Compact Policy Routing

Gábor Rétvári, András Gulyás, Zalán Heszberger, Márton Csernai, József J. Bíró

High Speed Networks Laboratory

Department of Telecommunications and Media Informatics

Budapest University of Technology and Economics

Email: {retvari, gulyas, heszi, csernai, biro}@tmit.bme.hu

MŰEGYETEM 1782

ACM PODC'11, June 6–8, 2011, San Jose, California, USA

Routing tables in the Internet grow rapidly

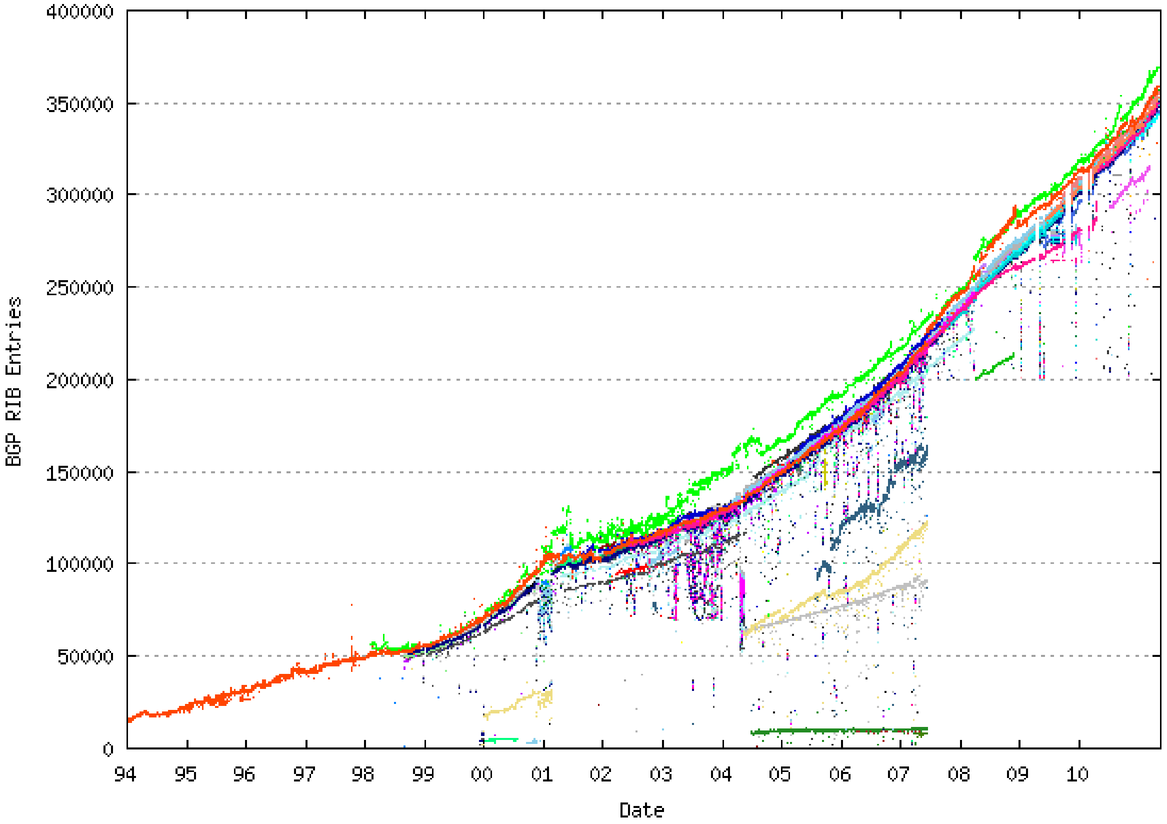

- There are many operational reasons behind this
- Are there any deep theoretical reasons as well?
- What can we do about it?

Compact routing and the Internet

- The theory dealing with fundamental routing scalability is called **compact routing**
- All our scalability results are for shortest paths
- But **Internet routing does not use shortest-paths**
- Instead, it is business interests, political pursuits and other operational concerns that shape Internet paths
  - business relations and SLA-s (BGP)
  - reliability and resilience (constraint-based routing)
  - available bandwidth (QoS routing, traffic engineering)
- Commonly called **policy routing**

This talk

- How do different policy routing architectures scale?
  - define a generic model for policy routing
  - characterize the memory requirements for implementing the models
  - discover the memory size–stretch trade-off
  - obtain tight bounds for policies important in practice

Routing algebras

- Abstract away **weight composition** ($\oplus$) and **weight comparison** ($\preceq$)
- A routing algebra $\mathcal{A} = (W, \phi, \oplus, \preceq)$ is a totally ordered commutative semigroup (Sobrinho, Griffin)
- Given a path $p = (v_1, v_2, \ldots, v_k)$, the weight $w(p)$ of p is $w(p) = \bigoplus_{i=1}^{k-1} w(v_i, v_{i+1})$
- A **preferred** $s - t$ **path** p^* over $\mathcal{A}$ is the one with the smallest weight: $p^* : w(p^*) \preceq w(p), \forall p \in \mathcal{P}_{st}$
- Examples
  - shortest-path routing: $\mathcal{S} = (\mathbb{R}^+, \infty, +, \leq)$
  - widest-path routing: $\mathcal{W} = (\mathbb{R}^+, 0, \min, \geq)$
  - most-reliable-path routing: $\mathcal{R} = ((0, 1], 0, *, \geq)$
  - various sorts of BGP policies

