

Geometry of Network Security

Edmond Jonckheere and Poonsuk Lohsoonthorn

Abstract—The main point of this paper is that network security has a geometric component, in the sense that some architectures promote some aspects of security. Such security issues closely related to the topological architecture of the network graph are multi-path routing to mitigate “eavesdropping” or “packet sniffing,” worm propagation and defense, and Distributed Denial of Service (DDoS) attack mitigation. Those geometric aspects relevant to network security are encapsulated in the concept of graph curvature. An architecture that promotes, in some sense, security is the negative curvature of the graph, which is shown to hold in several physical and logical graphs and in the well know “scale free” model.

I. INTRODUCTION

The objective of this paper is to define small scale, medium scale, and large scale curvature concepts for graphs and to show their significance in such security issues as robust routing, distributed denial of service mitigation, epidemiology of worms, worm defense, etc. We further look at the curvature of some physical and logical graphs, as well as such model as the scale free model, and show their negative curvature by Monte Carlo simulations. This paper involves the relatively new field of coarse geometry [20], one aspect of which is that graphs and manifolds are treated as the same mathematical objects, and, as such, a graph can be given a curvature.

II. TOPOLOGICAL SURFACE EMBEDDING

For the concept of graph curvature to be more palatable, it is convenient to think a graph as embedded in a surface, so that if the embedding is isometric the curvature of the graph could be defined as that of the surface. Consider a graph G , which we write on the sphere S^2 , possibly with some edges crossing. (In fact, the only situation where no edge crossings occur is when the graph is planar.) For each edge crossing, “pull a handle” and draw one of the edges on the handle rather than on the sphere. After pulling a handle for every pair of crossing edges, the graph is written—without edge crossings—on a sphere with g handles, that is, the compact orientable surface S_g of genus g . The computational implementation of this process relies on the so-called *Hefter-Edmonds rotation system* [16, Sec. 3.2]: intuitively, a surface encompassing a graph induces a permutation π_v , or *rotation*, of the edges flowing out of each vertex v . Conversely, assume that we have a *rotation system* $\{\pi_v : v \in G^0\}$, that is, $\forall v \in G^0$, we are given a

permutation π_v of the edges flowing out of v . This yields an embedding $e : G \rightarrow S_g$, in which the boundaries $e(G)$ of the faces are given by the π walks [16, Sec. 3.2]: Start at a vertex v^0 and proceed along an outflowing edge e^0 until we reach another vertex v^1 from which we proceed along the edge $\pi_{v^1}(e^0)$, etc., until we come back to v^0 along an edge e such that $\pi_{v^0}(e) = e^0$; this closed path bounds a face and the faces are “glued” along the edges to yield S_g . This embedding process is of course nonunique, but among all such processes there is one that leads to a minimum number of handles, called the *genus of the graph* $g(G)$ (see [16, Sec. 3.4]). A fundamental result says that those rotation systems that yield the minimum genus embedding are cellular embeddings [16, Prop. 3.4.1], that is, $S_g \setminus e(G)$ breaks as the disjoint union of (*acyclic*) *cells*, i.e., homeomorphs of the open disk. The advantage of a *cellular* embeddings is that the traffic flow on the edges can be extended to the cells. Of course, this extension is nonunique, but has such invariants as the index of the flow field. We can then conceive the traffic as flowing on a continuous geometric structure.

III. SMALL SCALE, LOW DIMENSIONAL ISOMETRIC EMBEDDING

The above embedding is purely topological and disregards the metric structure of the graph. To check whether the graph can be *isometrically* embedded in the surface requires the comparison theory [6, Chap. 4]. The comparison theory was historically the first attempt at defining curvature of metric spaces and is based on the fact that *any* metric triangle can be isometrically embedded in *any* standard Riemannian manifold \mathbb{M}_κ^r of constant sectional curvature κ and dimension r . Given a geodesic triangle Δuvw in some geodesic metric space (X, d) , the *comparison triangle* [6, Def. 4.1.8], [5, Chap. II.1] in the standard constant curvature space $(\mathbb{M}_\kappa, \bar{d})$ is a triangle $\bar{u}\bar{v}\bar{w}$ such that $\bar{d}(\bar{u}, \bar{v}) = d(u, v)$, $\bar{d}(\bar{v}, \bar{w}) = d(v, w)$, $\bar{d}(\bar{w}, \bar{u}) = d(w, u)$. The metric space (X, d) to be a Cartan-Alexandrov-Toponogov [12, Sec. 3.2] or CAT($\kappa < 0$)-space, i.e., has curvature bounded from above by $\kappa < 0$, if $\forall z \in [uv]$, $\forall y \in [uw]$ along with their comparison points in \mathbb{M}_κ , we have $d(z, y) \leq \bar{d}(\bar{z}, \bar{y})$ (see [6, Sec. 4.1.4]).

