{1}{\text{vol}(B_R(0))^2} \int_{S_0^{n-1}} \int_0^R \int_0^R \text{sh}_k^{n-1}(x+y) dx dy d\vec{v}.$$

Thus $\mu_k(\|\mathbf{x}\|)$ reaches its maximum at $\mathbf{x} = 0$. Furthermore, it is obviously symmetric for the subgroup of $\text{SO}^+(n, 1)$ that fixes 0. Finally, from the appendix, $\mu_k(\|\mathbf{x}\|)$ is monotone decreasing with $\|\mathbf{x}\|$.

It remains to show that the optimal way to distribute the volume allocated to X is in a ball around 0. First, we show that if $\inf \lambda_t(X)$ is achieved for some X^* that does not enjoy the symmetry under the action of the subgroup of $\text{SO}^+(n, 1)$ that fixes 0, then the same $\lambda_t(X^*)$ can be achieved for another subset that has the symmetry but cannot be optimal.

Decompose the big ball as $B_R(0) = \sqcup_{i=0}^{m-1} A[r_i, r_{i+1})$, where $A[r_i, r_{i+1})$ is the annulus $\{\mathbf{x} : r_i \leq d(0, \mathbf{x}) < r_{i+1}\}$. Clearly, there exists an annulus $A[\rho_i, \rho_{i+}) \subseteq A[r_i, r_{i+1})$ such that $\lambda_t(X \cap A[r_i, r_{i+1})) = \lambda_t(A[\rho_i, \rho_{i+}))$. Thus

$$\lambda_t(X) = \sum_i \lambda_t(A[\rho_i, \rho_{i+})) \quad \text{and} \quad \lambda_t(X) = \sum_i \int_{\rho_i}^{\rho_{i+}} \mu_k(r) T(r) dr,$$

where $T(r)$ is the transverse measure such that $\text{vol}(A[r, r + dr]) = T(r)dr$, that is, $T(r)$ is the “area” of the sphere at 0 with radius r . Consider now the constrained optimization problem

$$\inf_{\rho} \sum_i \int_{\rho_i}^{\rho_{i+}} \mu_k(r) T(r) dr$$

subject to

$$\sum_i \int_{\rho_i}^{\rho_{i+}} T(r) dr = \text{vol}(X).$$

Again, a Lagrange multiplier argument proves that optimality cannot hold with some i such that $\rho_{i+} < \rho_{i+1}$. (The intuition is that since the density μ_k is monotone decreasing, optimality would require ρ_{i+1} to drop down to ρ_{i+} .) Thus $\rho_0 = 0$ and $\rho_{i+} = \rho_{i+1}$, that is, the presumed optimal annulus collapses to a ball. \square

When D has a nontrivial symmetry group that fixes some point 0, the above argument can easily be extended to prove that the maximum traffic occurs for a subset X having the same symmetry. Hence in this more general situation the center of mass and the maximum traffic point are collocated.

7. Differential Geometry versus Real Network Congestion

It is argued that no matter how theoretical our model $\Lambda_t(B_r(0))$ of the traffic load is, it is remarkably accurate at predicting how the “load at the center,” here 0, scales with the number of vertices, N , in a *real* network.

In [Narayan and Saniee 09], the traffic load at the “center” of a sample of networks from the Rocketfuel database [Spring et al. 04] has been numerically found to scale as N^2 . Using scaled Gromov analysis [Jonckheere et al. 08], it was asserted in [Narayan and Saniee 09] that the networks from that database are negatively curved. In [Narayan and Saniee 09], the “center” is defined as the “network core” or the “set of nodes that are at the intersection of the majority of geodesics.” The latter concept is unmistakably the same as our theoretical *center of mass* concept. In the same paper, the authors provide experimental evidence that the traffic load at the center of Watts–Strogatz small-world 2-dimensional networks scales as $N^{1.5}$. It is commonly admitted, and it has been proved using scaled Gromov δ -analysis [Jonckheere et al. 08], that those networks are not hyperbolic; they are rather Euclidean, even positively curved, in a certain range of the connectivity and rewiring parameters [Lohsoonthorn 03, Section 6.4.2].

