LIP6

COMPACT ROUTING Main results and techniques

LaBRI - Université de Bordeaux - CNRS

25 septembre 2015

CO-AUTHORS

Cyril Gavoille





Nicolas Hanusse



David Ilcinkas



- Models and results.
- How to build a compact routing scheme.
- Compact routing in internet-like graphs.







What is compact routing, why does it exist?

COMPACT ROUTING

MOTIVATIONS FOR COMPACT ROUTING

- Routing in large scale networks
 - having better scaling capabilities
 - forward packets faster
 - maintain routing tables up to date efficiently
- May help to save energy









MEMORY # entries stored per node





MEMORY # entries stored per node





MEMORY # entries stored per node

Route length (u,v) = 5Distance (u,v) = 4

Stretch (u,v) = 5/4

STRETCH maximal stretch over all pairs



DISTRIBUTED ROUTING



Partial knowledge of the graph (local)

⇒ need of communication to compute routing tables





DISTRIBUTED ROUTING



COMMUNICATION COST # of small* messages exchanged

*polylogarithmic in the size of the network

Partial knowledge of the graph (local)

⇒ need of communication to compute routing tables





COMPACT ROUTING



Increased energy consumption

Forwarding packets faster Saving energy



Communication cost



- Cost efficient maintenance
- Saving energy



NAME-INDEPENDENT VS. LABELED

- Two models are considered for nodes naming (addresses)
 - Name-independent routing: use arbitrary routing addresses
 - Labeled routing: the designer chooses nodes' names

odes naming (addresses) e arbitrary routing addresses chooses nodes' names



F

Name independent routing





F

Name independent routing



Query route to « A »



F

Name independent routing





Query route to « A »



F

Name independent routing



Labeled routing



Query route to « A »



F

Name independent routing



Labeled routing



Query route to « A »

Query route to « A, 5 »



	Stretch	Memory	Communication cost
[AGM'06a]	< 2k+1	$\Omega(n^{1/k})$	any

lower bounds



		Stretch	Memory	Communication cost
	[AGM'06a]	< 2k+1	$\mathbf{\Omega}(\mathbf{n}^{1/k})$	any
OME	[GGHI'13]	1	any	$\Omega(n^2)$



		Stretch	Memory	Communication cost
	[AGM'06a]	< 2k+1	$\mathbf{\Omega}(\mathbf{n}^{1/k})$	any
owe	[GGHI'13]	1	any	$\Omega(n^2)$
	BFS-tree	1	Õ(n)	Õ(nm)



		Stretch		Memory	Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3	$\Omega(n^{1/k})$ $\Omega(n)$	any
OWE	[GGHI'13]	1		any	$\Omega(n^2)$
	BFS-tree	1	optir	mal Õ(n)	Õ(nm)



		Stretch	Memory	Communication cost
	[AGM'06a]	< 2k+1 < 3	$\Omega(n^{1/k})$ $\Omega(n)$	any
owe	[GGHI'13]	1	any	$\Omega(n^2)$
	BFS-tree	1 opti	mal Õ(n)	Õ(nm)
	[AR'93]	1	Õ(n)	$\tilde{O}(n^2)$



		Stretch	Memory	Communication cost
r bounds	[AGM′06a]	< 2k+1 < 3	$\mathbf{\Omega}(n^{1/k})$ $\mathbf{\Omega}(n)$	any
IOWe	[GGHI'13]	1	any	$\Omega(n^2)$
	BFS-tree	1 opti	mal Õ(n)	Õ(nm)
	[AR'93]	1	Õ(n)	$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)	$\tilde{O}(n^{1/k})$	centralised
DO				

upper



		Stretch		Memory		Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3	$\Omega(n^{1/k})$	Ω(n)	any
owe	[GGHI'13]	1		any		$\mathbf{\Omega}(n^2)$
	BFS-tree	1	optir	mal Õ(n)		Õ(nm)
	[AR'93]	1		$\tilde{O}(n)$		$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)		$\tilde{O}(n^{1/k})$		centralised
noq	[AGM'08]	3		$\tilde{O}(\sqrt{n})$		centralised
upper						



