

On Space-Stretch Trade-Offs: Lower Bounds

Ittai Abraham¹ **Cyril Gavoille**² Dahlia Malkhi^{1,3}

¹Hebrew University of Jerusalem, Israel

²University of Bordeaux, France

³Microsoft Research

SPAA 2006

The Compact Routing Problem

Input: a network G (a weighted connected graph)

Output: a **routing scheme** for G

A routing scheme is a distributed algorithm that allows **any** source node to route messages to **any** destination node, given the destination's network identifier

The Compact Routing Problem


Input: a network G (a weighted connected graph)

Output: a **routing scheme** for G

A routing scheme is a distributed algorithm that allows **any** source node to route messages to **any** destination node, given the destination's network identifier


Goal: to minimize the size of the routing tables

Example: Grid with X,Y-coordinates



Routing algorithm: X,Y-routing

Example: Grid with X,Y-coordinates



Routing algorithm: X,Y-routing

Complexity Measures: Space & Stretch

Space = size of the largest local routing tables

Complexity Measures: Space & Stretch

Space = size of the largest local routing tables
(more precisely, size of the smallest local routing algorithm including all constants and data-structures)

In the example: space = $O(\log n)$ bits

Complexity Measures: Space & Stretch

Space = size of the largest local routing tables
(more precisely, size of the smallest local routing algorithm including all constants and data-structures)

In the example: space = $O(\log n)$ bits

Stretch = ratio between length of the route and distance

$$|\text{route}(x, y)| \leq \text{stretch} \cdot \text{dist}(x, y)$$

In the example: stretch = 1 (shortest path)

Complexity Measures: Space & Stretch

Space = size of the largest local routing tables
(more precisely, size of the smallest local routing algorithm including all constants and data-structures)

In the example: space = $O(\log n)$ bits

Stretch = ratio between length of the route and distance

$$|\text{route}(x, y)| \leq \text{stretch} \cdot \text{dist}(x, y)$$

In the example: stretch = 1 (shortest path)

Question: for a given family of graphs, find the best space-stretch trade-off

Two variants: Name-independent vs. Labeled

The destination enters the network with its **name**, which is determined by either the designer of the routing scheme (labeled), or an adversary (name-independent).

Labeled: the designer is free to name the nodes according to the topology and the edge weights of the graph

Name-independent: the input is a graph with fixed node names

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)
---------	---------	-------------------

arbitrary	1	$\tilde{O}(n)$	[folk]
-----------	---	----------------	--------

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)
---------	---------	-------------------

arbitrary	1	$\tilde{O}(n)$	[folk]
	$4k - 5$	$\tilde{O}(n^{1/k})$	[Thorup,Zwick]

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)	
arbitrary	1	$\tilde{O}(n)$	[folk]
	$4k - 5$	$\tilde{O}(n^{1/k})$	[Thorup,Zwick]
tree	1	$\tilde{O}(1)$	[TZ/Fraigniaud,G.]

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)	
arbitrary	1	$\tilde{O}(n)$	[folk]
	$4k - 5$	$\tilde{O}(n^{1/k})$	[Thorup,Zwick]
tree	1	$\tilde{O}(1)$	[TZ/Fraigniaud,G.]
doubling- α dim.	$1 + \varepsilon$	$\log \Delta$	[Talwar/Slivkins]
		$\tilde{O}(1)$	[Chan et al./Abraham et al.]

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)	
arbitrary	1	$\tilde{O}(n)$	[folk]
	$4k - 5$	$\tilde{O}(n^{1/k})$	[Thorup,Zwick]
tree	1	$\tilde{O}(1)$	[TZ/Fraigniaud,G.]
doubling- α dim.	$1 + \varepsilon$	$\log \Delta$	[Talwar/Slivkins]
		$\tilde{O}(1)$	[Chan et al./Abraham et al.]
planar	$1 + \varepsilon$	$\tilde{O}(1)$	[Thorup]

An overview: Labeled model

Labels are of polylogarithmic size

$$\tilde{O}(f(n)) = f(n) \cdot \text{polylog}(n)$$

network	stretch	space/node (bits)	
arbitrary	1	$\tilde{O}(n)$	[folk]
	$4k - 5$	$\tilde{O}(n^{1/k})$	[Thorup,Zwick]
tree	1	$\tilde{O}(1)$	[TZ/Fraigniaud,G.]
doubling- α dim.	$1 + \varepsilon$	$\log \Delta$	[Talwar/Slivkins]
		$\tilde{O}(1)$	[Chan et al./Abraham et al.]
planar	$1 + \varepsilon$	$\tilde{O}(1)$	[Thorup]
H -minor-free	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham,G.]

