Theoretical Computer Science 639 (2016) 72–90

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Data center interconnection networks are not hyperbolic ☆

David Coudert [a,b], Guillaume Ducoffe [b,a,*]

[a] Inria, France
[b] Univ. Nice Sophia Antipolis, CNRS, I3S, UMR 7271, 06900 Sophia Antipolis, France

ARTICLE INFO

Article history:
Received 7 June 2015
Received in revised form 17 May 2016
Accepted 18 May 2016
Available online 24 May 2016
Communicated by R. Klasing

Keywords:
Greedy routing scheme
Metric embedding
Graph endomorphism
Gromov hyperbolicity
Cayley graph
Data center interconnection network

ABSTRACT

Topologies for data center interconnection networks have been proposed in the literature through various graph classes and operations. A common trait to most existing designs is that they enhance the symmetric properties of the underlying graphs. Indeed, symmetry is a desirable property for interconnection networks because it minimizes congestion problems and it allows each entity to run the same routing protocol. However, despite sharing similarities these topologies all come with their own routing protocol. Recently, generic routing schemes have been introduced which can be implemented for any interconnection network. The performances of such universal routing schemes are intimately related to the *hyperbolicity* of the topology. Roughly, graph hyperbolicity is a metric parameter which measures how close is the shortest-path metric of a graph from a tree metric (the smaller the gap the better). Motivated by the good performances in practice of these new routing schemes, we propose the first general study of the hyperbolicity of data center interconnection networks. Our findings are disappointingly negative: we prove that the hyperbolicity of most data center interconnection topologies scales linearly with their diameter, that is the worst-case possible for hyperbolicity. To obtain these results, we introduce original connection between hyperbolicity and the properties of the endomorphism monoid of a graph. In particular, our results extend to all vertex and edge-transitive graphs. Additional results are obtained for de Bruijn and Kautz graphs, grid-like graphs and networks from the so-called Cayley model.

1. Introduction

The network topologies that are used to interconnect the computing units of large-scale facilities (e.g., super computers, data centers hosting cloud applications, etc.) are designed to optimize various constraints such as equipment cost, deployment time, capacity and bandwidth, routing functionalities, reliability to equipment failures, power consumption, etc. This large variety of (conflicting) criteria has yielded numerous proposals of interconnection networks. See for instance [1–10] for the most recent ones. A common feature of the proposed constructions is to design network topologies offering a high-level of *symmetries*. Indeed, it is easier to balance the traffic load, and hence to minimize the congestion, on network topologies with a high-level of symmetry. Furthermore, it simplifies the initial wiring of the physical infrastructure and it ensures that each router node can run the protocol.

☆ This work is partially supported by ANR project Stint under reference ANR-13-BS02-0007 and ANR program "Investments for the Future" under reference ANR-11-LABX-0031-01.

* Corresponding author.
E-mail address: guillaume.ducoffe@inria.fr (G. Ducoffe).

http://dx.doi.org/10.1016/j.tcs.2016.05.025

However, despite sharing properties, interconnection networks rely on specific routing algorithms that are optimized for each topology. As a novel step toward efficient and topology agnostic routing schemes, the authors in [11–13] proposed to use greedy routing schemes based on an embedding of the topology into certain metric space such as the hyperbolic metric space, and more recently the word metric space. This approach has been shown particularly efficient for Internet-like graphs [14,15] where routes with low stretch are obtained. One explanation of this good behavior is that Internet-like graphs have low *hyperbolicity* [16,17], a graph parameter providing sharp bounds on the stretch (or distortion) of the distances in a graph when it is embedded into an edge-weighted tree.

In this paper, we characterize or give upper and lower-bounds on the hyperbolicity of a broad range of interconnection network topologies. These bounds can be used to analyze the worst-case behavior of greedy routing schemes in these topologies. Before we present our results, let us further put in context the role they play in routing and in other distance-related problems.