The comparison triangle also allows a concept of angle to be defined *solely in terms of the distance, independently of the concept of inner product*. The *Alexandrov angle* [5, Def. 1.12], [6, 4.3] at the vertex u of a geodesic triangle Δuvw is the $\overline{\lim}_{\epsilon \rightarrow 0} (\underline{\lim}_{\epsilon \rightarrow 0})$ of the angle at the vertex \bar{u} of the comparison triangle $\bar{u}\bar{v}_\epsilon\bar{w}_\epsilon$ of $\Delta uv_\epsilon w_\epsilon$ in the standard negative (positive) curvature space \mathbb{M}_κ , where $v_\epsilon \in [uv]$, $w_\epsilon \in [uw]$, and $d(u, v_\epsilon) = \epsilon d(u, v)$, $d(u, w_\epsilon) = \epsilon d(u, w)$.

This work was supported by DARPA Contract N66001-00-C-8044.

E. Jonckheere is with the Department of Electrical Engineering-Systems, University of Southern California, Los Angeles, CA 90089-2563, USA jonckhee@usc.edu

P. Lohsoonthorn is with the Department of Electrical Engineering-Systems, University of Southern California, Los Angeles, CA 90089-2563, USA lohsoont@hotmail.com

The angle $\angle \bar{v}_\epsilon \bar{u} \bar{w}_\epsilon$ depends on the metric of \mathbb{M}_κ , but remarkably the limit as $\epsilon \downarrow 0$ does not depend on what comparison space is chosen [5, Prop. 2.9]. The Alexandrov angle is a *small scale* concept, and as such it does not directly apply to graphs. Indeed if u is an arbitrary vertex in a graph, the Alexandrov angle between the links $[uv]$ and $[uw]$ would be 180 deg. At such a small scale, the geometry of the graph is completely lost so that this result is not surprising. For a graph, the lowest possible scale that can possibly provide some geometric insight in the neighborhood of a vertex u is the scale of vertices directly linked to u . At that scale, a nontrivial Alexandrov angle can be defined.

Armed with the concept of Alexandroff angle, we now can come back to the local curvature of a graph G in a surface S_g . Assume that the graph has been topologically embedded using the rotation system $\{\pi_x : x \in G^0\}$. Consider the vertex x , choose an edge e incident upon x , and define vertices v^1, \dots, v^n such that $uv^1 = e$, $uv^2 = \pi_u(e)$, $uv^3 = \pi_u^2(e)$, $uv^n = \pi_u^{n-1}(e)$, and finally $\pi_u(uv^n) = \pi_u(uv^1)$. Observe that v^i, v^{i+1} need not be connected by an edge, but that there is a distance $d(v^i, v^{i+1})$ associated with them. Assume that the system of points u, v^1, \dots, v^n and distances $d(u, v^i), d(v^i, v^{i+1})$ is embeddable in \mathbb{E}^3 (see [4]). In this case, we can connect v^i, v^{i+1} with an edge of length $d(v^i, v^{i+1})$. The cone with A as apex and the polygonal line $v^1 v^2 \dots v^n$ as base is a singular surface and we strive to define the curvature at its apex. Skipping the details, it is not hard to see using the Gauss-Bonnet theorem that the curvature at the apex A is positive if $\sum_i \angle v^i u v^{i+1} < 2\pi$ while the same curvature is negative if $\sum_i \angle v^i u v^{i+1} > 2\pi$.

Now we come back to the isometric embedding question in the low scale limit $v^i \rightarrow u$. If $\sum_i \angle v^i u v^{i+1} = 2\pi$, then the subgraph with vertices u, v^i can be isometrically embedded into S_g . If, however, $\sum_i \angle v^i u v^{i+1} > 2\pi$, then the isometric embedding of the subgraph yields a *pleat* singularity at A and the surface is said to have a *singular hyperbolic metric* (see [12, Chap. 14]). If, on the other hand, $\sum_i \angle v^i u v^{i+1} < 2\pi$, then the isometric embedding of the subgraph must have a *conical singularity* at u (see [12, Sec. 61, 6.2]), for indeed the sum of the angles at the apex of a cone is $< 2\pi$.

