





















The authors are supported by EPSRC (UK) grant EP/F002718/1. We would also like to thank Randy Bush, Ken Calvert, Pierre François, Jennifer Rexford, Michael Schapira, Philip Taylor and Laurent Vanbever for their comments and advice on this paper.

## REFERENCES

- [1] L. Blunk, J. Damas, F. Parent, and A. Robachevsky. RFC 2012: Routing Policy Specification Language next generation (RPSLNg), 2005.
- [2] B. Carré. *Graphs and Networks*. Oxford University Press, 1979.
- [3] Cisco Systems. Enhanced interior gateway routing protocol. White Paper 16406.
- [4] D. Donaldson and J. A. Weymark. A quasiordering is the intersection of orderings. *Journal of Economic Theory*, 78(2):382–387, 1998.
- [5] B. Dushnik and E. W. Miller. Partially ordered sets. *American Journal of Mathematics*, 63(3):600–610, 1941.
- [6] O. Filip, P. Machek, M. Mares, and O. Zajicek. *BIRD Internet Routing Daemon User's Guide*. <http://bird.network.cz>.
- [7] L. Gao and J. Rexford. Stable internet routing without global coordination. *IEEE/ACM Transactions on Networking*, pages 681–692, December 2001.
- [8] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica. Pathlet routing. In *Proc. SIGCOMM*, pages 111–122, 2009.
- [9] M. Gondran and M. Minoux. *Graphs and Algorithms*. Wiley, 1984.
- [10] M. Gondran and M. Minoux. *Graphs, Dioids, and Semirings : New Models and Algorithms*. Springer, 2008.
- [11] R. Govindan. Time-space tradeoffs in route-server implementation. *Journal of Internetworking: Research and Experience*, 6, 1995.
- [12] R. Govindan, C. Alaettinoglu, K. Varadhan, and D. Estrin. Route servers for inter-domain routing. *Comput. Netw. ISDN Syst.*, 30(12), 1998.
- [13] T. G. Griffin. The stratified shortest-paths problem. In *Proc. COM-SNETS*, 2010.
- [14] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, April 2002.
- [15] A. J. T. Gurney. *Construction and verification of routing algebras*. PhD thesis, University of Cambridge, 2009.
- [16] A. J. T. Gurney and T. G. Griffin. Lexicographic products in metarouting. In *Proc. Inter. Conf. on Network Protocols*, October 2007.
- [17] E. Jasinska and C. Malayer. (Ab)Using route servers, 2010. Presentation at NANOG 48 (<http://www.nanog.org/meetings/nanog48>).
- [18] V. Van den Schrieck, P. François, and O. Bonaventure. BGP add-paths: The scaling/performance tradeoffs. *IEEE Journal on Selected Areas in Communications*, 2010. (in press).
- [19] L. Vanbever, P. François, O. Bonaventure, and J. Rexford. Customized BGP route selection using BGP/MPLS VPNs, 2009. Presentation at Cisco Systems Routing Symposium ([http://inl.info.ucl.ac.be/system/files/Cisco\\_NAG\\_2009\\_ns\\_bgp.pdf](http://inl.info.ucl.ac.be/system/files/Cisco_NAG_2009_ns_bgp.pdf)).
- [20] P. Verkaik, D. Pei, T. Scholl, A. Shaikh, A. C. Snoeren, and J. E. van der Merwe. Wrestling control from BGP: scalable fine-grained route control. In *Proc. USENIX ATC*, pages 1–14, 2007.
- [21] D. Walton, E. Retana, E. Chen, and J. Scudder. Advertisement of multiple paths in BGP, 2009. IETF Internet-Draft (expired). <http://tools.ietf.org/html/draft-walton-bgp-add-paths-06>.
- [22] Y. Wang, I. Avramopoulos, and J. Rexford. Design for configurability: rethinking interdomain routing policies from the ground up. *IEEE Journal on Selected Areas in Communications*, 27(3):336–348, 2009.
- [23] Y. Wang, M. Schapira, and J. Rexford. Neighbor-specific BGP: more flexible routing policies while improving global stability. In *Proc. 11th SIGMETRICS*, pages 217–228, 2009.
- [24] W. Xu and J. Rexford. MIRO: multi-path interdomain routing. In *Proc. SIGCOMM*, pages 171–182, 2006.
- [25] X. Yang, D. Clark, and A. W. Berger. NIRA: a new inter-domain routing architecture. *IEEE/ACM Transactions on Networking*, 15(4):775–788, 2007.
- [26] X. Yang and D. Wetherall. Source selectable path diversity via routing deflections. In *Proc. SIGCOMM*, pages 159–170, 2006.