Related work. Greedy routing schemes based on an embedding into the hyperbolic space have been introduced by Kleinberg in [14]. Since then, various authors explored further this approach [15,18,19]). In particular, they showed that the graphs of the Autonomous Systems of the Internet embed better into a hyperbolic space than into an Euclidean space.[1] It was only recently in [21] that a formal relationship between the performances of hyperbolic embeddings and the hyperbolicity was proved. Namely, the authors proved that the over-delay for such routing schemes, or equivalently the stretch of the routing, depends on the hyperbolicity. In [22], the authors proved that similar results hold for greedy routing schemes based on an embedding of the topology into some word metric space (*e.g.*, see [23] for more information). More precisely, they use hyperbolicity to upper-bound the complexity of their routings, as well as to bound the size of the automata that are involved in their routing schemes.

Their results add up to prior worst-case analysis of graph heuristics that already pointed out the important role played by the hyperbolicity. For instance, there are approximation algorithms for problems related to distances in graphs — like diameter and radius computation [24], and minimum ball covering [25] — whose approximation constant depends on the hyperbolicity. Sometimes the approximation factor is a universal constant but the algorithm relies on a data-structure whose size is proportional to the hyperbolicity of the network topology [26]. Geometric routing schemes in [15,18,19] do not make exception and so have a stretch *lower-bounded* by the hyperbolicity (the bound is reached by some of them).

There have been measurements to confirm that complex networks such as the graphs of the Autonomous Systems of the Internet, social networks and phylogenetic networks all have a low hyperbolicity. We refer to [16,27–31] for the most important studies in this area. Additional related work in [32,33] shows that the low hyperbolicity of complex networks may be a consequence of some preferential attachment mechanisms. However, we are not informed of any study on the hyperbolicity of data center interconnection networks. In this paper, we aim to fill in this gap through a theoretical study of their underlying graphs.

Our contributions. In an attempt to confront with the diversity of interconnection network topologies proposed in the literature, we relate hyperbolicity with a few graph properties that are frequently encountered in these topologies. Indeed, we do not aim to provide a — long and non-exhaustive — listing of unrelated results for each network, but rather to exhibit a small number of their characteristics that are strongly related with their metric invariants. In particular, we relate hyperbolicity with the symmetries of a graph.

- We prove in Section 3 that for graphs whose center is a k-distance dominating set for some small value of k, the hyperbolicity scales linearly with the diameter. This class of graphs strictly contains graphs whose diameter equals the radius, *a.k.a.* the *self-centered graphs* [34,35]. In particular, it comprises all vertex-transitive graphs (a strict subclass of self-centered graphs), as well as edge-transitive graphs. A main consequence of our result is that every interconnection network whose topology is based on a *Cayley graph* has large hyperbolicity.[2]
- In addition, we prove that similar results hold for graphs admitting an *endomorphism* such that the distance between any vertex and its symmetric image is large. On the way to prove these results, we define a new graph invariant which is called *weak mobility*, that generalizes the so-called graph mobility (*e.g.*, see [37,38]). We use these new results to improve our lower-bounds on the hyperbolicity of several interconnection networks.
- For completeness, we also characterize the hyperbolicity of other "symmetric" networks such as de Bruijn, Kautz and grid-like graphs. More precisely, we apply different techniques that are based on their shortest-paths distribution so that we can prove in Section 4 that they also have a large hyperbolicity. The techniques that are involved in the proofs have been introduced in previous papers [39,40], but to the best of our knowledge the way we use them in this work is new.

[1] In fact, it follows from [20] that for any n-vertex graph G there is an embedding φ of G into the Euclidean space (with unbounded dimension) such that for every $u, v \in V(G)$ we have $d(\varphi(u), \varphi(v)) \leq \mathcal{O}(\sqrt{\log\log n}) \cdot d_G(u, v) + \hat{\mathcal{O}}(\delta(G) \cdot \log n)$, with $\delta(G)$ being the hyperbolicity (the $\hat{\mathcal{O}}$-notation suppresses the polyloglog factors). However, it does not seem that hyperbolicity is the most relevant parameter in the study of Euclidean embeddings.

[2] Independently from this work, the authors in [36] proved that for any vertex-transitive graph, the hyperbolicity scales linearly with the diameter. However, their proof relies on another definition of hyperbolicity, and it is unclear whether the proof can be extended to other graph classes. By contrast, our proof yields a tighter lower-bound for hyperbolicity, and it relies on a much simpler and more general argument (*i.e.*, see Theorem 4). Especially, it also applies to edge-transitive graphs.