We can now perceive the bigger picture: At the large scale the graph $G \subseteq S_{g>2}$ would be hyperbolic, with local patches of singular negative curvature and local patches of positive curvature.

IV. ISOMETRIC EMBEDDING IN 3 AND HIGHER DIMENSIONAL SPACES

A problem with the previous analysis, along with its definition of positively and negatively curved graphs, is that it clings on the assumption that the graph has been topologically embedded in a surface. Here we indicate the way to carry over this analysis to graphs embedded in 3-manifolds.

There is a dramatic transition as we go from 2-D to 3-D, because indeed, while a metric triangle, that is, a triple of points with their distances satisfying the triangle inequality, is embeddable in any space of any curvature, the basic 3-D building block, the metric tetrahedron, that is, a quadruple of points such that every of its triple satisfies the triangle inequality, is *not always* isometrically embeddable in an arbitrarily curved space (see Sec. V). In addition, the 2-D geometry fact that the sum of the internal angles of a triangle dictates the curvature is more complicated in 3-D. If $\alpha_{ij}, i \neq j$, is the dihedral angle of the edge $x^i x^j$, and if we agree that $\alpha_{ii} = \pi$, then what dictates the geometry is the Gram matrix $\Gamma = \{-\cos \alpha_{ij}\}$ (see [19], [15]). Precisely, the geometry is hyperbolic iff $\lambda_1(\Gamma) < 0 < \lambda_2(\Gamma) \leq \lambda_3(\Gamma) \leq \lambda_4(\Gamma)$ and all (i, j) cofactors are positive; the geometry is Euclidean iff $\lambda_1(\Gamma) = 0 < \lambda_2(\Gamma) \leq \lambda_3(\Gamma) \leq \lambda_4(\Gamma)$ with the same condition on the cofactors; finally, the geometry is spherical iff $\Gamma > 0$.

Now take a graph, let its complete n -subgraphs be the n -simplexes, $n \leq 3$, of a 3-D complex, and *assume* that this simplicial complex is the triangulation of a 3-manifold \mathbb{M}^3 . Then the (sectional) curvature (of the complex and hence the graph) can be defined as the function $k(x^i x^j) = 2\pi - \sum_{kl} \alpha_{ij}^{kl}$, where α_{ij}^{kl} is the dihedral angle around $x^i x^j$ in the tetrahedron $x^i x^j x^k x^l$ (see [14]). This definition extends trivially to $n > 3$. A confirmation of the validity of this definition is provided by the 3-manifold fact that, if $k(x^i x^j) > 0$, then \mathbb{M}^3 has a spherical metric [15]. A scalar curvature can also be defined as $S(x^i) = \sum_{jkl} \kappa(x^i x^j) \text{vol}(x^i x^j x^k x^l)$ (see [14]). The Yamabe flow problem asserts that the scalar curvature of a Riemannian manifold can be deformed to a constant one [25], [22], [1], [21]. However, this is not in general true for the *combinatorial* Yamabe problem [14], a fact that has unfortunate consequences for network congestion (see Sec. IV).

Using a simplified approach to local combinatorial curvature, Eckmann-Moses [7] showed that the curvature of the World Wide Web graph follows a power law distribution across negative curvature.

V. MEDIUM SCALE EMBEDDING IN CONSTANT CURVATURE SPACE

If we strip a weighted graph (G, w) from those links that are too costly to communicate along, the graph becomes uniquely characterized by its node set x^1, \dots, x^n and a distance $d(x^i, x^j)$ for every pairs of nodes. It can be shown [4, Th. 63.1] that the graph viewed as a metric space (G, d) can be irreducibly isometrically embedded in the standard constant curvature $\kappa > 0$ space of dimension r iff the diameter of the set of points does not exceed $\frac{\pi}{\sqrt{\kappa}}$ and the $n \times n$ matrix $\Gamma = \{\cos(d(x^i, x^j)\sqrt{\kappa})\}$ is positive semidefinite with eigenvalues

$$0 = \lambda_1 = \dots = \lambda_{n-r-1} < \lambda_{n-r} \leq \dots \leq \lambda_n$$

Likewise [4, Th. 106.1 and Cor.], the same graph is embeddable in the standard hyperbolic space of curvature $\kappa < 0$

if only if the $n \times n$ matrix $\Gamma = \{\cosh(d(x^i, x^j)\sqrt{-\kappa})\}$ has eigenvalues

$$\lambda_1 \leq \dots \leq \lambda_{r+1} < 0 = \lambda_{r+2} = \dots = \lambda_{n-1} < \lambda_n$$