		Stretch			Memory		Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3 < 5		$\Omega(n^{1/k})$	Ω(n) Ω(√n)	any
owe	[GGHI'13]	1			any		$\Omega(n^2)$
	BFS-tree	1	opti	mal	Õ(n)		Õ(nm)
	[AR'93]	1			$\tilde{O}(n)$		$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)			$\tilde{O}(n^{1/k})$		centralised
noq .	[AGM'08]	3	opti	mal	$\tilde{O}(\sqrt{n})$		centralised
upper							



		Stretch			Memory		Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3 < 5		$\Omega(n^{1/k})$	Ω(n) Ω(√n)	any
owe	[GGHI'13]	1			any		$\Omega(n^2)$
	BFS-tree	1	opti	mal	Õ(n)		Õ(nm)
	[AR'93]	1			$\tilde{O}(n)$		$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)			$\tilde{O}(n^{1/k})$		centralised
bou	[AGM'08]	3	opti	mal	$\tilde{O}(\sqrt{n})$		centralised
pper	[SGF+10]	7			$\tilde{O}(\sqrt{n})$		NC.



		Stretch			Memory		Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3 < 5		$\Omega(n^{1/k})$	Ω (n) Ω (√n)	any
ove	[GGHI'13]	1			any		$\Omega(n^2)$
	BFS-tree	1	opti	mal	$\tilde{O}(n)$		Õ(nm)
	[AR'93]	1			$\tilde{O}(n)$		$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)			$\tilde{O}(n^{1/k})$		centralised
noq .	[AGM'08]	3	opti	mal	$\tilde{O}(\sqrt{n})$		centralised
pper	[SGF+10]	7			$\tilde{O}(\sqrt{n})$		NC.
D	[GGHI'13]	7			$\tilde{O}(\sqrt{n})$		O(m√n)



		Stretch			Memory		Communication cost
r bounds	[AGM'06a]	< 2k+1	< 3 < 5		$\Omega(n^{1/k})$	Ω (n) Ω (√n)	any
owe	[GGHI'13]	1			any		$\Omega(n^2)$
	BFS-tree	1	opti	mal	$\tilde{O}(n)$		Õ(nm)
	[AR'93]	1			$\tilde{O}(n)$		$\tilde{O}(n^2)$
nds	[AGM'06b]	O(k)			$\tilde{O}(n^{1/k})$		centralised
noq.	[AGM'08]	3	opti	mal	$\tilde{O}(\sqrt{n})$		centralised
pper	[SGF+10]	7			$\tilde{O}(\sqrt{n})$		NC.
D	[GGHI'13]	7			$\tilde{O}(\sqrt{n})$		O(m√n)
	my thesis	5			$\tilde{O}(\sqrt{n})$		O(m√n)





Memory	Communication cost
$\Omega(n^{1/k})$	any



		Stretch	Memory	Communication cost
lower bounds	[AGM'06]	< 2k+1	$\Omega(n^{1/k})$	any
	[GGHI'13]	1	any	Ω (n ²)





Memory	Communication cost	
$\Omega(n^{1/k})$	any	
any	$\Omega(n^2)$	

All name-independent results hold,





	Memory	Communication cost			
	$\Omega(n^{1/k})$	any			
	any	$\Omega(n^2)$			
ep	ependent results hold,				

|--|



		Stretch	Memory	Communication cost		
sounds	[AGM'06]	< 2k+1	$\Omega(n^{1/k})$	any		
ower k	[GGHI'13]	1	any	$\Omega(n^2)$		
S	All name-independent results hold,					
upper bound	[TZ'01]	4k-5	$\tilde{O}(n^{1/k})$	centralised		
	[TZ'01]	3	$\tilde{O}(\sqrt{n})$	centralised		



HOW TO BUILD A ROUTING SCHEME PART I : Labeled Compact Routing in Trees

LABELED COMPACT ROUTING IN TREES Two candidates

- I will not talk about the difference, they have very similar performances:
 - "Thorup-Zwick" [TZ'01]
 - "Fraigniaud-Gavoille" [FG'01]
- Achieve labeled compact routing in rooted trees with:
- Memory space: O(log² n)*
 - Stretch: 1
 - (Constant query time)

*[FG'01] actually have a O(log² n/loglog n) memory space


LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Heavy child ► GOAL, for every node u: a in the tree.

- - compute the compact routing label l(u) which represent path from u to r
- First step, heavy children nodes:
 - in the example the heavy child is always the rightmost one.



LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Nodes names a h

path	r	a	b	С	U





LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Nodes names a h

path	r	a	b	С	U
path*	r	*	*	С	*





LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Nodes names a h

path	r	a	b	С	U
path*	r	*	*	С	*
cpath_u	r		2	С	1



LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Nodes names a

path	r	а	b	С	U
path*	r	*	*	С	*
cpath_u	r		2	С	1
b_u	1	()	1	0



LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Nodes names a b

path	r	а	b	С	U
path*	r	*	*	С	*
cpath_u	r		2	С	1
b_u	1	()	1	0

Finally we set the name of node u to $\ell(u) = (u, cpath_u, b_u)$



LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Routing tables l(V) destination labels.



Routing tables are very small, every node stores two entries:

- link to the parent
- link to the heavy child
- Routing is based on "prefix comparisons" of the source and



Fraignaud-Gavoille - Routing algorithm (roughly)



- b_u is not a prefix of $b_v \Rightarrow \ell(v)$ is not a descendent
- route toward parent of $\ell(u)$ (info. in routing table)



Fraignaud-Gavoille - Routing algorithm (roughly)





Fraignaud-Gavoille - Routing algorithm (roughly)



- therefore $\ell(v)$ is a descendent of $\ell(u'')$
- since b_v ends with 0, $\ell(v)$ is in the heavy side,
- route using the routing table of $\ell(u'')$



Fraignaud-Gavoille - Routing algorithm (roughly)

 $\ell(\mathbf{v})$

 $\ell(\mathbf{w})$





Fraignaud-Gavoille - Routing algorithm (roughly)





Routing:

 $\ell(\mathbf{v})$

l(w

• similarly as before, but in this case the routing information is retrieved in the cpath of $\ell(v)$.

LABELED COMPACT ROUTING IN TREES Fraignaud-Gavoille - Labels size

- two routing entries per node.
- But what about the nodes names (labels) sizes?



Now we know that routing can be achieved giving nodes labels, with

Fraignaud-Gavoille - Labels size

path	r	а	b	С	U
path*	r	*	*	С	*
cpath_u	r	2)	С	1
b_u	1	()	1	0



Fraignaud-Gavoille - Labels size

path	r	а	b	С	U
path*	r	*	*	С	*
cpath_u	r)	С	1
b_u	1	()	1	0

Size of path can be n.
What about cpath size?



Fraignaud-Gavoille - Labels size

path	r	а	b	С	U
path*	r	*	*	С	*
cpath_u	r	2)	С	1
b_u	1	()	1	0

- Size of path can be n.
- What about cpath size?
 - (clue: the worst-case-tree is the binary one, why?)



What is the maximum number of 1's in b_u? #of non-heavy children?

Fraignaud-Gavoille - Labels size

path	r	а	b	С	U
path*	ľ	*	*	С	*
cpath_u	r	2)	С	1
b_u	1	()	1	0

- Size of path can be n.
- What about cpath size?
 - What is the maximum number of 1's in b_u? #of non-heavy children? (clue: the worst-case-tree is the binary one, why?)
 - By definition, path* stars are grouped if they are consecutive, therefore the number of O's in b_u is also bounded by log n.



Fraignaud-Gavoille - Labels size

path	r	а	b	С	U
path*	ľ	*	*	С	*
cpath_u	r	2)	С	1
b_u	1	()	1	0

- Size of path can be n.
- What about cpath size?
 - What is the maximum number of 1's in b_u? #of non-heavy children? (clue: the worst-case-tree is the binary one, why?)
 - By definition, path* stars are grouped if they are consecutive, therefore the number of O's in b_u is also bounded by log n.
- Every item in cpath is of size O(log n) (either a length or a node id), therefore, the total size of l(u) is O(log² n).



LABELED COMPACT ROUTING IN TREES Two candidates

- Performances of [FG'01] for rooted trees:
 - Memory space: O(log² n)
 - Stretch: 1 (shortest path)
 - Labeled scheme



HOW TO BUILD A ROUTING SCHEME PART II : Labeled compact routing for general graphs





Using "reference node(s)"
 we call them, Landmark(s)





Using "reference node(s)"
 we call them, Landmark(s)





In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.







- In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.
- But not every graph is a tree ...







- In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.
- But not every graph is a tree ...







- In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.
- But not every graph is a tree ...







