

Randomized Compact Routing in Decomposable Metrics

Goran Konjevod ^{*} Andréa W. Richa [†] Donglin Xia [‡] Ling Zhou [‡]

Abstract

We study the compact routing problem in networks whose shortest path metrics are decomposable. Decomposable metrics are more general than doubling metrics, growth-bounded metrics, and metrics induced by graphs excluding $K_{r,r}$ as a minor. In this work, we present both name-dependent and name-independent constant stretch compact routing schemes for bounded decomposable metrics with polylogarithmic storage requirements at each node and polylogarithmic packet headers. Our work is the first to design compact routing schemes with constant stretch for networks as general as decomposable metrics.

KEY WORDS

Distributed computing, compact routing, decomposable metrics, randomized algorithms.

1 Introduction

Decomposable metrics have been used as a fundamental tool in the design of randomized algorithms on metric spaces in recent years. An r -bounded δ -padded decomposition of a metric (X, d) is a distribution over X 's partitions such that (i) for each random-generated partition, all of its clusters have radius no more than r ; and (ii) the probability that the distance from any node $x \in X$ to the “boundary” of the partition is no less than δr is no less than $1/2$. We say a metric (X, d) is α -decomposable, if for all $r > 0$ there exists an r -bounded $1/\alpha$ -padded decomposition.

A routing scheme is a distributed algorithm that deliver packets from any source to any destination node. Making the routing storage requirement scalable with the network size while minimizing the routing stretch is one of the fundamental trade-offs for routing scheme design, where the *stretch* is the maximum ratio of the routing path length to the shortest path length over all source destination pairs. We say a routing scheme is compact if its storage requirement at each node and for each packet header is polylogarithmic on the number of nodes. There are two variants of routing scheme design: (i) name-dependent (or labeled) routing and (ii) name-independent routing. Labeled routing allows the scheme designer to label the nodes with additional routing information. In name-independent routing, the scheme must use solely the (arbitrary) original naming.

In this paper, we present both labeled and name-independent compact routing with polylogarithmic storage and constant stretch for networks with bounded decomposable metrics. Compact routing has been studied in various networks, such as growth-bounded networks [3], power-law networks [4], and doubling metrics [1, 9]. It is known [7, 10] that a bound on the doubling dimension implies a bound on the padded decomposability. In [8, 11, 5], it is shown that every metric space X induced by an edge-weighted graph excluding $K_{r,r}$ as a minor is $O(r^2)$ -decomposable. To the best of our knowledge, our work is the to provide randomized schemes for compact routing with constant stretch in networks as general as bounded decomposable metrics.

¹konjevod1@llnl.gov, Lawrence Livermore National Laboratory, CA, USA

²aricha@asu.edu, Arizona State University, Tempe, AZ, USA

³{doxia,lizho}@microsoft.com, Microsoft, Redmond, WA, USA

2 Randomized Compact Routing Schemes

Given an edge-weighted n -node Δ -diameter graph $G = (V, E)$ whose shortest path metric is α -decomposable, we present (i) a 4α -stretch labeled routing scheme w.h.p.¹ with $O(\log \Delta \log^3 n / \log \log n)$ -bit routing labels, $O(\log \Delta \log^3 n / \log \log n)$ -bit routing tables and $O(\log^2 n / \log \log n)$ -bit packet headers; and (ii) a 22α -stretch name-independent routing scheme w.h.p. with $O(\log \Delta \log^5 n / (\log \log n)^2)$ -bit routing labels, and $o(\log^2 n)$ -bit packet headers.

Since G is α -decomposable, there are a bundle of 2^i -bounded $1/\alpha$ -padded decompositions, $\forall i \in [\log(\alpha\Delta)]$ ². For each $i \in [\log \alpha\Delta]$, we randomly pick up $\log n$ partitions from the 2^i -bounded $1/\alpha$ -padded decomposition. The key observation is that, for any source-destination pair (u, v) with $d(u, v) \leq 2^i/\alpha$, the probability that u and v are contained in the same cluster for one of the $\log n$ random 2^i -bounded partitions is no less than $1 - (1/2)^{\log n} = 1 - 1/n$. Thus with the help of local compact routing schemes at each cluster at each level, searching for the destination from lower to higher levels $i \in [\log \alpha\Delta]$ could result in constant stretch (depending on α) with high probability. In the following, we build upon this idea in our labeled and name-independent routing schemes.

2.1 Labeled Routing Scheme

Each cluster of all the $\log n$ random 2^i -bounded partitions, for each $i \in [\log \alpha\Delta]$, maintains a labeled routing scheme as in the following lemma on the shortest path tree of the cluster with its center as the root.