Finally [4, Th. 42.3], the graph is embeddable in \mathbb{E}^r , $r \leq n-1$, iff the $(n+1) \times (n+1)$ Cayley-Menger matrix [4, Th. 41.1, 42.1],

$$\begin{pmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & 0 & d(x^1, x^2)^2 & \dots & d(x^1, x^n)^2 \\ 1 & d(x^2, x^1)^2 & 0 & \dots & d(x^2, x^n)^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & d(x_n, x^1)^2 & d(x^n, x^2)^2 & \dots & 0 \end{pmatrix}$$

has eigenvalues

$$\lambda_1 \leq \dots \leq \lambda_{r+2} < 0 = \lambda_{r+3} = \dots = \lambda_{n-1} < \lambda_n$$

Using the above eigenvalue reformulation, it is possible to check embeddability of basic graph structures in standard spaces [10]. The complete graph K_n with uniform link weights is embeddable in \mathbb{E}^{n-1} , in $\mathbb{M}_{\kappa < 0}^{n-1}$, and in $\mathbb{M}_{\kappa > 0}^{n-1}$ provided

$$\kappa \leq \left(\frac{\cos^{-1}\left(-\frac{1}{n-1}\right)}{d(x^i, x^j)} \right)^2$$

Furthermore, the same complete graph K_n is embeddable in $\mathbb{M}_{\kappa > 0}^{n-2}$ if the above holds with *equality*. This leaves us in a quandary as to what geometry K_n has. Since K_n is bounded while $\mathbb{M}_{\kappa < 0}$ is not, the complete graph has spherical geometry. A star structure with uniform link weight is embeddable in a sufficiently negatively curved manifold. To check embeddability of such a core concentric structure as the ISP graph, we construct the model shown in Fig 1: We start with the complete graph K_n with unit link weight and, to each vertex, we attach a long edge of length ℓ . We can be shown that this model is embeddable in $\mathbb{M}_{\kappa < 0}$ provided $\ell\sqrt{-\kappa}$ is beyond a threshold. This is already an example as to how large scale geometric issues can be dealt with in this setting. In fact, large scale geometry appears more vividly if we attempt to embed a tree in constant curvature space. Consider a tree with uniform degree and link weight w . Embeddability in a *finite* curvature space is not possible; however, if we let $w\sqrt{-\kappa} \rightarrow \infty$, and extract the dominant part of the matrix $\{\cosh(d(x^i, x^j)\sqrt{-\kappa})\}$, then the embeddability condition is satisfied. This means that the tree is embeddable in $\mathbb{M}_{\kappa < 0}$ from the large scale viewpoint or that the tree is embeddable in $\mathbb{M}_{-\infty}$.

VI. LARGE SCALE GEOMETRY

We now proceed to large scale curvature by looking at theoretically infinite diameter graphs. Because positive curvature manifolds have their diameter bounded, it follows that at large scale, the only relevant curvature concept is nonpositive. To define such a concept, it is instructive to go back to Riemannian geometry. Let \mathbb{M} be a Riemannian




Fig. 1. Left: The Internet Service Provider (ISP) graph: Each node is an ISP; two ISP's are linked if one is the next-hop of the other in the Border Gateway Protocol (BGP); each link is assigned a weight equal to the number of paths observed between the two ISP's by a traceroute-like routine; the weight of a link is color coded; light gray represents low weight while dark black represents high weight; high degree nodes are at the center while low degree nodes are at the periphery; observe that the geodesic triangle $\triangle ABC$ has low δ_F compared with its diameter. Right: The idealized model to show embeddability in hyperbolic space.

manifold of sectional curvature $\kappa < 0$. The *fatness* of a geodesic triangle $\triangle uvw$ is defined as

$$\delta_F(\triangle uvw) := \inf \left\{ d(x, y) + d(y, z) + d(z, x) : \begin{array}{l} x \in [vw] \\ y \in [uw] \\ z \in [uv] \end{array} \right\} \quad (1)$$

and the bounded fatness property [20, pp. 84-85] is that

$$\delta_F(\mathbb{M}) := \sup\{\delta_F(\triangle uvw) : u, v, w \in \mathbb{M}\} < \frac{6}{\sqrt{-\kappa}} \quad (2)$$