Fact 7.1. [Narayan and Saniee 09] *Let $\beta_c(v)$ be the betweenness centrality of the vertex v in a network.*

1. *In the Rocketfuel database [Spring et al. 04] of real networks, which are scaled-Gromov hyperbolic by the definition of [Jonckheere et al. 08], the maximum traffic rate scales as*

$$\max_v \beta_c(v) = \Theta(N^2).$$

2. *For synthetic 2-dimensional Euclidean lattice networks, the maximum traffic rate scales as*

$$\max_v \beta_c(v) = \Theta(N^{1.5}).$$

To proceed to a continuous-geometry justification of the above, the graphs are embedded in appropriate manifolds, using the Bonk–Schramm theorem for a Gromov hyperbolic graph, or using the trivial Euclidean embedding for the Euclidean lattice graphs. The Bonk–Schramm embedding $f : G \rightarrow D \subseteq \mathbb{H}^n$ maps the geodesic flow on the graph G to a quasigeodesic flow on the manifold D . Consider a vertex v on the graph. We clearly have $\Lambda_t(\text{star}(v)) \approx \Lambda_t(B_r(f(v)))$, for some r of order the mesh of the graph. But as argued in Section 2, $\Lambda_t(\text{star}(v))$ scales as $\tau(v)$, which scales as $\beta_c(v)$. The continuous-geometry model of the traffic metric of Fact 7.1 is therefore $\Lambda_t(B_r(\mathbf{x}))$ for some $\mathbf{x} \in D$. Obviously, the continuous-geometry equivalent of N is

$$N = \text{vol}(D).$$

Fact 7.2. (Theorems 6.1 and 6.2.) *Let $\Lambda_t(B_r(\mathbf{x}))$ be the traffic load in a small metric ball $B_r(\mathbf{x})$ embedded in a large metric ball $B_R(0)$, where $R \gg r$.*

- *In a hyperbolic ball $B_R(0) \subset \mathbb{H}^n$, we have*

$$\max_{\mathbf{x}} \Lambda_t(B_r(\mathbf{x})) = \Theta(\text{vol}(B_R(0))^2) = \Theta(N^2).$$

- *In a Euclidean ball $B_R(0) \subset \mathbb{E}^n$, we have*

$$\max_{\mathbf{x}} \Lambda_t(B_r(\mathbf{x})) = \Theta\left(\frac{\text{vol}(B_R(0))^2}{R^{n-1}}\right) = \Theta(N^{1+1/n}).$$

In view of the striking consistency between Facts 7.1 and 7.2, our model correctly predicts how the maximum traffic load scales as a function of N or the volume of the manifold.

By the same token, we have extended the Euclidean results of [Narayan and Sanjeev 09] to the case $n > 2$. This generalization allows us to observe that as the dimension increases, the Euclidean congestion decreases and the gap between traffic loads in hyperbolic and Euclidean spaces increases.

While there are computational methods to check the Gromov property of a network [Ariaei et al. 08, Jonckheere et al. 10, Jonckheere et al. 08, Lohsoonthorn 03], associating a dimension with a complex network is another matter entirely. The remarkable feature is that the asymptotic traffic analysis in negatively curved spaces transcends the dimension.

8. The Case of Positive Curvature

The case of a Riemannian manifold positively curved as $0 < k_1^2 < \kappa < k_2^2$ is significantly different from that of a nonpositively curved manifold, because the diameter of the former is bounded as π/k_1 . Furthermore, by the sphere theorem [do Carmo 92, Chapter 13], if $k_1 = k_2/2$, the manifold is homeomorphic to a sphere. Thus, in contrast to the nonpositive curvature case, we cannot take a very large manifold with fixed positive curvature; so we will have to take an arbitrarily large manifold of positive curvature decreasing to 0. The first part of the following theorem is completely trivial; the second part is easily proved by bounding the integrand of (10.2) and taking the upper limit of the two inner integrals to be $\pi R = \pi/k$.

Theorem 8.1. *Let M be a connected n -dimensional Riemannian manifold with constant curvature $\kappa = k^2 > 0$. Then the following hold:*

1. *Both the normalized traffic λ_t and the inertia $\phi_M^{(2)}$ over the whole manifold are uniform.*
2. *$\lambda_t(B_r(0)) = \Theta(1/R^{2n-2})$, where $B_{\pi R}(0) = M \setminus \{\text{antipodal point of } 0\}$.*

Comparing with Theorem 6.2, it is clear that the normalized traffic in positively curved space decays with n even faster than in the Euclidean case.