- In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.
- But not every graph is a tree ...
- The stretch can be unbounded.







- In a tree, with [FG'01], we can do shortest path with constant memory in the labeled model.
- But not every graph is a tree ...
- The stretch can be unbounded.
- Solution: use "vicinity balls"



 $\ell(v')$





- For nodes closer than L e.g., $\ell(v)$, use classical shortest path routing
- For others e.g., $\ell(\mathbf{v}')$, use [FG'01] routing
 - This way the stretch is bounded by 3.
 - But ... vicinity balls can be huge, and so as the routing tables size.













With \sqrt{n} landmarks

 $\ell(u)$





L(u)

With \sqrt{n} landmarks vicinity balls have $\tilde{O}(\sqrt{n})$



With \sqrt{n} landmarks vicinity balls have size $\tilde{O}(\sqrt{n})$

 $\ell(u)$

For \sqrt{n} landmarks nodes:

- Memory: $\tilde{O}(\sqrt{n})$
 - Stretch: 3 at most
 - Similar to the scheme by Thorup & Zwick [TZ'01]





For \sqrt{n} landmarks nodes:

l(u)

- Memory: $\tilde{O}(\sqrt{n})$
 - Stretch: 3 at most
 - Similar to the scheme by Thorup & Zwick [TZ'01]



 $\ell(v)$

HOW TO BUILD A ROUTING SCHEME PART III : From Labeled to Name-independent routing

For \sqrt{n} landmarks nodes:

l(u)

- Memory: $\tilde{O}(\sqrt{n})$
 - Stretch: 3 at most
 - Similar to the scheme by Thorup & Zwick [TZ'01]



V
LABELED COMPACT ROUTING SCHEME

For \sqrt{n} landmarks nodes:

l(u)

- Memory: $\tilde{O}(\sqrt{n})$
 - Stretch: 3 at most
 - Similar to the scheme by Thorup & Zwick [TZ'01]



V

How to adapt this algorithm in the name-independent model?

- When routing from U to V, if $\ell(U)$ and $\ell(V)$ can be retrieved then we are done.
 - otherwise memory would me O(n).
 - 1. every node U could store its own label $\ell(U)$ (constant memory) 2. every node U cannot store every destination label
- The routing label $\ell(v)$ can be retrieved using local collaboration amongst vicinity nodes.





Vicinity label-sharing

all n labels

Identifier	comp routing-
V1	$\ell(v)$
V 2	$\ell(v_2)$
V 3	$\ell(v)$
V4	$\ell(\mathbf{v})$
• • •	• • •
$v_i = v$	l(v
• • •	• • •
Vn	$\ell(v_r)$
	Identifier V_1 V_2 V_3 V_4 V_4 $V_i = V$ V_i



pactng-label

- V_1
- (v_2)
- (v_3)
- (v_4)

- (\mathbf{v})
- \ $v_n)$

Vicinity label-sharing

	all n	labels compact-		Identifier	compo routing-l
	Identitier	routing-label		V 2	$\ell(v_2)$
	V1	$\ell(v_1)$		V3	$\ell(v_3)$
U3	V 2	$\ell(v_2)$	hash	Vi	$\ell(v_i)$
u 🔵	V 3	$\ell(v_3)$	function		
	V4	$\ell(v_4)$	$n \rightarrow \sqrt{n}$		compo
	• • •	• • •		Identifier	routina-l
U5	$v_i = v$	$\ell(v)$			
	• • •	0 0 0		V 1	
	Vn	$\ell(v_n)$		V4	E (V4)
				V9	$\ell(V_9)$







Vicinity label-sharing

				l	16
	all n Identifier	labels compact-		Identifier	compo routing-l
	Idennier	routing-label		V 2	$\ell(v_2)$
	V1	$\ell(v_1)$		V 3	$\ell(v_3)$
U3	V 2	$\ell(v_2)$	hash	Vi	$\ell(v_i)$
u	V 3	$\ell(v_3)$	function		
	V 4	$\ell(v_4)$	$n \rightarrow \sqrt{n}$		
U U U U U U U U U U U U U U U U U U U	• • •	• • •		Identifier	compo routing-l
U U 5	$V_i = V$	$\ell(v)$			
	• • •			V1	$\mathcal{E}(V_1)$
	Vp	$\ell(v_{p})$		V 4	$\ell(v_4)$
	V []			V 9	$\ell(v_9)$







How to adapt this algorithm in the name-independent model?