Clearly, because the above definition of $\delta_F(\cdot)$ relies only on a metric structure, it can be extended to geodesic metric spaces, in particular to graphs. Therefore, we will say that a graph G is δ_F -hyperbolic if $\delta_F(G) < \infty$. By this definition, all finite diameter graphs are hyperbolic, so that an amended definition relevant to *very large, yet finite* graphs is warranted. It can be shown that, for $\epsilon > 0$,

$$\sup_{\substack{\triangle \subseteq \mathbb{M}_{\kappa < 0}^r \\ \text{diam}(\triangle) \geq \epsilon}} \frac{\delta_F(\triangle)}{\text{diam}(\triangle)} < \sup_{\substack{\triangle \subseteq \mathbb{E}^r \\ \text{diam}(\triangle) \geq \epsilon}} \frac{\delta_F(\triangle)}{\text{diam}(\triangle)} < \frac{3}{2}$$

and the inequalities becomes equalities as $\epsilon \downarrow 0$. (The first inequality is proved via the comparison theory of Sec. III and the second is proved from the fact that $\delta_F(\triangle) = \frac{2a^2b^2 + 2a^2c^2 + 2b^2c^2 - a^4 - b^4 - c^4}{2abc}$ for a triangle with edge lengths a, b, c along with some elimination method.) It follows that a necessary condition for a finite graph to be a $\text{CAT}(\kappa < 0)$ space is that the maximum of δ_F/diam be less than $3/2$. It can be shown [10] that this maximum occurs on the vertices.

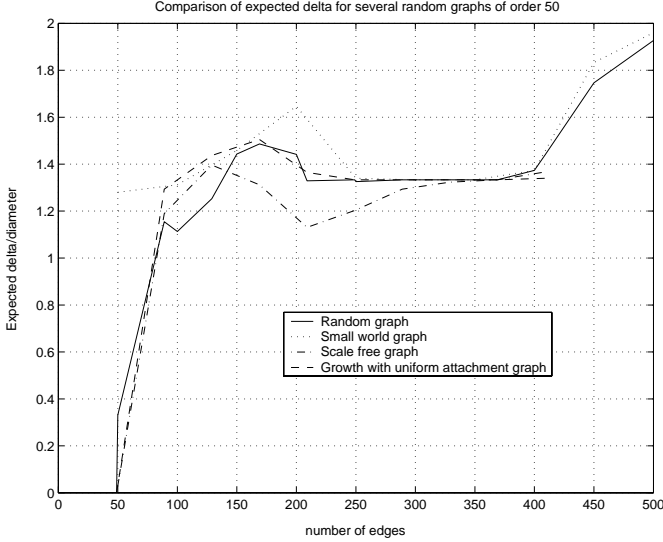


Fig. 2. The mathematical expectation of $\max \delta_F/\text{diam}$ versus the total number of edges for all 4 graph generators. Observe that the scale free graph is the most hyperbolic.

Another computationally more efficient but less intuitive measure is based on the fact that the distances $d(x^i, x^j)$ among four points, x^1, x^2, x^3, x^4 , carry curvature information. Let us order the vertices so that, if we define $L(x^i) = d(x^1, x^4) + d(x^2, x^3)$, $M(x^i) = d(x^1, x^2) + d(x^3, x^4)$, and $S(x^i) = d(x^1, x^3) + d(x^2, x^4)$, we have $S \leq M \leq L$. Define

$$\delta_G(X) := \sup \left\{ \frac{L(x^1, x^2, x^3, x^4) - M(x^1, x^2, x^3, x^4)}{2} : x^i \in X \right\}$$

Then it can be shown that X is δ_F -hyperbolic iff $\delta_G(X) < \infty$. For a finite graph, it can be shown that the real issue is $\delta_G(u, v, w, x)/\text{diam}(u, v, w, x) < (\sqrt{2} - 1)/\sqrt{2}$, where u, v, w, x are vertices sufficiently spaced.

