L11: Algebraic Path Problems with applications to Internet Routing Lecture 8 Two versions of distributed Bellman-Ford: Distance Vector and Path Vector

Timothy G. Griffin

timothy.griffin@cl.cam.ac.uk
Computer Laboratory
University of Cambridge, UK

Michaelmas Term, 2018

Our basic iterative algorithm for solving equation $\mathbf{L} = \mathbf{A} \triangleright \mathbf{L} \oplus \mathbf{B}$

$$\begin{aligned} \mathbf{A} \triangleright^{\langle 0 \rangle} \mathbf{B} &= \mathbf{B} \\ \mathbf{A} \triangleright^{\langle k+1 \rangle} \mathbf{B} &= \mathbf{A} \triangleright (\mathbf{A} \triangleright^{\langle k \rangle} \mathbf{B}) \oplus \mathbf{B} \end{aligned}$$

A closer look ...

$$(\mathbf{A} \triangleright^{\langle k+1 \rangle} \mathbf{B})(i,\ d) \quad = \quad \mathbf{B}(i,\ d) \oplus \bigoplus_{(i,u) \in E} \mathbf{A}(i,\ u) \triangleright (\mathbf{A} \triangleright^{\langle k \rangle} \mathbf{B})(u,\ d)$$

This is the basis of distributed Bellman-Ford algorithms (as in RIP and BGP) — a node i computes routes to a destination d by applying its “policies” $\mathbf{A}(i,\ \_)$ to the routes learned from its immediate neighbors.

What if we start iteration in an arbitrary state?

Suppose that we have solved (via iteration) the equation

$$\mathbf{L}_{\mathrm{old}} = \mathbf{A}_{\mathrm{old}} \rhd \mathbf{L}_{\mathrm{old}} \oplus \mathbf{B}_{\mathrm{old}}$$

and then there is a change in the toplopy from $\mathbf{A}_{\mathrm{old}},\ \mathbf{B}_{\mathrm{old}}$ to $\mathbf{A}_{\mathrm{new}},\ \mathbf{B}_{\mathrm{new}}$ and then the iteration continues starting at the "stale" state $\mathbf{L}_{\mathrm{old}}$.

$$\begin{aligned} \mathbf{A}_{\mathrm{new}} \rhd^{\langle 0 \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}} &= \mathbf{L}_{\mathrm{old}} \\ \mathbf{A}_{\mathrm{new}} \rhd^{\langle k+1 \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}} &= \mathbf{A}_{\mathrm{new}} \rhd (\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}}) \oplus \mathbf{B}_{\mathrm{new}} \end{aligned}$$

This represents a simplified (synchronous) model of what happens in the (asynchronous) real-world of routing — routing protocols "iterate forever" while the topology changes.

What if we start iteration in an arbitrary state?

Theorem

For $1 \leqslant k$,

$$\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}} = (\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}}) \oplus (\mathbf{A}_{\mathrm{new}} \rhd^{\langle k-1 \rangle} \mathbf{B}_{\mathrm{new}})$$

If the new system is q-stable and $q < k$, then

$$\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}} = (\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}}) \oplus (\mathbf{A}_{\mathrm{new}} \rhd^{*} \mathbf{B}_{\mathrm{new}})$$

Big problem

Why should $\mathbf{A}_{\mathrm{new}} \rhd^{*} \mathbf{B}_{\mathrm{new}}$ win?

RIP-like (“distance vector”) example (see RFC 1058)