How to adapt this algorithm in the name-independent model?

all labels

Identifier	cor routi
V1	l
V 2	l
V 3	l
V 4	l
• • •	
$v_i = v$	
• • •	
Vn	l
	Identifier V_1 V_2 V_3 V_4 V_4 $V_i = V$ \dots V_n



- mpacting-label
- $\ell(v_1)$
- $\ell(v_2)$
- $\ell(v_3)$
- $\ell(v_4)$
- *l*(v)
- $\ell(v_n)$

• • •

How to adapt this algorithm in the name-independent model?

all labels

Identifier	cor routi
V1	l
V 2	l
V 3	l
V 4	l
• • •	
$v_i = v$	
• • •	
Vn	l
	Identifier V_1 V_2 V_3 V_4 V_4 $V_i = V$ \dots V_n



- mpacting-label
- $\ell(v_1)$
- $\ell(v_2)$
- $\ell(v_3)$
- $\ell(v_4)$
- $\ell(v)$
- $\ell(v_n)$

• • •

 Moreover, in its vicinity routing table, node U stores the colours of the nodes.

How to adapt this algorithm in the name-independent model?

all labels

Identifier	cor routi
V1	l
V 2	l
V 3	l
V 4	l
• • •	
$v_i = v$	
• • •	
Vn	l
	Identifier V_1 V_2 V_3 V_4 V_4 $V_i = V$ \dots V_n



- mpacting-label
- $\ell(v_1)$
- $\ell(v_2)$
- $\ell(v_3)$
- $\ell(v_4)$
- $\ell(v)$

• • •

 $\ell(\mathbf{v}_n)$

• • •

- Moreover, in its vicinity routing table, node U stores the colours of the nodes.
- The memory used for label sharing is #labels/#colors, for example it can be:

 $n/\sqrt{n} = \sqrt{n}$













• Every node U stores two types of entries:





- Every node U stores two types of entries:
 - Next-hop and color of nodes in its vicinity ball, ie., $\tilde{O}(\sqrt{n})$ entries.





- Every node U stores two types of entries:
 - Next-hop and color of nodes in its vicinity ball, ie., $\tilde{O}(\sqrt{n})$ entries.
 - Compact routing labels of nodes whose hash value is blue, ie., $\tilde{O}(\sqrt{n})$ entries.





- Every node U stores two types of entries:
 - Next-hop and color of nodes in its vicinity ball, ie., $\tilde{O}(\sqrt{n})$ entries.
 - Compact routing labels of nodes whose hash value is blue, ie., $\tilde{O}(\sqrt{n})$ entries.
- Therefore, the total memory per node is $\tilde{O}(\sqrt{n})$

)





• Hyp.: d(u,v) = 1



Hyp.: d(u,v) = 1 $\Rightarrow \text{Radius} \leq 1$

Hyp.: d(u,v) = 1
⇒ Radius ≤ 1

$$d(v,L_v) \leq d(v,L_u) \leq 2$$

- Hyp.: d(u,v) = 1 $\Rightarrow \text{Radius} \leq 1$
- ► $d(v,L_v) \le d(v,L_u) \le 2$ • $d(u,Lv) \le d(v,Lv)+d(u,v) \le 3$

- Hyp.: d(u,v) = 1 $\blacktriangleright \Rightarrow \text{Radius} \leq 1$
- ► $d(v,L_v) \le d(v,L_u) \le 2$ • $d(u,Lv) \leq d(v,Lv)+d(u,v) \leq 3$

 $"u \rightarrow w \rightarrow u \rightarrow Lv \rightarrow v \leq 7"$

- Hyp.: d(u,v) = 1 $\blacktriangleright \Rightarrow \text{Radius} \leq 1$
- ► $d(v,L_v) \le d(v,L_u) \le 2$
- ► $d(u,Lv) \le d(v,Lv)+d(u,v) \le 3$
- $"u \rightarrow w \rightarrow u \rightarrow Lv \rightarrow v \leq 7"$
- Therefore the stretch is ≤ 7 .

	Stretch	Memory	Communication cost
[AGM'08]	3 opti	mal $\tilde{O}(\sqrt{n})$	centralised
[GGHI'13]	7	$\tilde{O}(\sqrt{n})$	O(m√n)
my thesis	5	$\tilde{O}(\sqrt{n})$	O(m√n)
To be done (hard)	3	$\tilde{O}(\sqrt{n})$	o(mn)

- complexities.

