

Internet as a doubling metric

Laurent Viennot, Emmanuelle Lebhar, Pierre Fraigniaud

Équipe-projet GANG Inria – Liafa

Internet Graph

- AS graph, Routers graph
- Power laws : Faloutsos et al. [SIGCOMM'99]
- Connectivity : Broido and claffy [SPIE'01]
- Not exactly power laws : Chen et al. [INFOCOM'02]

Modeling and estimating internet delays

- Triangulation with landmarks : Francis et al. [INFOCOM'99]
- Euclidean space embedding (GNP) : Ng and Zhang [INFOCOM'02]
- Hyperbolic space embedding : Shavitt and Tankel [INFOCOM'04]
- Non metric model (min-plus) : Key et al. [INFOCOM'08]

Closest nodes

- Delays as distances.
- Nearest server, through distance estimation (landmarks, GNP, many others...).
- Nearest neighbor (distributed) : Karger and Ruhl [STOC'2002]
- Working testbed : Wong et al. [SIGCOMM'2005]

Routing

- Efficient overlay network : Abraham et al. [SODA'04]
- Routing scheme design : Krioukov et al. [INFOCOM'04], Thorup et Zwick [STOC'2001]
- Sparse spanner : Chan and Gupta [SODA'06]

Definition

Latency distance : $d(u, v) = RTT(u, v)$.


(The Round Trip Time is the time to send a packet from u to v and to receive an acknowledgment back from u to v).

Algorithms assumptions

- d is a metric.
- Bounded growth (more general than euclidean spaces) [Plaxton et al. 99].
- Doubling metric (more general than bounded growth) [Gupta et al. 03, Kleinberg et al. 04, Slivkins 05].

Example of algorithmic use

- Triangle inequality for low stretch overlay routing.



Skewed triangle

- $d(w, v) \leq \delta d(u, v)$.
- Triangle inequality :
 $d(u, w) \leq d(u, v) + d(v, w) \leq (1 + \delta)d(u, v)$.
- Routing from u to v through w results in a stretch $1 + 2\delta$.

Definition

Latency distance : $d(u, v) = RTT(u, v)$.

(The Round Trip Time is the time to send a packet from u to v and to receive an acknowledgment back from v to u).

Is this a doubling metric ?

- $d(u, u) \approx 0$
- should be symmetric (up to time variations),
- triangle inequality ?
- ball growth ?
- doubling property (more general than ball growth) ?

Metric

- Triangle inequality [Francis et al. 01, Zhang 06]
- Symmetry [Pathak et al. 07]

Bounded growth

- On average [Zhang 06]

Doubling dimension

- Next slides...

Ping

- Direct RTT measurements.
- Traceroutes from skitter data $20 \times 400\,000$ [CAIDA]
- All-sites-ping 200×200 , every hour [Yoshikawa and Stribling]

King

- Inter-nameservers (DNS) delays through recursive requests [Gummadi et al. 02].
- P2PSim 1700×1700 [Gil et al.]
- Meridian project 2300×2300 [Wong et al. 05]

Acknowledgments

Thanks to all for providing their data.

Ping

- Direct RTT measurements.
- Traceroutes from skitter data $20 \times 400\,000$ [CAIDA]
- All-sites-ping 200×200 , every hour [Yoshikawa and Stribling]

King

- Inter-nameservers (DNS) delays through recursive requests [Gummadi et al. 02].
- P2PSim 1700×1700 [Gil et al.]
- Meridian project 2300×2300 [Wong et al. 05]

Acknowledgments

Thanks to all for providing their data.

Ping

- Direct RTT measurements.
- Traceroutes from skitter data $20 \times 400\,000$ [CAIDA]
- All-sites-ping 200×200 , every hour [Yoshikawa and Stribling]


King

- Inter-nameservers (DNS) delays through recursive requests [Gummadi et al. 02].
- P2PSim 1700×1700 [Gil et al.]
- Meridian project 2300×2300 [Wong et al. 05]