Composing algebras

- Given two routing algebras $\mathcal{A}$ and $\mathcal{B}$

  $$\mathcal{A} = (W_{\mathcal{A}}, \phi_{\mathcal{A}}, \oplus_{\mathcal{A}}, \preceq_{\mathcal{A}}) \quad \text{and} \quad \mathcal{B} = (W_{\mathcal{B}}, \phi_{\mathcal{B}}, \oplus_{\mathcal{B}}, \preceq_{\mathcal{B}}) \ ,$$

  their **lexicographic product** $\mathcal{A} \times \mathcal{B} = (W, \phi, \oplus, \preceq)$:

  $$W = W_{\mathcal{A}} \times W_{\mathcal{B}}$$

  $$(w_1, v_1) \oplus (w_2, v_2) = (w_1 \oplus_{\mathcal{A}} w_2, v_1 \oplus_{\mathcal{B}} v_2)$$

  $$(w_1, v_1) \preceq (w_2, v_2) = \begin{cases} v_1 \preceq_{\mathcal{B}} v_2 & \text{if } w_1 =_{\mathcal{A}} w_2 \\ w_1 \preceq_{\mathcal{A}} w_2 & \text{otherwise} \end{cases}$$

- Examples
  - widest-shortest path: $\mathcal{WS} = \mathcal{S} \times \mathcal{W}$
  - shortest-widest path: $\mathcal{SW} = \mathcal{W} \times \mathcal{S}$

Some useful algebraic properties

- A routing algebra $\mathcal{A}$ is **regular**, if it is both
  - **monotone** (M): $w_1 \preceq w_2 \oplus w_1$ for all $w_1, w_2 \in W$, and

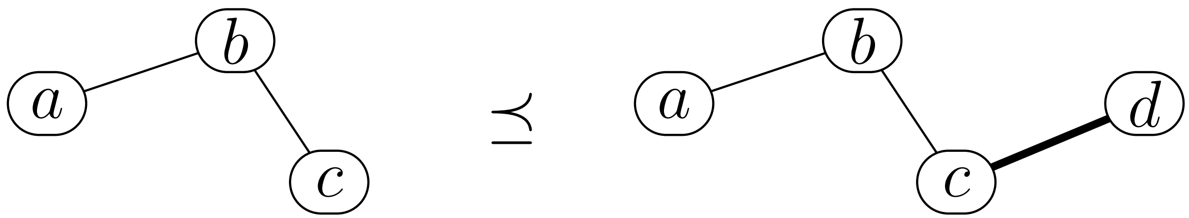

  - **isotone** (I): $w_1 \preceq w_2 \Rightarrow w_3 \oplus w_1 \preceq w_3 \oplus w_2$ for all $w_1, w_2, w_3 \in W$

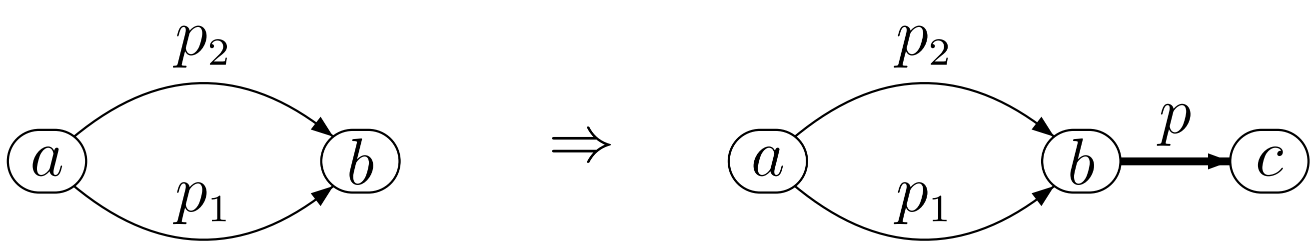

$w(p_1) \preceq w(p_2)$ $\qquad$ $w(p_1) \oplus w(p) \preceq w(p_2) \oplus w(p)$

- Strict versions exist with $\prec$ instead of $\preceq$

Routing process model

- The **local memory requirement** for implementing $\mathcal{A}$ is

$$M_{\mathcal{A}} = \max_{G \in \mathcal{G}_n} \min_{R \in \mathcal{R}} \max_{u \in V} M_{\mathcal{A}}(R, u) \ ,$$

  where $M_{\mathcal{A}}(R, u)$ is the number of bits needed to encode the local routing function R_u at some node u
- $\mathcal{A}$ is **incompressible** if there is at least one graph in which at least one node needs linear information for routing according to $\mathcal{A}$, i.e., $M_{\mathcal{A}}$ is $\Omega(n)$
- **Compressible** otherwise