VII. MONTE CARLO SIMULATIONS

To bring extra evidence of our claim, also formulated independently by Baryshnikov [3], that the Internet is δ -hyperbolic, we analyze the curvature properties of the various graph models of Internet build-up. In the early Erdős-Rényi $\mathcal{R}(n, M)$ model, a random graph was characterized by n nodes and M links distributed uniformly at random. Another model is the small world $\mathcal{W}(n, d, \beta)$ model of Watts and Strogatz [24]. In this model, we start with n nodes arranged in a lattice, that is, every node has the same degree d , and then every link is rewired with probability β . The $\mathcal{R}(n, M)$ and $\mathcal{W}(n, d, \beta)$ models were, however, not able to reproduce the typical heavy tail feature of network graphs, and for this reason the scale free model, $\mathcal{F}(n, m)$, also referred to as growth/preferential attachment model [2], was introduced. In this model, the network starts up with a core of n_0 nodes, linked as a random spanning tree. Then new nodes v^1, \dots, v^n are recursively brought and




Fig. 3. The mathematical expectation of $\max \delta_G/\text{diam}$ versus the total number of edges M for all 4 graph generators. Observe that the scale free graph is the most hyperbolic.

each new node has m links attached to the previous nodes with a probability proportional to the degree, referred to as *preferential attachment* $\mathcal{F}(n, m)$. A slight variant is the uniform attachment probability model $\mathcal{U}(n, m)$. (See [11] for a comparison.) It can be shown [2] that the resulting preferential attachment graph $\mathcal{F}(n, m)$ has heavy tail degree distribution as $n \rightarrow \infty$ and that the uniform attachment graph does not.

We took $n_0 = 10$ and constructed the graphs all the way through $n = 50$ nodes, for all 4 graph generators. To draw a fair comparison among the 4 generators, we took the total number of links M as parameter and we plotted

$$E \left(\max_{u, v, w, x \in G^0} \frac{\delta_F(\Delta uvw)}{\text{diam}(\Delta uvw)} \right), E \left(\max_{u, v, w, x \in G^0} \frac{\delta_G(u, v, w, x)}{\text{diam}(u, v, w, x)} \right)$$

versus M for all 4 generators, as shown in Figs 2, 3, respectively. Observe that the two curves are quite consistent in their overall shape.

The curves can be explained as follows: The backbone graph is a random tree visiting all nodes; when links are being added, the δ initially increases rapidly because the tree is being “fattened,” while the diameter remains roughly constant; then the δ decreases because the new links start creating some shortcuts, while the diameter still remains constant; then we reach the most interesting region, half-way between the minimum and the maximum M , because the geometry is hyperbolic without trivial tree structure; after that, there are too many links, which has the effect of decreasing the diameter while the δ remains constant, hence the curve goes up, to approach the complete graph situation $\delta_F(K_{50})/\text{diam}(K_{50}) = 2$. Another important observation is that the scale free model is the most hyperbolic.

VIII. APPLICATIONS

A. Multipath Routing

One of the many security concerns in modern data networks is eavesdropping, that is, unauthorized packet interception along a link with the potential of reconstructing the full message—if all packets are sent along the same optimum path from source to destination, as TCP does under normal conditions. One of the proposed patches to such a security breach is to send packets in a randomized fashion along different equal-cost [9], [23] or nonoptimal routes [8], [11]. The problem when packets travel along nonoptimal routes is that they incur different delays, arrive out-of-order at the destination, and TCP is very limited in its ability to reorder the packet to their original sequence. In fact, most of the time, TCP will simply drop the out-of-order packets. A robustified version [8], referred to as TCP-MP, where MP stands for Multiple-Path, has been developed and, if implemented in the protocol suite of the source and destination, is able to reconstruct the packet sequence to its original order. However, the worse the out-of-order, the more the overhead, so that it is still desirable to maintain the delay discrepancy under control by sending packets along near optimal routes, subject to a multiplicative tolerance $\lambda \geq 1$ and formalized under λ -quasi-geodesics [5, Def. I.8.22]. A λ -quasi-geodesic is, by definition, a λ -quasi-isometric embedding, that is, a map $\tilde{\gamma} : [a, b] \rightarrow X$ subject to

$$\frac{1}{\lambda}|s - s'| \leq d(\tilde{\gamma}(s), \tilde{\gamma}(s')) \leq \lambda|s - s'|$$

In [8], the near optimal paths were computed by a “brute force” search. The remarkable property of hyperbolic graphs is that if $\tilde{\gamma}$ is a λ -quasi-geodesic with the same end points as the geodesic γ , the Hausdorff distance $d_H(\gamma, \tilde{\gamma})$ is bounded by a quantity depending only on (λ, δ) (see [5, Th. III.H.1.7]).

Even though the quasi-optimal paths can be sought, with reasonable efficiency, within this neighborhood of the optimum path, we prefer, however, to use yet another property of hyperbolic space—mainly that the quasi-geodesics can be computed as k -local geodesics, that is, locally optimum paths over k hops [5, Th. III.H.1.13]. A quasi-optimal route discovery protocol can be devised based on this property [13].