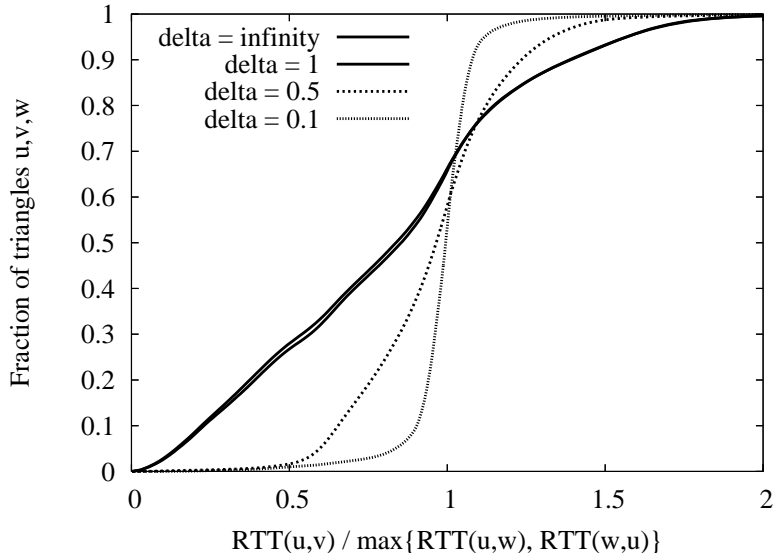
Acknowledgments

Thanks to all for providing their data.

Triangle inequality

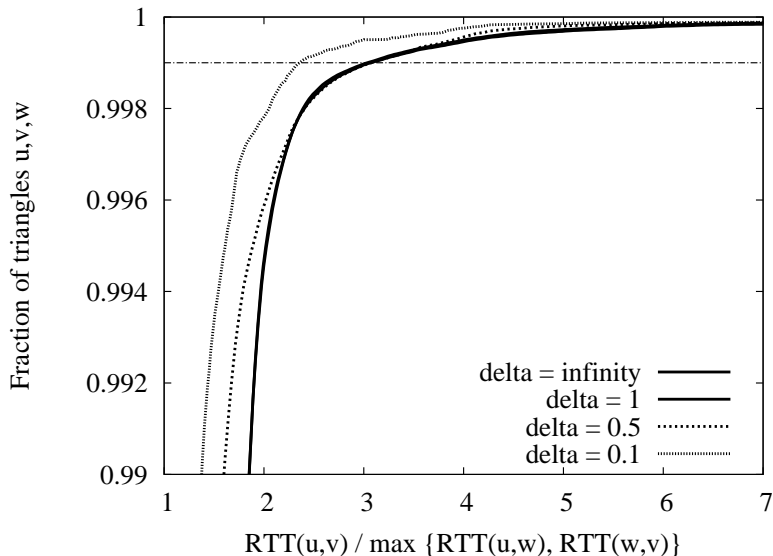


Max triangle inequality



- Less than $1 + \delta$ for a classical metric.

Max triangle inequality zoomed



- Less than $1 + \delta$ for a classical metric.


Definition

d is an inframetric if it verifies metric axioms where triangle inequality is replaced by :


$$d(u, v) \leq \rho \max\{d(u, w), d(w, v)\}$$

for some fixed ρ .

Triangles



Skewed triangles



Skewed triangle


- $d(w, v) \leq \delta d(u, v)$.
- Triangle inequality :
 $d(u, w) \leq d(u, v) + d(v, w) \leq (1 + \delta)d(u, v)$.

Definition

An inframetric d is (ρ_s, δ) -skewed if for all δ -skewed triangle u, v, w ,

$$d(u, v) \leq \rho_s \max\{d(u, w), d(w, v)\}$$

Skewed triangles



Definition

Let d be a distance function. The ball of radius r centered on a node u is

$$B_u(r) = \{v \mid d(u, v) \leq r\}$$

Definition

A distance function has growth γ if $|B_u(2r)| \leq \gamma |B_u(r)|$ for all u, r .

Some facts

- The size of $B_u(r)$ increases polynomially with r .
- Any metric of an euclidean space of dimension k has growth $O(k)$.
- Same thing for the shortest path distance of a grid with dimension k .