An overview: Name-independent model

network	stretch	space/node (bits)	
bounded growth	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham et al.]

An overview: Name-independent model

network	stretch	space/node (bits)	
bounded growth	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham et al.]
doubling- α dim.	$9 + \varepsilon$	$\tilde{O}(1)$	[Konjevod et al./Abraham et al.]

An overview: Name-independent model

network	stretch	space/node (bits)	
bounded growth	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham et al.]
doubling- α dim.	$9 + \varepsilon$	$\tilde{O}(1)$	[Konjevod et al./Abraham et al.]
H -minor-free (unweighted)	$O(1)$	$\tilde{O}(1)$	[Abraham et al.]

An overview: Name-independent model

network	stretch	space/node (bits)	
bounded growth	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham et al.]
doubling- α dim.	$9 + \varepsilon$	$\tilde{O}(1)$	[Konjevod et al./Abraham et al.]
H -minor-free (unweighted)	$O(1)$	$\tilde{O}(1)$	[Abraham et al.]
trees	$2^k - 1$	$\tilde{O}(n^{1/k})$	[Laing]

An overview: Name-independent model

network	stretch	space/node (bits)	
bounded growth	$1 + \varepsilon$	$\tilde{O}(1)$	[Abraham et al.]
doubling- α dim.	$9 + \varepsilon$	$\tilde{O}(1)$	[Konjevod et al./Abraham et al.]
H -minor-free (unweighted)	$O(1)$	$\tilde{O}(1)$	[Abraham et al.]
trees	$2^k - 1$	$\tilde{O}(n^{1/k})$	[Laing]
arbitrary	$O(k^2 2^k)$	$\tilde{O}(n^{1/k})$	[Arias et al./Awerbuch, Peleg]
	$O(k)$	$\tilde{O}(n^{1/k})$	[next talk]

Lower bounds for name-independent

Rem: lower bound for labeled \Rightarrow lower bound for name-indep

Lower bounds for name-independent

Rem: lower bound for labeled \Rightarrow lower bound for name-indep

network	stretch	space/node (bits)
arbitrary	< 1.4	$\Omega(n \log n)$ [G.,Pérennès.]
	< 3	$\Omega(n)$ [G.,Gengler]
(only $k = 1, 2, 3, 5$)	$< 2k + 1$	$\Omega(n^{1/k})$ [Thorup,Zwick]

Lower bounds for name-independent

Rem: lower bound for labeled \Rightarrow lower bound for name-indep

network	stretch	space/node (bits)
arbitrary	< 1.4	$\Omega(n \log n)$ [G.,Pérennès.]
	< 3	$\Omega(n)$ [G.,Gengler]
(only $k = 1, 2, 3, 5$)	$< 2k + 1$	$\Omega(n^{1/k})$ [Thorup,Zwick]
trees	≤ 3	$\Omega(\sqrt{n})$ [Laing,Rajaraman]
	$\leq 9 - \varepsilon$	$\Omega(n^{(\varepsilon/60)^2})$ [Konjevod et al.]

Lower bounds for name-independent

Rem: lower bound for labeled \Rightarrow lower bound for name-indep

network	stretch	space/node (bits)
arbitrary	< 1.4	$\Omega(n \log n)$ [G.,Pérennès.]
	< 3	$\Omega(n)$ [G.,Gengler]
	(only $k = 1, 2, 3, 5$) $< 2k + 1$	$\Omega(n^{1/k})$ [Thorup,Zwick]
trees	≤ 3	$\Omega(\sqrt{n})$ [Laing,Rajaraman]
	$\leq 9 - \varepsilon$	$\Omega(n^{(\varepsilon/60)^2})$ [Konjevod et al.]
for all $k \geq 1$	$< 2k + 1$	$\Omega((n \log n)^{1/k})$ [this paper]


Theorem

- 1 Any name-indep. routing scheme using $< (n \log n)^{1/k}$ bits/node has a **max stretch** $\geq 2k + 1$ for some graph.
- 2 Any name-indep. routing scheme using $< (n/k)^{1/k}$ bits/node has an **average stretch** $\geq k/4$ for some graph.