B. DDoS Attack Mitigation

Core concentric networks, which we would rename as hyperbolic networks, have also been proposed as an architecture that mitigates Distributed Denial of Service (DDoS) attacks. Indeed, the attack of the “zombies” to a destination would have to transit through some “cores” where the actions of the various “zombies” would destroy each other.

C. Positive Curvature Graphs

The problem with a hyperbolic space is that the good behavior of the geodesics might be detrimental in Information Warfare (IW). Indeed, the quasi-geodesics, along

which we could reroute, are guaranteed to be close to the geodesic—probably too close by security standards. Indeed, to fool the enemy attempting “eavesdropping” or “packet sniffing,” it might be desirable to have the quasi-geodesics geographically far away from the geodesic that the enemy is wiretapping. In other words, sometimes, it is desirable to have a wild behavior of the quasi-geodesics. This bad behavior of the quasi-geodesics typically occurs in a positively curved or spherical space.

D. Worm Propagation and Defense

While there are many worm propagation schemes, they can be unified under the theme that a worm propagates on a graph, from one infected node to the neighboring nodes. The propagation graph depends on the nature of the worm and as such could have varying topological, curvature, and random properties. For example, a mail worm propagates on the logical mail graph. While the mail graph appears to fall short of heavy tail, it has the growth, sublinear preferential attachment model [18]. In general, a worm exploits some software vulnerability and does a random scanning of the 32-bit address space every clock tick T . The random scanning is in general programmed using a linear congruence, well known to generate a uniform probability distribution, in which case the worm can be thought of as propagating on the Erdős-Rényi $\mathcal{R}(n, p)$ random graph. However, in the Sapphire/Slammer worm, the hacker miscoded the linear congruence [17] with the consequence that the “congruence” was cycling over some subset of IP addresses itself triggered by a “seed.” If it is agreed that the probability of choosing a particular cycle is directly proportional to the size of the cycle if the initial seed is selected uniformly at random, it follows that the probability distribution of the IP addresses is not uniform. The propagation graph hence appears to be of the core-concentric type, with the long cycles in the core and the small cycles at the periphery.

In the absence of defense strategy to contain a worm, simulation reveal a propagation pattern, that is, a plot of the number of infected machines versus time, that is heavily graph dependent. This is another indication of the crucial role played by the curvature of the graph. An easy mail worm defense scheme would be to shut down the server when two, or three messages from the same sender are received at close interval. If we implement this scheme, Monte Carlo simulations indicate that the speed of propagation of the worm definitely slows down, but that the overall shape of the curve remains the same. The invariance of the shape of the curve under worm defense strategy is a fundamental limitation as to what the elementary defense strategy can do. This is reminiscent of the fundamental limitations on the achievable feedback performance, well known in linear control theory. The explanation of this limitation is as follows: Assume the graph is infinite and hyperbolic. A finite worm defense strategy implemented by finitely many shutdowns of mail servers would be no more

than a quasi-isometry on the graph, and it is well-known in coarse geometry that a quasi-isometry is unable to destroy the δ -hyperbolic property of the graph [5, Th. III.H.1.9], hence the limitation as to what the defense strategy could accomplish.

E. Congestion Interpretation of Curvature

Since many worms do not contain malicious code, but disable the network by congesting it, it is important to understand what geometric network properties make it prone to congestion.

Consider a graph specified by vertices u, v^1, \dots, v^5 and distances $d(u, v^i) = 1, d(v^i, v^{i+1}) = 1$ where the other distances $d(v^i, v^j), i \neq j$ are derived from the previous ones. From the previous analysis, the local curvature at the vertex u is positive. Assume a message has to transit from v^i to v^j . Since the base is a pentagon, the communication cost between v^i and v^j is at most two, so that most of the traffic will transit around the base rather than through the apex. Hence, *positive local curvature at u means light traffic through u .*

Consider now a similar situation except that the base v^1, \dots, v^n has at least seven ($n \geq 7$) vertices, which by the previous analysis makes the local curvature at u is negative. In this case, even when v^i, v^j are nearly diametrically opposed on the base, their communication cost cannot exceed 2, if the traffic transits through the apex. Therefore, *strong negative curvature means congestion.* It appears therefore that it might be a good idea to keep the curvature constant; however, the combinatorial Yamabe problem reveals that not even the scalar curvature can be made constant.