Finally, using the comparison formula of the appendix, the following is easily derived:

Theorem 8.2. *Let M be a connected Riemannian manifold with curvature bounded as $k_1^2 < \kappa \leq k_2^2$. Then*

$$\mu_{k_2}^+(\mathbf{x}) < \frac{d\lambda_t(\mathbf{x})}{d\mathbf{x}} \leq \mu_{k_1}^+(\mathbf{x}).$$

Thus, under varying but bounded curvature, the traffic density in a small neighborhood of \mathbf{x} remains bounded between the densities in fixed curvature. (Observe that by symmetry, $\mu_{k_i}^+(\mathbf{x})$, $i = 1, 2$, are independent of \mathbf{x} .)

9. Conclusion

We have provided a mathematical justification of the experimentally observed fact that negatively curved networks—even with uniform curvature—driven by uniformly distributed demand have small areas of very high traffic concentration. Nonnegatively curved networks, on the other hand, do not exhibit this phenomenon as dramatically as negatively curved networks. In fact, uniformly positively curved networks have uniform traffic distribution; more generally, bounds on the curvature implies bounds on the maximum traffic.

Since the root cause of congestion in a network is its negative curvature, load balancing could be achieved by controlling the curvature to become and remain positive [Lou 08], despite outages and varying demand.

The areas of maximum traffic have been narrowed down to areas of low inertia. It has been proved that networks with enough symmetry have colocated maximum traffic and minimum inertia. But as already stated, for general networks, bounding the distance between the two points is a challenging problem.

The traffic dealt with here is that driven by a uniformly distributed demand. The extension to nonuniformly distributed demand is not a major hurdle. It suffices to redefine the inertia as $\phi_G^{(2)}(v) = \sum_{v^i} d^2(v, v^i)(\sum_j \Lambda_d(v^i, v^j))$.

Finally, the present analysis is a spatial one. The temporal component would bring the dynamics of packet drops and retransmission into the picture. Early ns-2 simulations have shown that UDP traffic has its maximum packet drop at the point of maximum traffic and minimum inertia [Lou 08].

From a broader perspective, this paper points to the curvature as being the all-encompassing parameter regulating network behavior. Certainly, as the present paper demonstrates, curvature regulates the distribution of the load. In addition, some recent independent work [Sarkar et al. 09] has pointed to the curvature as the crucial parameter determining whether greedy forwarding in a sensor network would be successful.

10. Appendix

10.1. Nonpositively Curved Spaces

Let D be a convex domain in an n -dimensional Riemannian manifold M . Here convex means that for every \mathbf{s} and \mathbf{t} in D , there is a unique shortest geodesic, or shortest path, $[\mathbf{s}, \mathbf{t}]$ going from \mathbf{s} to \mathbf{t} and contained in D .

For a subset X of D , we are interested in

$$\lambda_t(X) = \frac{1}{\text{vol}(D)^2} \int_{D \times D} \ell(X \cap [\mathbf{s}, \mathbf{t}]) \, d\mathbf{s} \, d\mathbf{t},$$

where the integral is with respect to the square of the n -dimensional volume of D .

From a dynamical point of view, assuming uniform traffic between pairs of points of D , and assuming that this traffic travels at unit speed along geodesics, $\lambda_t(X)$ measures the average of the amount of time spent in X .

We provide an estimate for $\lambda_t(X)$ when M has nonpositive curvature, and an exact computation when the curvature is constant.

Consider a point \mathbf{s} in D . For every unit vector \vec{v} based at \mathbf{s} , draw the geodesic $g_{\vec{v}}$ emanating from \mathbf{s} in the direction of \vec{v} , and let $e(\vec{v})$ be the distance from \mathbf{s} to the point where the geodesic exits D . (This exit point is unique by the convexity of D .) We can similarly consider the geodesic $g_{-\vec{v}}$ emanating from \mathbf{s} in the opposite direction $-\vec{v}$, and the corresponding distance $e(-\vec{v})$.