Propositions

- **Proposition:** a routing policy $\mathcal{A}$ can be implemented by a destination-based routing function on any graph, if and only if it is **regular** (Sobrinho)
- Gives a trivial upper bound for regular algebras: if $\mathcal{A}$ is regular then $M_{\mathcal{A}}$ is $O(n \log d)$
- Question is, whether this trivial bound is tight
- **Proposition:** $\mathcal{S} = (\mathbb{R}^+, \infty, +, \leq)$ is incompressible (Fraigniaud, Gavoille, Pérennès)
- First, we ask whether similar characterization exists for other algebras

Local memory requirements

- Policies with „maximum-type” weight composition scale well
- **Def.:** $\mathcal{A}$ is **selective**, if $w_1 \oplus w_2 \in \{w_1, w_2\}$ for each $w_1, w_2 \in W$
- **Theorem:** if $\mathcal{A}$ is selective and monotone, then it is compressible
- Under the above conditions, a spanning tree exists in which the only path between any s and t is a preferred $s - t$ path
- On the other hand, „shortest-path-like” policies do not scale
- **Theorem:** if $\mathcal{A}$ is strictly monotone, then it is incompressible
- In fact, it is enough if $\mathcal{A}$ contains a subalgebra isomorphic to shortest path routing

Local memory requirements

Algebra	Definition	Prop.	$M_{\mathcal{A}}$
Shortest path	$\mathcal{S} = (\mathbb{R}^+, \infty, +, \leq)$	SM, I	$\Theta(n)$
Widest path	$\mathcal{W} = (\mathbb{R}^+, 0, \min, \geq)$	S, I, M	$\Theta(\log n)$
Most reliable path	$\mathcal{R} = ((0,1], 0, *, \geq)$	SM, I	$\Theta(n)$
Usable path	$\mathcal{U} = (\{1\}, 0, *, \geq)$	S, I, M	$\Theta(\log n)$
Widest-shortest path	$\mathcal{WS} = \mathcal{S} \times \mathcal{W}$	SM, I	$\Theta(n)$
Shortest-widest path	$\mathcal{SW} = \mathcal{W} \times \mathcal{S}$	SM, $\neg I$	$\Omega(n)$
BGP valley-free	$\mathcal{B}_1$	S	$\Theta(\log n)$
BGP local pref.	$\mathcal{B}_2$	$\neg M$, $\neg I$	$\Theta(n)$
BGP AS path length	$\mathcal{B}_3$	$\neg M$, $\neg I$	$\Theta(n)$

- Tight bounds for all regular algebras
- For $\mathcal{SW}$, all we know is that $M_{\mathcal{SW}}$ is $\Omega(n)$ and $O(n^2)$
- BGP policy routing is incompressible, even for a very minimalistic model

Algebraic compact routing

- Many important policies turn out incompressible
- With shortest path routing, a standard trick to decrease worst-case memory size is to allow routes to be somewhat longer
- Path length increase is upper bounded by some constant **stretch**
- Can we play out the same trick for policy routing?
- **Def.:** a routing scheme is of stretch k over algebra $\mathcal{A}$, if for any path p_{st} selected by the scheme:

$$w(p_{st}) \preceq \underbrace{w(p^*_{st}) \oplus w(p^*_{st}) \oplus \ldots \oplus w(p^*_{st})}_{\text{k times}} \ ,$$

where p^*_{st} is some preferred $s-t$ path in $\mathcal{A}$

Algebraic compact routing

- Regular algebras admit a compact implementation
- **Theorem:** if $\mathcal{A}$ is regular, then there is a stretch-3 compact routing scheme for $\mathcal{A}$
- Basically, we can generalize landmark-based compact routing schemes to regular algebras
- Local memory is $O(n^{2/3})$ (due to Cowen) or $O(n^{1/2})$ (Thorup and Zwick)

Algebraic compact routing

- What can we say when regularity fails?
- Turns out, when isotonicity fails in a particular way no compact implementation with constant stretch exists
- **Theorem:** if a monotone algebra $\mathcal{A}$ contains a set of $p \geq 2$ weights $\{w_1, w_2, \ldots, w_p\}$:

  $\forall i, j \in \{1, \ldots, p\}, i \neq j : \; w_i \oplus w_j \succ w_i^{2k}$ and $w_i \oplus w_j \succ w_j^{2k}$

  for some $k \geq 1$, then there is no stretch-k routing scheme with sublinear memory requirement
- **Corollary:** no constant stretch compact routing scheme exists for shortest-widest path routing

Conclusions

- Scaling limits for shortest path routing well researched
- But shortest path routing does not matter in the Internet
- The fundamental scalability of policy routing has so far not been studied
- We took the first steps in this direction
  - gave an algebraic compressibility characterization
  - identified the algebraic requirements for effectively trading between path preference and memory
  - classified most practically relevant policies
- The main message is that with regularity, we are on the safe side, but when regularity fails, we might not even hope for a constant stretch compact routing algorithm
- Are there any theoretical reasons behind the recent explosion of routing table sizes? Answer seems affirmative