IX. CONCLUSION AND FURTHER PROSPECTS

Beyond security, it is claimed that many graph theoretic issues are encapsulated in that single concept—graph curvature. The “core-concentric” wired networks are on the negative curvature side, whereas the wireless “meshed” networks are rather on the positive curvature side. With the advent of adaptive networking and the control plane, it is desirable to have stable geodesics, which is achieved with a hyperbolic network. If, on the other hand, we want to use multipath routing to the extreme by having wildly varying quasi-geodesics, this can be achieved on a positively curved network, but at the expense that the geodesics will be unstable as well, hence compounding the fluttering problem. Clearly, the curvature reflects the various trade-offs.

X. ACKNOWLEDGMENTS

Many thanks to Dr. N. Badr for many insightful discussions on the issue of meshed versus core-concentric networks.

REFERENCES

[1] T. Aubin. Equations différentielles non linéaires et problème de Yamabe concernant la courbure scalaire. *J. Math. Pures Appl.*, 55(9)(3):269–296, 1976.

[2] Albert-Laszlo Barabasi, Reka Albert, and Hawoong Jeong. Mean-field theory for scale free random networks. *Physica A*, 272:173–187, 1999. published by Elsevier.

[3] Y. Baryshnikov. On the curvature of the internet. In *Workshop on Stochastic Geometry and Teletraffic*, Eindhoven, The Netherlands, April 2002.

[4] L. M. Blumenthal. *Theory and Applications of Distance Geometry*. Oxford at the Clarendon Press, London, 1953.

[5] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *A Series of Comprehensive Surveys in Mathematics*. Springer, New York, NY, 1999.

[6] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*, volume 33 of *Graduate Study in Mathematics*. American Mathematical Society, Providence, Rhode Island, 2001.

[7] J.-P. Eckmann and E. Moses. Curvature of co-links uncovers hidden thematic layers in the world wide web. *PNAS*, 99(9), April 2002.

[8] Joo Pedro Hespanha and Stephan Bohacek. Preliminary results in routing games. In *Proc. of the 2001 Amer. Contr. Conf.*, June 2001.

[9] C. Hoppps. Analysis of an equal-cost multi-path algorithm. *Request for Comments (RFC) 2992*, November 2000.

[10] E. Jonckheere and P. Lohsoonthorn. *Geometric Topology of Networks*. Springer?, 200?. book project under development; copy available upon request.

[11] 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. FA5-1.

[12] Michael Kapovich. *Hyperbolic Manifolds and Discrete groups*, volume 183 of *Progress in Mathematics*. Birkhauser, Boston, MA, 2001.

[13] P. Lohsoonthorn. *Hyperbolic Geometry of Networks*. PhD thesis, University of Southern California, 2003.

[14] Feng Luo. Combinatorial Yamabe flow on surfaces. arXiv:math.GT/0306167 v1, 10 Jun 2003.

[15] Feng Luo. On a problem of Fenchel. *Geometriae Dedicata*, 64:277–282, 1997.

[16] B. Mohar and C. Thomassen. *Graphs on Surfaces*. The Johns Hopkins University Press, Baltimore, London, 2001.

[17] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. The spread of the sapphire/slammer worm. <http://www.cs.berkeley.edu/~nweaver/sapphire/>, 2003. Cooperative Association for Internet Data Analysis and University of California, San Diego.

[18] M. E. J. Newman, S. Forrest, and J. Balthrop. Email network and the spread of computer viruses. *Physical Review E*, 66:035101–1–4, 2002.

[19] I. Rivin. Some observations on the simplex. arXiv:math.MG/0308239v1, 2003.

[20] John Roe. *Index Theory, Coarse Geometry, and Topology of Manifolds*. Number 90 in Conference Board of the Mathematical Sciences (CBMS); Regional Conference Series in Mathematics. American Mathematical Society, Providence, RI, 1996.

[21] R. Schoen. Conformal deformation of a Riemannian metric to constant scalar curvature. *J. Differential Geometry*, 20:479–495, 1984.

[22] N. S. Trudinger. Remarks concerning the conformal deformation of Riemannian structures on compact manifolds. *Ann. Scuola Norm. Sup. Pisa*, 22(2):265–274, 1968.

[23] C. Villamizar. Ospf optimized multipath (ospf-omp). *draft-ietf-ospf-omp-03*, page 46, June 1999.

[24] D. J. Watts and S. H. Strogatz. Collective dynamics of small world networks. *Nature*, 393:440–442, 1998.

[25] H. Yamabe. On a deformation of Riemannian structures on compact manifolds. *Osaka Math. J.*, 12:21–37, 1960.