For $k > 0$, define

$$\mu_k(\mathbf{s}) = \frac{1}{\text{vol}(D)^2} \int_{S_{\mathbf{s}}^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \frac{1}{k^{n-1}} \sinh^{n-1}(kx + ky) \, dx \, dy \, d\vec{v}, \quad (10.1)$$

where the outer integral is with respect to the $(n - 1)$ -dimensional volume of the unit sphere S_s^{n-1} consisting of all unit vectors \vec{v} based at \mathbf{s} . The inner double integral can of course be computed by elementary calculus.

For $k = 0$, let

$$\mu_0(\mathbf{s}) = \frac{1}{\text{vol}(D)^2} \int_{S_s^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} (x + y)^{n-1} dx dy d\vec{v}.$$

Theorem 10.1. *If the curvature of M is constant and equal to $-k^2$, then*

$$\lambda_t(X) = \int_X \mu_k(\mathbf{x}) d\mathbf{x}.$$

More generally, if the curvature of M is everywhere bounded between two non-positive constants $-k_1^2$ and $-k_2^2$ with $0 \leq k_1 \leq k_2$, then

$$\int_X \mu_{k_1}(\mathbf{x}) d\mathbf{x} \leq \lambda_t(X) \leq \int_X \mu_{k_2}(\mathbf{x}) d\mathbf{x}.$$

In other words, the quantity $\lambda_t(X)$ is estimated by the integral of the density functions μ_k .

Proof. Let \mathcal{T} be the set of triples $(\mathbf{s}, \mathbf{t}, \mathbf{x})$ where $\mathbf{s}, \mathbf{t}, \mathbf{x}$ are points of D such that \mathbf{x} is located on the geodesic $[\mathbf{s}, \mathbf{t}]$ and inside of X . There are two natural measures that can be put on \mathcal{T} .

The first one is the product $ds dt d\ell$ of the volume in \mathbf{s} , the volume in \mathbf{t} , and the arc-length parameter ℓ along the geodesic $[\mathbf{s}, \mathbf{t}]$. The reason we are considering \mathcal{T} with this measure is that

$$\int_{D \times D} \ell(X \cap [\mathbf{s}, \mathbf{t}]) ds dt = \int_{\mathcal{T}} ds dt d\ell.$$

The second one requires a different description of \mathcal{T} . If we consider the unit vector \vec{v} based at \mathbf{x} and pointing in the direction of \mathbf{s} (so that $-\vec{v}$ points in the direction of \mathbf{t}), the distance s from \mathbf{x} to \mathbf{s} , and the distance t from \mathbf{x} to \mathbf{t} , then the three points $(\mathbf{s}, \mathbf{t}, \mathbf{x})$ are completely determined by the \mathbf{x} -based vector \vec{v} and by the numbers s and $t \geq 0$.

Recall that the *unit tangent bundle* $T^1 X$ consists of all unit vectors \vec{v} based at points of X . This is a $(2n - 1)$ -dimensional manifold (n dimensions for the base point, $n - 1$ for the direction), and the metric of X naturally lifts to a Riemannian metric on $T^1 X$ via the Levi-Civita connection.

The above construction identifies the set \mathcal{T} of triples $(\mathbf{s}, \mathbf{t}, \mathbf{x})$ such that $\mathbf{x} \in X \cap [\mathbf{s}, \mathbf{t}]$ with the subset of $T^1 X \times \mathbb{R} \times \mathbb{R}$ consisting of those (\mathbf{v}, s, t) such that \mathbf{v} is a point of $T^1 X$ and such that $0 \leq s \leq e(\vec{v})$ and $0 \leq t \leq e(-\vec{v})$, with $e(\vec{v})$ the

distance between \mathbf{x} and the point $e(\vec{v})$ where the geodesic starting at \mathbf{x} along the direction \vec{v} exits D (see Figure 4, right panel). Considering the volume form $d\mathbf{v}$ on T^1X now provides another measure $d\mathbf{v} ds dt$ on \mathcal{T} .

When the curvature of M is bounded between $-k_1^2$ and $-k_2^2$, the two measures $ds dt dl$ and $d\mathbf{v} ds dt$ on \mathcal{T} can be compared by the standard Riemannian arguments on the variation of geodesics (using Jacobi fields):

Lemma 10.2.

$$\operatorname{sh}_{k_1}^{n-1}(s+t) d\mathbf{v} ds dt \leq ds dt dl \leq \operatorname{sh}_{k_2}^{n-1}(s+t) d\mathbf{v} ds dt,$$

where $\operatorname{sh}_k(x) = \frac{1}{k} \sinh(kx)$ if $k > 0$ and $s_0(x) = x$.