Ping

- Direct RTT measurements.
- Traceroutes from skitter data $20 \times 400\,000$ [CAIDA]
- All-sites-ping 200×200 , every hour [Yoshikawa and Stribling]

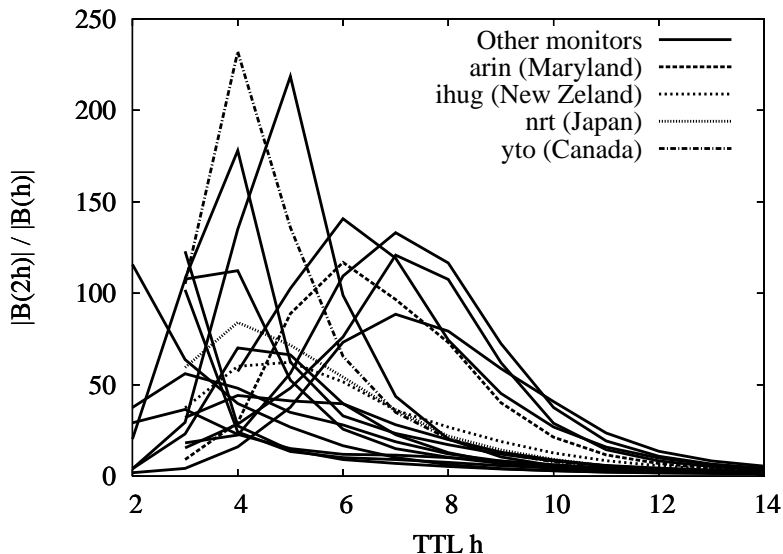
King

- Inter-nameservers (DNS) delays through recursive requests [Gummadi et al. 02].
- P2PSim 1700×1700 [Gil et al.]
- Meridian project 2300×2300 [Wong et al. 05]

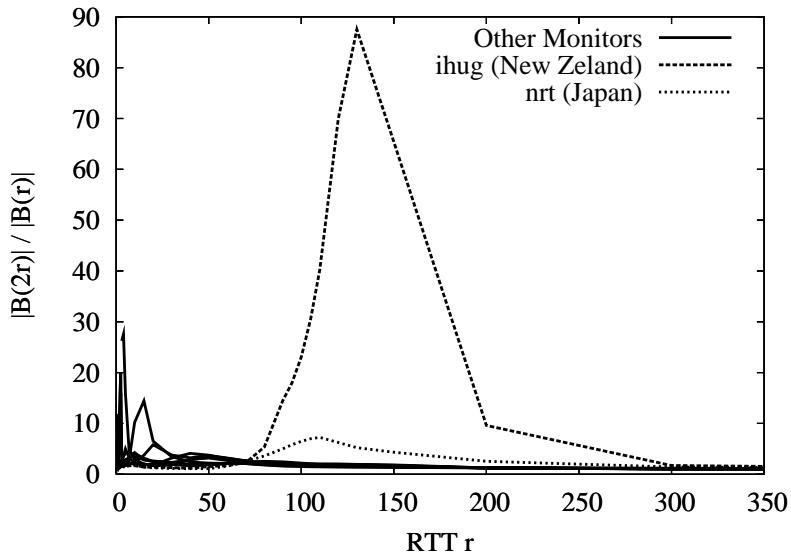
Acknowledgments

Thanks to all for providing their data.