Lemma 2.1 ([6]) *For every weighted tree T on n nodes, there exists a labeled routing scheme that, given any destination label, routes optimally on T from any source to the destination. The storage per node, the label size, and header size are $O(\log^2 n / \log \log n)$ bits.*

The label of a node $x \in V$ consists of the local routing labels given by Lemma 2.1 for all clusters containing x over all random partitions and all levels. Thus the size of a label is $O(\log \Delta \log^3 n / \log \log n)$ bits. Similarly, the total size of the routing table at each node is also $O(\log \Delta \log^3 n / \log \log n)$ bits.

The routing algorithm from u to v , $\forall u, v \in V$, is described as follows. Node u identifies the minimal $i \in [\log(\alpha\Delta)]$ such that u and v have the local labels from the same cluster of one of the $\log n$ random partitions P at level i . Then u uses the local tree labeled routing scheme on P to route to v by the local label of v .

Stretch Analysis W.l.o.g.³, assume that $2^{i-1}/\alpha < d(u, v) \leq 2^i/\alpha$. Thus with probability $(1 - 1/n)$, nodes u and v are in the same cluster of one of the $\log n$ partition, i.e. node u could route the message to v with cost 2^{i+1} . Therefore, by $2^{i-1}/\alpha < d(u, v)$, we have that the routing stretch is at most 4α w.h.p..

3 Name-independent Routing Scheme

Each cluster of all the $\log n$ random 2^i -bounded partitions, for each $i \in [\log \alpha\Delta]$, maintains a name-independent tree routing scheme as in the following lemma on the shortest path tree of the cluster with its center as the root.

Lemma 3.1 ([2]) *Every weighted rooted tree with n nodes has a single-source name-independent routing scheme (in the designer port model) such that the distance traveled between the root r*

¹With high probability, i.e., with probability at least $1 - 1/n^c$, where $c > 0$ is a constant.

²For any integer $x > 0$, let $[x]$ denote the set $\{0, 1, \dots, x - 1\}$.

³Without loss of generality.

and a destination v is at most $d(r, v) + 2d(T)$, where $d(T)$ is the height of the tree, and the error-report to the root costs at most $4d(T)$ if the destination is not in the tree. Moreover, only $O(\log^4 n / (\log \log n)^2)$ bits are needed per node and headers have size $o(\log^2 n)$.

Since for each $i \in [\log(\alpha\Delta)]$ we pick up $\log n$ partitions, the routing table at each node has size $O(\log \Delta \log^5 n / (\log \log n)^2)$ bits.

We now present the name-independent routing algorithm from any source node u to any destination v . For i from 0 to $\log(\alpha\Delta)$, node u repeatedly searches for the destination v using the single source name-independent routing algorithm at the cluster containing u of one of the $\log n$ partitions at level i such that u is $2^i/\alpha$ -padded in the partition, where the probability that such partition exists is $1 - 1/n$.

Stretch Analysis Note that for each iteration i , the cost of an error report is $d(u, c) + 4d(T) + d(c, u) \leq 6 \cdot 2^i$, where c is the center of the cluster with u and T is the shortest path tree rooted at c and spanning the cluster. For the last iteration where the destination is reached, the cost is $d(u, c) + 2d(c, v) + 2d(T) \leq 5 \cdot 2^i$. W.l.o.g., we assume that $2^{i-1}/\alpha < d(u, v) \leq 2^i/\alpha$. Thus with probability $(1 - 1/n)$, the routing cost is at most $\sum_{j=0}^{i-1} 6 \cdot 2^j + 5 \cdot 2^i < 11 \cdot 2^i$. Therefore, $2^{i-1}/\alpha < d(u, v)$, and thus the routing stretch is at most 22α w.h.p.

References

- [1] I. Abraham, C. Gavoille, A. V. Goldberg, and D. Malkhi. Routing in networks with low doubling dimension. In *Proc. 26th ICDCS*, page 75, 2006.
- [2] I. Abraham, C. Gavoille, and D. Malkhi. Routing with improved communication-space trade-off. In *Proc. 18th DISC*, pages 305–319, 2004.
- [3] I. Abraham and D. Malkhi. Name independent routing for growth bounded networks. In *Proc. 17th SPAA*, pages 49–55, 2005.
- [4] W. Chen, C. Sommer, S.-H. Teng, and Y. Wang. Compact routing in power-law graphs. In *Proc. 23rd DISC*, pages 379–391, 2009.
- [5] J. Fakcheroenphol and K. Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *APPROX*, pages 36–46, 2003.
- [6] P. Fraigniaud and C. Gavoille. Routing in trees. In *Proc. 28th ICALP*, pages 757–772, 2001.
- [7] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals and low-distortion embeddings. In *Proc. 44th FOCS*, pages 534–543, 2003.
- [8] P. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proc. 25th STOC*, pages 682–690, 1993.
- [9] G. Konjevod, A. W. Richa, and D. Xia. Optimal scale-free compact routing schemes in networks of low doubling dimension. In *Proc. 18th SODA*, pages 939–948, 2007.
- [10] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. In *Proc. 45th FOCS*, pages 434–443, 2004.
- [11] S. Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proc. 15th SCG*, pages 300–306, 1999.