Various mathematical objects are involved in our modelling of routing. We summarize their standard definitions here.

A *preorder*  $(S, \preceq)$  is a set together with a binary relation that is reflexive ( $x \preceq x$  for all  $x$  in  $S$ ) and transitive (if  $x \preceq y$  and  $y \preceq z$  then  $x \preceq z$ ).

A preorder that is antisymmetric (if  $x \preceq y \preceq x$  then  $x = y$ ) is a *partial order*. If in addition it is total (for all  $x$  and  $y$ , either  $x \preceq y$  or  $y \preceq x$ ) then it is a *total order* or *linear order*.

If  $\preceq$  is a preorder then we write  $\prec$  for the corresponding strict relation:  $x \prec y$  if and only if  $x \preceq y$  and  $\neg(y \preceq x)$ . The equivalence relation  $\sim$  is given by  $x \sim y$  if and only if both  $x \preceq y$  and  $y \preceq x$ .

A *semigroup*  $(S, \oplus)$  is a set with an associative binary operation. It may be commutative ( $x \oplus y = y \oplus x$  for all  $x$  and  $y$ ) or idempotent ( $x = x \oplus x$  for all  $x$ ) but need not be either.

If we have a set  $F$  of functions over  $S$ , then we can combine this with a preorder or a semigroup to make a *preorder transform*  $(S, \preceq, F)$  or a *semigroup transform*  $(S, \oplus, F)$ .

A graph  $G = (V, E)$  can be weighted over a preorder transform by providing

- 1) a function  $s : V \rightarrow S$ , and
- 2) a function  $w : E \rightarrow F$ .

Then, the weight of a path from node  $i$  to node  $j$  in  $G$ , which uses the arcs  $e_1$  through  $e_k$ , can be calculated as

$$(w(e_k) \circ w(e_{k-1}) \circ \cdots \circ w(e_2) \circ w(e_1))(s(i)).$$

The  $s$  function thus supplies an originated value for each node, and each arc function alters this value in some way. The path weight is, however, still a value in  $S$ , which can be compared with other such values via  $\preceq$  as expected.

Graph weightings over semigroup transforms are defined in the same way: the only difference is in whether a binary operator or a relation is used to determine the best weight.

Two orders (indeed, any two relations) on the same set can be combined via intersection. The intersection of  $\preceq_1$  and  $\preceq_2$  is the order  $\preceq$  where

$$x \preceq y \iff x \preceq_1 y \wedge x \preceq_2 y.$$

This operation naturally extends to the intersection of more than two orders.

A function  $f$  over an order  $(S, \preceq)$  is said to be *inflationary* when

$$\forall x \in S : x \preceq f(x).$$

If  $S$  has a unique greatest element  $\top$ , then  $f$  is *strictly inflationary* when

$$\forall x \in S : x \prec f(x) \vee x = f(x) = \top$$

Otherwise, the definition of a strictly inflationary function only requires that  $x \prec f(x)$  for all  $x$  in  $S$ . These definitions extend to sets of functions in the obvious way: a set of functions is (strictly) inflationary if every function in the set is (strictly) inflationary.