Therefore,

$$\begin{aligned} \lambda_t(X) &= \frac{1}{\operatorname{vol}(D)^2} \int_{\mathcal{T}} ds dt dl \\ &\leq \frac{1}{\operatorname{vol}(D)^2} \int_{\mathcal{T}} \operatorname{sh}_{k_2}^{n-1}(s+t) ds dt d\mathbf{v} \\ &= \frac{1}{\operatorname{vol}(D)^2} \int_{T^1X} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \operatorname{sh}_{k_2}^{n-1}(s+t) ds dt d\mathbf{v} \\ &= \frac{1}{\operatorname{vol}(D)^2} \int_X \int_{S_{\mathbf{x}}^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \operatorname{sh}_{k_2}^{n-1}(s+t) ds dt d\vec{v} d\mathbf{x} \\ &= \int_X \mu_{k_2}(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

Observe that on the fourth line, $d\vec{v}$ stands for the $(n-1)$ -dimensional volume on the sphere $S_{\mathbf{x}}^{n-1}$ of unit vectors based at \mathbf{x} , whereas $d\mathbf{v}$ represents the $(2n-1)$ -dimensional volume of T^1X in the two previous lines.

The inequality

$$\lambda_t(X) \geq \int_X \mu_{k_1}(\mathbf{x}) d\mathbf{x}$$

is proved by the same argument. \square

For instance, consider the case that D and X are two concentric balls in an n -manifold of constant curvature $-k^2 \leq 0$, of respective radii R and r with $r \ll R$. Then, $e(\vec{v}) \asymp R$ for every vector \vec{v} based at a point of X .

Theorem 10.3. *If D and X are two concentric balls of respective radii R and r , with $r \ll R$, in the Euclidean space of dimension n , then the proportion $\lambda_t(X)$*

of traffic in D that transits through X is approximately

$$\lambda_t(X) \asymp c_0(n) \frac{r^n}{R^{n-1}},$$

where $c_0(n)$ is an explicit constant depending on the dimension n .

If instead, D and X are concentric balls in n -dimensional hyperbolic space (where the curvature is -1), then

$$\lambda_t(X) \asymp c_1(n)r^n.$$

Proof. Under these hypotheses, $\lambda_t(X)$ is of order $\mu_k(\mathbf{0})\text{vol}(X)$. Write $\mu_k(\mathbf{0}) = M_k(\mathbf{0})/\text{vol}(D)^2$, where $M_k(\mathbf{0})$ is the triple integral. Then $e(\vec{v}) \asymp R$ for every vector \vec{v} based at a point of X . Therefore, in the Euclidean case, $M_0(\mathbf{0})$ is of order R^{n+1} , $\text{vol}(D)$ is of order R^n , and $\lambda_t(X)$ is of order r^n/R^{n-1} . In the hyperbolic case, $M_1(\mathbf{0})$ is of order $e^{2(n-1)R}$, $\text{vol}(D)$ is of order $e^{(n-1)R}$, and thus $\lambda_t(X)$ is of order $\text{vol}(X)$, which is itself of order r^n for r small. \square

Note the independence of R in the negatively curved case.

Proof of Lemma 10.2. The Jacobi field equation [do Carmo 92] in the hyperplane orthogonal to the geodesic $[\mathbf{x}, \mathbf{e}(\vec{v})]$ in the n -manifold M of constant curvature $-k^2$ reads $\frac{d^2 J(s)}{ds^2} - k^2 J(s) = 0$. If $dx_i = J(0)$, $i = 1, \dots, n-1$, is a linear element orthogonal to the geodesic at \mathbf{x} and $d\theta_i = J'(0)$ is the elementary angle on the sphere $S_{\mathbf{x}}^{n-1}$, then the solution to the Jacobi field equation at s reads, with a similar solution for $[x, \mathbf{e}(-\vec{v})]$ at t ,

$$\begin{aligned} ds_i &= \frac{1}{k} \sinh(ks) d\theta_i + \cosh(ks) dx_i, \\ dt_i &= -\frac{1}{k} \sinh(kt) d\theta_i + \cosh(kt) dx_i. \end{aligned}$$