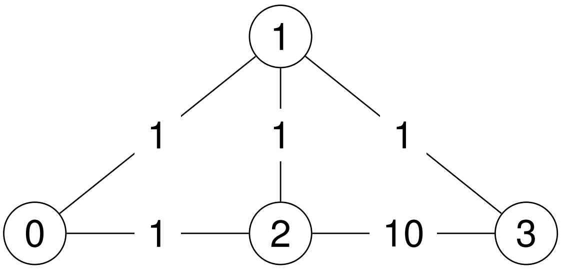

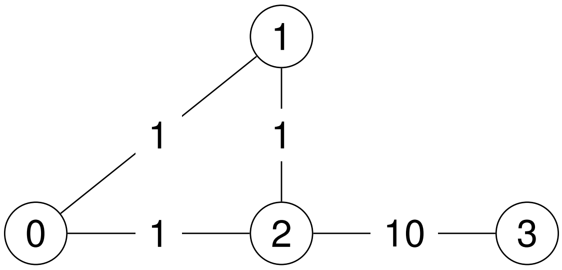

Adjacency matrix $\mathbf{A}_{\mathrm{old}}$

$$\begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & \infty & 1 & 1 & \infty \\ 1 & 1 & \infty & 1 & 1 \\ 2 & 1 & 1 & \infty & 10 \\ 3 & \infty & 1 & 10 & \infty \end{array}$$

Adjacency matrix $\mathbf{A}_{\mathrm{new}}$

$$\begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & \infty & 1 & 1 & \infty \\ 1 & 1 & \infty & 1 & \infty \\ 2 & 1 & 1 & \infty & 10 \\ 3 & \infty & \infty & 10 & \infty \end{array}$$

Using $\mathrm{AddZero}(\infty,\ (\mathbb{N},\ \min,\ +))$ but ignoring inl and inr

RIP-like example — counting to convergence (2)

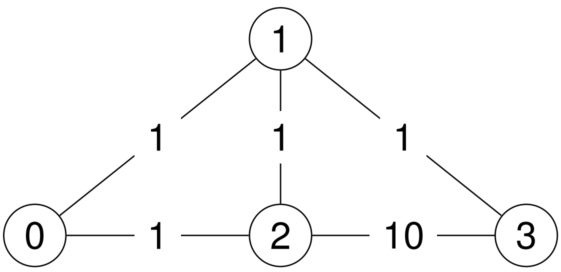

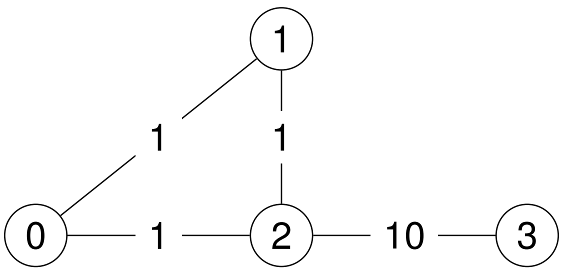

The solution $\mathbf{A}^*_{\mathrm{old}}$

$$\begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 2 \\ 1 & 1 & 0 & 1 & 1 \\ 2 & 1 & 1 & 0 & 2 \\ 3 & 2 & 1 & 2 & 0 \end{array}$$

The solution $\mathbf{A}^*_{\mathrm{new}}$

$$\begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 11 \\ 1 & 1 & 0 & 1 & 11 \\ 2 & 1 & 1 & 0 & 10 \\ 3 & 11 & 11 & 10 & 0 \end{array}$$

RIP-like example — counting to convergence (3)

The scenario: we arrived at state $\mathbf{A}^*_{\text{old}}$, but then links $\{(1,3),\ (3,1)\}$ fail. So we start iterating using the new matrix $\mathbf{A}_{\text{new}}$.

Let $\mathbf{N}_K$ represent $\mathbf{A}_{\text{new}} \rhd^{\langle k \rangle}_{\text{new}} \mathbf{I}$ starting in state $\mathbf{A}^*_{\text{old}}$.