Theorem

- 1 Any name-indep. routing scheme using $< (n \log n)^{1/k}$ bits/node has a **max stretch** $\geq 2k + 1$ for some graph.
- 2 Any name-indep. routing scheme using $< (n/k)^{1/k}$ bits/node has an **average stretch** $\geq k/4$ for some graph.

Rem 1: All previous lower bounds for labeled case (Peleg,Upfal / G.,Pérennès / G.,Gengler / Kranakis,Krizanc / Thorup,Zwick) are based on the construction of **dense large girth** graphs



if $\text{stretch} < 2k + 1$, then u is forced to "know" the edge (u, v)

Theorem

- 1 Any name-indep. routing scheme using $< (n \log n)^{1/k}$ bits/node has a **max stretch** $\geq 2k + 1$ for some graph.
- 2 Any name-indep. routing scheme using $< (n/k)^{1/k}$ bits/node has an **average stretch** $\geq k/4$ for some graph.

Erdős Conjecture: \exists graph of girth $2k + 2$ with $\Omega(n^{1+1/k})$ edges (proved only for $k = 1, 2, 3, 5$). So, the extra $(\log n)^{1/k}$ term **cannot** be obtained with a girth approach.

Theorem

- 1 Any name-indep. routing scheme using $< (n \log n)^{1/k}$ bits/node has a **max stretch** $\geq 2k + 1$ for some graph.
- 2 Any name-indep. routing scheme using $< (n/k)^{1/k}$ bits/node has an **average stretch** $\geq k/4$ for some graph.

Rem 2: It makes a clear separation between labeled and name-independent routing.

In the **label** model, $O(\text{polylog}(n))$ space and $O(1)$ average stretch exists for every graph! [Abraham, Bartal, Chan, Gupta, Kleinberg et al. (FOCS05)]


In the **name-indep** model, if space is $O(\text{polylog}(n))$, then the average stretch must be $\Omega(\log n / \log \log n)$ for some graphs.

Proof: The model

(only point 1)

$$R(x_i, h_i, q_i) = (h_{i+1}, p_{i+1})$$

$R(x_i, \cdot, \cdot)$ describes the routing algorithm in x_i




Proof: The model

(only point 1)

$$R(x_i, h_i, q_i) = (h_{i+1}, p_{i+1})$$


$R(x_i, \cdot, \cdot)$ describes the routing algorithm in x_i




Definition (Kolmogorov Complexity)

Given an object P , let $\mathcal{K}(P)$ denote the length of the smallest program that prints P and halts.

Proof: The counter-example




Proof: The counter-example



Lemma

$\exists L \subset \{1, \dots, n\}$ with $|L| = n/2$ such that every $P \subset L$ satisfies $|P| \leq \mathcal{K}(P) + 2 \log n$.

Proof: The counter-example




Lemma

$\exists L \subset \{1, \dots, n\}$ with $|L| = n/2$ such that every $P \subset L$ satisfies $|P| \leq \mathcal{K}(P) + 2 \log n$.

Edge weight is 1 or k : $w(e) = 1$ iff $\text{port}(e) \in L$.

Proof: The counter-example




Lemma

$\exists L \subset \{1, \dots, n\}$ with $|L| = n/2$ such that every $P \subset L$ satisfies $|P| \leq \mathcal{K}(P) + 2 \log n$.

Edge weight is 1 or k : $w(e) = 1$ iff $\text{port}(e) \in L$.

Node naming: light (=red) edges lead to name $\leq n/2$

Proof: The counter-example



Lemma


$\exists L \subset \{1, \dots, n\}$ with $|L| = n/2$ such that every $P \subset L$ satisfies $|P| \leq \mathcal{K}(P) + 2 \log n$.

Edge weight is 1 or k : $w(e) = 1$ iff $\text{port}(e) \in L$.

Node naming: light (=red) edges lead to name $\leq n/2$

Proof: Analysis 1/3


Assume any R is given (no limits on the headers), and uses $\leq M$ bits at every node



Idea: To analyze all the walks from the root to all light destinations (names $\leq n/2$).