PERFORMANCES

If the route use the best landmark from $\{L_{u}, L_{v}\}$ the stretch can be improved to 5.

This can even be done in a distributed way without any impact on the asymptotic

COMPACT ROUTING FOR INTERNET-LIKE GRAPHS

INTERNET-LIKE GRAPHS

- Power-law degree distribution
 - Low diameter (also logarithmic)
 - Sparse graph (logarithmic average degree)
- Looks like a tree with a dense cluster as a root
- Mimicked by a synthetic graph model, called RPLG:
 - Random Power Law Graph
 - Parametrised by the power-law exponent, t

INTERNET-LIKE GRAPHS (RPLG)

Power-law exponent t = 2.1

MPROVING MEMORY PERFORMANCES

We are actually very used in routing with these settings, lets think about this

- Low memory
- Low stretch
- Dense center
- Very sparse outside the center
- Most of the destinations are outside the center

MPROVING MEMORY PERFORMANCES

about Public transportation

- Low memory New city "there's no one around and your phone is dead"* Low stretch - Its holidays, you can't wait in transports Dense center - City center is easily accessible Very sparse outside the center - Suburban public transportation ...

- Most of the destinations are outside the center Every analogy limps ...

How to use the structure to save some entries?

We are actually very used in routing with these settings, lets think

IMPROVING MEMORY PERFORMANCES

MPROVING MEMORY PERFORMANCES

IMPROVING MEMORY PERFORMANCES





MPROVING MEMORY PERFORMANCES

How to use the structure to save some entries?

What is the hash value of v?





IMPROVING MEMORY PERFORMANCES





MPROVING MEMORY PERFORMANCES

How to use the structure to save some entries?

What is the label of node v? $\ell(v)$





IMPROVING MEMORY PERFORMANCES

















THEORETICAL RESULTS





Compared to older results for similar settings on RPLG

Memory				
average	maximum			
$O(n^{1/(t-1)+t-3})$	O(n ^{1-1/(t-1)})			
O(n ^{1/110})	O(n ^{1/11})			
$O(n^{(t-2)/(2t-3)})$				
$O(n^{1/12})$				
O(√n)	O(n)			

Upper bounds for t = 2.1

IN DEPTH [CSWT'12] VS. [GGHI'15]

Theoretical comparison for various exponential values

- when t < 2.707 [GGHI'15] avg. memory is</p> smaller
- [CSTW'12]:
 - has a smaller theoretical max. memory
 - but is a labeled scheme!
- For small values of t, hidden constants may play an important role

What are the performances in practice?





	avg stretch	memory		
		average	maximum	
[AGMNT'08]	1.56	396	1 1 4 3	
ed [CSTW'12]	1.30	55.2	580	
[TZLL'13]	1.24	404	1 877	
[GGHI'15]	1.75	6.47	228	
	[AGMNT'08] ed [CSTW'12] [TZLL'13] [GGHI'15]	avg stretch [AGMNT'08] 1.56 ed [CSTW'12] 1.30 [TZLL'13] 1.24 [GGHI'15] 1.75	avg stretch me avg stretch average average average average 1.56 average 396 addition 1.30 average 1.30 average 1.24 average 404 average 1.75	

EXPERIMENTAL RESULTS For an RPLG with t=2.1



	For an AS-graph 16K nodes				
		avg stretch	memory		
			average	maximum	
Labele	[AGMNT'08]	1.74	465	1 261	
	ed [CSTW'12]	1.18	24	687	
	[TZLL'13]	1.52	106	2 3 2 4	
	[GGHI'15]	1.59	4.05	415	

EXPERIMENTAL RESULTS



FUTURE WORKS

- In sparse networks, for q=1/2, routing tables of size $O(n^q)$ • communication cost is $\tilde{O}(n^{1+q})$
- - can that be done for any q?
- Labeled = Name-independent? (True for stretch 3)* ie., name-independent routing with stretch 2k+1 and memory $O(n^{1/k})$
- Very compact routing schemes for internet-like graphs: "logarithmic memory is required for infinite scaling"

*Disclaimer: No prize is offered for solving this.





MERCI