In the above, ds_i , dt_i are the variations of the geodesic in the orthogonal hyperplane at \mathbf{s} , \mathbf{t} , respectively. The volume defined by the variation of \mathbf{s} at s is $d\mathbf{s} = (\wedge_{i=1}^{n-1} ds_i) ds$. Clearly, $\wedge_{i=1}^{n-1} ds_i$ is the transverse variation and ds is the longitudinal variation. A similar statement holds for $d\mathbf{t}$. Because of the skew-symmetry of exterior differential forms, it is convenient to introduce the notation

$\theta \stackrel{\sigma}{=} \omega$ to denote $\theta = (-1)^{n-1}\omega$. With this notation, we get

$$\begin{aligned}
 ds \wedge dt \wedge dl &\stackrel{\sigma}{=} \wedge_{i=1}^{n-1} (ds_i \wedge dt_i) \wedge ds \wedge dt \wedge dl \\
 &\stackrel{\sigma}{=} \wedge_{i=1}^{n-1} \left(\frac{1}{k} \sinh(k(s+t)) d\theta_i \wedge d\mathbf{x}_i \right) \wedge ds \wedge dt \wedge dl \\
 &\stackrel{\sigma}{=} \left(\frac{1}{k} \sinh(k(s+t)) \right)^{n-1} (\wedge_{i=1}^{n-1} d\theta_i \wedge d\mathbf{x}_i) \wedge ds \wedge dt \\
 &= \text{sh}_k^{n-1}(k(s+t)) (\wedge_{i=1}^{n-1} d\theta_i) \wedge ((\wedge_{i=1}^{n-1} d\mathbf{x}_i) \wedge dl) \wedge ds \wedge dt \\
 &= \text{sh}_k^{n-1}(k(s+t)) (d\vec{v} \wedge ((\wedge_{i=1}^{n-1} d\mathbf{x}_i) \wedge dl)) \wedge ds \wedge dt \\
 &= \text{sh}_k^{n-1}(k(s+t)) (d\vec{v} \wedge d\mathbf{x}) \wedge ds \wedge dt \\
 &= \text{sh}_k^{n-1}(k(s+t)) d\mathbf{v} \wedge ds \wedge dt.
 \end{aligned}$$

Observe that to go from the second to the third line, we need skew-commutativity of \wedge along with some elementary hyperbolic trigonometry. The inequality follows from the Rauch comparison argument [Kobayashi and Nomizu 96, VIII, Theorem 4.1], which in this context refers to the monotone increasing property of the solution $\text{sh}_k(x)$ to the Jacobi field equation with k . \square

We finish this section with two corollaries that are useful in the main body of the text. A notation needs to be made more explicit: $e_{\mathbf{x}}(\vec{v})$ denotes the length of the geodesic shot from \mathbf{x} in the direction \vec{v} and terminating at $\partial B_R(0)$.

Corollary 10.4. *For fixed \vec{v} , $e_{\mathbf{x}}(\vec{v}) + e_{\mathbf{x}}(-\vec{v})$ is strictly monotone decreasing with $\|\mathbf{x}\|$.*

Proof. Let $e_{\mathbf{x}}(\vec{v})$ be the point where the geodesic starting at \mathbf{x} with direction \vec{v} exits $B_R(0)$. With a mild abuse of notation, we will let \vec{v} denote the angle between the radial $[\mathbf{0}, \mathbf{x}]$ and the tangent to the geodesic at \mathbf{x} . The same definition applies to $e_{\mathbf{x}}(-\vec{v})$. Because \mathbb{H}^n is isotropic, the geodesic $[e_{\mathbf{x}}(-\vec{v}), e_{\mathbf{x}}(\vec{v})]$ is contained in a plane, which itself contains \mathbf{x} and $\mathbf{0}$. Let $[\mathbf{0}, \mathbf{p}]$ be the radial orthogonal to that geodesic. In the right triangle $\triangle \mathbf{0pe}_{\mathbf{x}}(\vec{v})$, we get

$$\cosh \left(\frac{1}{2} \|e_{\mathbf{x}}(\vec{v})e_{\mathbf{x}}(-\vec{v})\| \right) = \frac{\cosh R}{\cosh(\|\mathbf{0p}\|)}.$$