Proof: Analysis 1/3

Assume any R is given (no limits on the headers), and uses $\leq M$ bits at every node




Idea: To analyze all the walks from the root to all light destinations (names $\leq n/2$).

Basic fact: If stretch $< 2k + 1$, then no heavy edge is used

Proof: Analysis 2/3


Assume now the stretch of R is $< 2k + 1$



$P_i = \{\text{all ports in alive walks after } i\text{th routing decision of } r\}$

Proof: Analysis 2/3

Assume now the stretch of R is $< 2k + 1$




$P_i = \{\text{all ports in alive walks after } i\text{th routing decision of } r\}$

- $P_i \subset L$ for all i

Proof: Analysis 2/3

Assume now the stretch of R is $< 2k + 1$




$P_i = \{\text{all ports in alive walks after } i\text{th routing decision of } r\}$

- $P_i \subset L$ for all i
- $\mathcal{K}(P_1) \leq M$ (why?)

Proof: Analysis 2/3

Assume now the stretch of R is $< 2k + 1$




$P_i = \{\text{all ports in alive walks after } i\text{th routing decision of } r\}$

- $P_i \subset L$ for all i
- $\mathcal{K}(P_1) \leq M$ (why?)
- $\mathcal{K}(P_{i+1}) \leq M \cdot (|P_i| + 1)$ (why?)

Proof: Analysis 2/3

Assume now the stretch of R is $< 2k + 1$



$P_i = \{\text{all ports in alive walks after } i\text{th routing decision of } r\}$

- $P_i \subset L$ for all i
- $\mathcal{K}(P_1) \leq M$ (why?)
- $\mathcal{K}(P_{i+1}) \leq M \cdot (|P_i| + 1)$ (why?)

So, $\mathcal{K}(P_{i+1}) \leq M \cdot (\mathcal{K}(P_i) + 2 \log n + 1)$, and therefore

$$|P_{i+1}| \leq \mathcal{K}(P_{i+1}) + 2 \log n + 1 \leq (1 + o(1)) \cdot M^{i+1}$$

Proof: Analysis 3/3

$$|P_{i+1}| \leq (1 + o(1)) \cdot M^{i+1}$$

$W_t = \{ \text{light destinations reached after the } t\text{th routing decision of } r \}$

Proof: Analysis 3/3

$$|P_{i+1}| \leq (1 + o(1)) \cdot M^{i+1}$$

$W_t = \{ \text{light destinations reached after the } t\text{th routing decision of } r \}$

- $|W_t| \leq |P_1| + |P_2| + \dots + |P_t|$

Proof: Analysis 3/3

$$|P_{i+1}| \leq (1 + o(1)) \cdot M^{i+1}$$

$W_t = \{ \text{light destinations reached after the } t\text{th routing decision of } r \}$

- $|W_t| \leq |P_1| + |P_2| + \dots + |P_t|$
- If $|W_t| < n/2$, then stretch $\geq 2t + 1$.
So, stretch $< 2k + 1$ implies $|W_k| \geq n/2$.

Proof: Analysis 3/3

$$|P_{i+1}| \leq (1 + o(1)) \cdot M^{i+1}$$

$W_t = \{ \text{light destinations reached after the } t\text{th routing decision of } r \}$

- $|W_t| \leq |P_1| + |P_2| + \dots + |P_t|$
- If $|W_t| < n/2$, then stretch $\geq 2t + 1$.

So, stretch $< 2k + 1$ implies $|W_k| \geq n/2$.

Combining: $n/2 \leq \sum_{i=1}^k (1 + o(1)) \cdot M^i = O(M^k)$, i.e.,
 $M = \Omega(n^{1/k})$.

Proof: Analysis 3/3

$$|P_{i+1}| \leq (1 + o(1)) \cdot M^{i+1}$$

$W_t = \{ \text{light destinations reached after the } t\text{th routing decision of } r \}$

- $|W_t| \leq |P_1| + |P_2| + \dots + |P_t|$
- If $|W_t| < n/2$, then stretch $\geq 2t + 1$.

So, stretch $< 2k + 1$ implies $|W_k| \geq n/2$.

Combining: $n/2 \leq \sum_{i=1}^k (1 + o(1)) \cdot M^i = O(M^k)$, i.e.,
 $M = \Omega(n^{1/k})$.

Actually, a finer argument yields $M = \Omega((n \log n)^{1/k})$.

QED

Thank you!