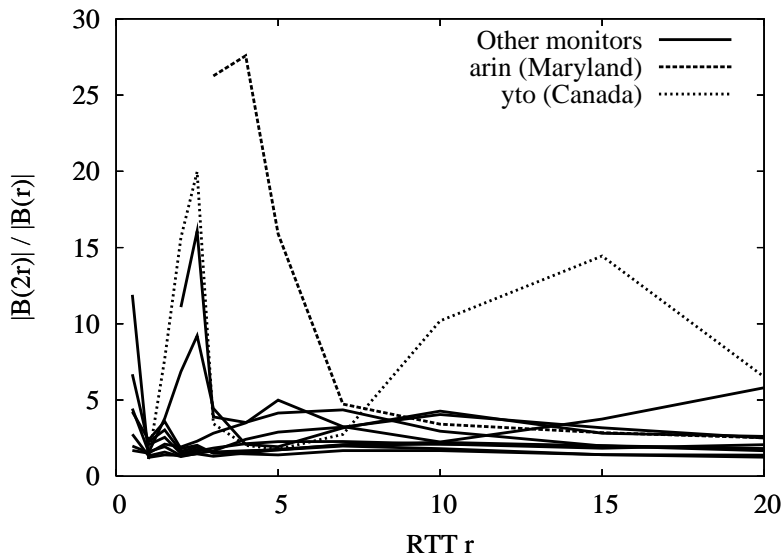
Hops ball growth




RTT ball growth



RTT ball growth




Doubling Metrics




- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics




- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics




- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics



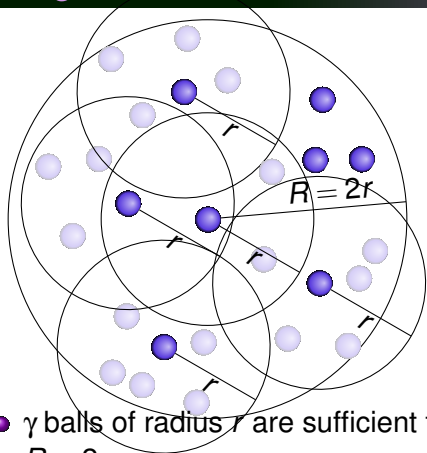
- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics




- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics



- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Doubling Metrics



- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.
- There exist compact routing schemes with $O(\gamma^{\log \frac{1}{\epsilon}} \log n)$ routing tables and stretch $1 + \epsilon$ [Slivkins 06, Abraham et al. 05].

Ping

- Direct RTT measurements.
- Traceroutes from skitter data $20 \times 400\,000$ [CAIDA]
- All-sites-ping 200×200 , every hour [Yoshikawa and Stribling]

King

- Inter-nameservers (DNS) delays through recursive requests [Gummadi et al. 02].
- P2PSim 1700×1700 [Gil et al.]
- Meridian project 2300×2300 [Wong et al. 05]

Acknowledgments

Thanks to all for providing their data.

Classical

- γ balls of radius r are sufficient to cover a ball of radius $R = 2r$.
- $\gamma^{\lceil \log_2 R/r \rceil}$ balls of radius r are sufficient to cover a ball of radius $R > r$.

Definition

A ρ -inframetric is γ -doubling if γ balls of radius r are sufficient to cover a ball of radius $R = \rho r$.

Definition

A ρ -inframetric has bounded growth γ if $|B_u(2r)| \leq \gamma |B_u(r)|$ for all u, r .

Theorem

A γ -doubling ρ -inframetric has growth $O(\gamma)$.

Extended definition

For $r \geq \tau$ and $R \geq \rho r$, $\beta \alpha^{\log_\rho R/r}$ balls of radius r are sufficient to cover a ball of radius $R \geq \rho r$.

(Equivalent for $\alpha = \gamma$, $\beta = 1$ and $\tau = 0$).

Definition

A ρ -inframetric has bounded growth γ if $|B_u(2r)| \leq \gamma |B_u(r)|$ for all u, r .

Theorem

A γ -doubling ρ -inframetric has growth $O(\gamma)$.

Extended definition

For $r \geq \tau$ and $R \geq \rho r$, $\beta \alpha^{\log_\rho R/r}$ balls of radius r are sufficient to cover a ball of radius $R \geq \rho r$.

(Equivalent for $\alpha = \gamma$, $\beta = 1$ and $\tau = 0$).

Voir figure

Most of the points lie between $2 \cdot (2.35)^{\log_2 R/r}$ and $25 \cdot (2.35)^{\log_2 R/r}$.

Theorem

Consider a ρ -inframetric (ρ_S, δ) -skewed (α, β, τ) -doubling inframetric. For any positive $\varepsilon < \delta/\rho_S$, there exist a compact routing scheme with stretch $(\rho_S/(1 - \rho\varepsilon), \tau\rho^2 \log n)$, table size $\beta\alpha^{3+\log_\rho(\frac{1}{\varepsilon})} \log_\rho(\Delta/\tau) \log n$ bits per node, and node label size $O(\log_\rho(\Delta/\tau) \log n)$ bits.

Corollary ((back to Slivkins 05))

For a classical metric, $\rho = 2$, $\rho_S = 1 + \varepsilon$. For any positive $\varepsilon < 1/2$, there exists a compact routing scheme with stretch $(\frac{1+\varepsilon}{1-2\varepsilon}, 4\tau \log n)$, table size $\beta\alpha^{3+\log_2(\frac{1}{\varepsilon})} \log_2(\Delta/\tau) \log n$ bits per node, and node label size $O(\log_2(\Delta/\tau) \log n)$ bits.

- Dynamics ?