On the other hand, in the right triangle $\triangle \mathbf{xp0}$, we have $\sinh \|\mathbf{0p}\| = \cosh(x) \sin \vec{v}$. It follows that

$$\cosh \left(\frac{1}{2} \|e_{\mathbf{x}}(\vec{v})e_{\mathbf{x}}(-\vec{v})\| \right) = \frac{\cosh R}{\sqrt{1 + \cosh^2 x \sin^2 \vec{v}}},$$

and from this the result is obvious. \square

Corollary 10.5. *For $k > 0$, the density function $\mu_k(\mathbf{s})$ is strictly monotone decreasing with $\|\mathbf{s}\|$.*

Proof. It is easily seen that the double inner integral in the definition (10.1) of μ_k is a polynomial in the variables $\exp(\pm ki(e_{\mathbf{s}}(\vec{v}) + e_{\mathbf{s}}(-\vec{v})))$, where i is some power. From the same definition, it is also obvious that the inner double integral is monotone increasing with $e_{\mathbf{s}}(\vec{v}) + e_{\mathbf{s}}(-\vec{v})$. But by the previous corollary, the latter is monotone decreasing with $\|\mathbf{s}\|$. Thus the double inner integral in (10.1) is monotone decreasing with $\|\mathbf{s}\|$, and so is the average over the sphere $S_{\mathbf{s}}^{n-1}$. \square

10.2. Positively Curved Spaces

The case of uniformly positive curvature $\kappa = k^2 > 0$ is treated in a way parallel to the preceding one. The difference resides in the Jacobi field equation, which now takes the form $d^2 J(s)/ds^2 + k^2 J(s) = 0$, with a solution involving ordinary trigonometric functions rather than hyperbolic trigonometric functions. Therefore, if for some bounded subset D of the space of curvature k^2 we define

$$\mu_k^+(\mathbf{s}) = \frac{1}{\text{vol}(D)^2} \int_{S_{\mathbf{s}}^{n-1}} \int_0^{e(\vec{v})} \int_0^{e(-\vec{v})} \frac{1}{k^{n-1}} \sin^{n-1}(k(x+y)) dx dy d\vec{v}, \quad (10.2)$$

the normalized traffic load in $X \subset D$ is given by

$$\lambda_t(X) = \int_X \mu_k^+(\mathbf{x}) d\mathbf{x}.$$

In a Riemannian manifold with its curvature bounded as $0 < k_1^2 \leq \kappa \leq k_2^2$, we have

$$\int_X \mu_{k_2}^+(\mathbf{x}) d\mathbf{x} \leq \lambda_t(X) \leq \int_X \mu_{k_1}^+(\mathbf{x}) d\mathbf{x}.$$

Acknowledgment

This research was supported by NSF Grant CNS-NetSE-1017881.

References

- [Ariaei et al. 08] F. Ariaei, M. Lou, E. Jonckheere, B. Krishnamachari, and M. Zuniga. "Curvature of Indoor Sensor Network: Clustering Coefficient." *EURASIP Journal on Wireless Communications and Networking*, 2008 (2008), 20 pages.

- [Berger 00] Marcel Berger. *Riemannian Geometry during the Second Half of the Twentieth Century*, University Lecture Series 17. Providence: American Mathematical Society, 2000.
- [Bonk and Schramm 00] M. Bonk and O. Schramm. “Embeddings of Gromov Hyperbolic Spaces.” *Geom. Funct. Analysis* 10 (2000), 266–306.
- [Bridson and Haefliger 99] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-positive Curvature*, A Series of Comprehensive Surveys in Mathematics 319. New York: Springer, 1999.
- [Cisco 05] Cisco Systems. “How Does Load Balancing Work?” Document ID: 5212, 2005.
- [DeVos and Mohar 07] M. DeVos and B. Mohar. “An Analogue of the Descartes–Euler Formula for Infinite Graph and Higuchi’s Conjecture.” *Transactions of the American Mathematical Society* 359 (2007), 3275–3286.
- [do Carmo 92] M. P. do Carmo. *Riemannian Geometry*. Boston: Birkhäuser, 1992.
- [Helgason 00] S. Helgason. *Differential Geometry and Symmetric Spaces*. Providence: Chelsea, 2000.
- [Higuchi 01] Y. Higuchi. “Combinatorial Curvature for Planar Graphs.” *J. of Graph Theory* 38 (2001), 220–229.
- [Jonckheere and Lohsoonthorn 02] E. A. Jonckheere and P. Lohsoonthorn. “A Hyperbolic Geometry Approach to Multi-path Routing.” In *Proceedings of the 10th Mediterranean Conference on Control and Automation (MED 2002)*, Lisbon, Portugal, July 2002.
- [Jonckheere and Lohsoonthorn 04] E. A. Jonckheere and P. Lohsoonthorn. “Geometry of Network Security.” In *Proceedings of the American Control Conference (ACC 2004)*, Boston, MA, June 2004.
- [Jonckheere et al. 08] E. Jonckheere, P. Lohsoonthorn, and F. Bonahon. “Scaled Gromov Hyperbolic Graphs.” *Journal of Graph Theory* 57 (2008), 157–180.
- [Jonckheere et al. 10] E. Jonckheere, F. Ariaei, and P. Lohsoonthorn. “Upper Bound on Scaled Hyperbolic Gromov Δ : 4-Point Condition.” Submitted to *Internet Mathematics*, 2010.
- [Jost 97] J. Jost. *Nonpositive Curvature: Geometric and Analytic Aspects*, Lectures in Mathematics. Boston: Birkhäuser, 1997.
- [Jost 98] J. Jost. *Riemannian Geometry and Geometric Analysis*, second edition. New York: Springer, 1998.
- [Kobayashi and Nomizu 96] S. Kobayashi and K. Nomizu. *Foundation of Differential Geometry*, volume 2. New York: Wiley, 1996.
- [Krioukov et al. 09] D. Krioukov, F. Papadopoulos, A. Vahdat, and M. Boguna. “Curvature and Temperature of Complex Networks.” *Physical Review E* 80 (2009), 035101(R).
- [Lohsoonthorn 03] P. Lohsoonthorn. “Hyperbolic Geometry of Networks.” PhD thesis, Department of Electrical Engineering–Systems, University of Southern California,

2003. (Available online at http://eudoxus.usc.edu/IW/MATTFINALTHESIS_MAIN.pdf.)
- [Lou 08] M. Lou. “Traffic Pattern Analysis in Negatively Curved Networks.” PhD thesis, University of Southern California, 2008. (Available online at <http://eudoxus.usc.edu/IW/Mingji-PhD-Thesis.pdf>.)
- [Narayan and Saniee 09] O. Narayan and I. Saniee. “The Large Scale Curvature of Networks.” arXiv:0907.1478v1, 2009.
- [Réti et al. 05] T. Réti, E. Bitay, and Z. Kosztolányi. On the Polyhedral Graphs with Positive Curvature. *Acta Polytechnica Hungarica* 2:2 (2005), 19–37.
- [Sarkar et al. 09] R. Sarkar, X. Yin, J. Gao, F. Luo, and X. D. Gu. “Greedy Routing with Guaranteed Delivery Using Ricci Flows.” In *Proceedings of the 8th International Symposium on Information Processing in Sensor Networks (IPSN)*, April 2009.
- [Spring et al. 04] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. “Measuring ISP Topologies with Rocketfuel.” *IEEE Transactions on Networking* 12:1 (2004), 2–16.
- [Sun and Yu 04] Liang Sun and Xingxing Yu. “Positively Curved Cubic Plane Graphs Are Finite.” *Journal of Graph Theory* 47:4 (2004), 241–274.
- [Zhao et al. 05] L. Zhao, Y.-C. Lai, K. Park, and N. Ye. “Onset of Traffic Congestion in Complex Networks.” *Physical Review E* 71 (2005), 026125–1–026125–8.

Edmond Jonckheere, Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 (jonckhee@usc.edu)

Mingji Lou, Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089 (mingjilou@gmail.com)

Francis Bonahon, Department of Mathematics, University of Southern California, Los Angeles, CA 90089 (fbonahon@math.usc.edu)

Yuliy Baryshnikov, Room MH 2C-361, Bell Laboratories, 600 Mountain Ave, Murray Hill, NJ 07974-0636 ([ymb@research.bell-labs.com](mailto:yymb@research.bell-labs.com))

Received November 2, 2009; accepted April 26, 2010.