RIP-like example — counting to convergence (4)

$$\mathbf{N}_0 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 2 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 2 \\ 2 & 1 & 2 & 0 \end{array}\right] \end{array} \qquad \mathbf{N}_3 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 4 \\ 1 & 0 & 1 & 4 \\ 1 & 1 & 0 & 4 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

$$\mathbf{N}_1 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 2 \\ 1 & 0 & 1 & 3 \\ 1 & 1 & 0 & 2 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array} \qquad \mathbf{N}_4 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 5 \\ 1 & 0 & 1 & 5 \\ 1 & 1 & 0 & 5 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

$$\mathbf{B}_2 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 3 \\ 1 & 0 & 1 & 3 \\ 1 & 1 & 0 & 3 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array} \qquad \mathbf{N}_5 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 6 \\ 1 & 0 & 1 & 6 \\ 1 & 1 & 0 & 6 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

RIP-like example — counting to convergence (5)

$$\mathbf{N}_6 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 7 \\ 1 & 0 & 1 & 7 \\ 1 & 1 & 0 & 7 \\ 2 & 1 & 2 & 0 \end{array}\right] \end{array} \qquad \mathbf{N}_9 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 10 \\ 1 & 0 & 1 & 10 \\ 1 & 1 & 0 & 10 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

$$\mathbf{N}_7 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 8 \\ 1 & 0 & 1 & 8 \\ 1 & 1 & 0 & 8 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array} \qquad \mathbf{N}_{10} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 11 \\ 1 & 0 & 1 & 11 \\ 1 & 1 & 0 & 10 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

$$\mathbf{N}_8 = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 9 \\ 1 & 0 & 1 & 9 \\ 1 & 1 & 0 & 9 \\ 11 & 11 & 10 & 0 \end{array}\right] \end{array}$$

RIP-like example — counting to infinity (1)

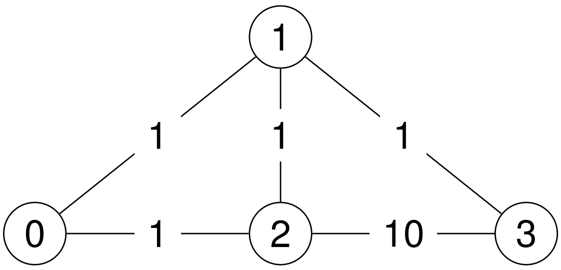

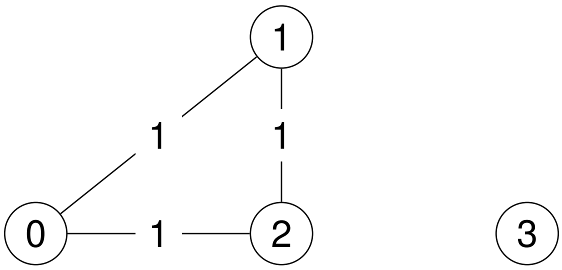

The solution $\mathbf{A}^*_{\text{old}}$ The solution $\mathbf{A}^*_{\text{new}}$

$$\begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & 2 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 2 \\ 2 & 1 & 2 & 0 \end{array}\right] \end{array} \qquad \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \end{array} \begin{array}{c} \begin{array}{cccc} 0 & 1 & 2 & 3 \end{array} \\ \left[\begin{array}{cccc} 0 & 1 & 1 & \infty \\ 1 & 0 & 1 & \infty \\ 1 & 1 & 0 & \infty \\ \infty & \infty & \infty & 0 \end{array}\right] \end{array}$$

Now let $\mathbf{N}_K$ represent $\mathbf{A}_{\text{new}} \rhd^{\langle k \rangle}_{\text{new}} \mathbf{I}$ starting in state $\mathbf{A}^*_{\text{old}}$.

RIP-like example — counting to infinity (2)

$$\mathbf{N}_0 = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 2 \\ 1 & 1 & 0 & 1 & 1 \\ 2 & 1 & 1 & 0 & 2 \\ 3 & 2 & 1 & 2 & 0 \end{array}$$

$$\mathbf{N}_1 = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 2 \\ 1 & 1 & 0 & 1 & 3 \\ 2 & 1 & 1 & 0 & 2 \\ 3 & \infty & \infty & \infty & 0 \end{array}$$

$$\mathbf{N}_2 = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 3 \\ 1 & 1 & 0 & 1 & 3 \\ 2 & 1 & 1 & 0 & 3 \\ 3 & \infty & \infty & \infty & 0 \end{array}$$

$\vdots \quad \vdots \quad \vdots$

$$\mathbf{N}_{376} = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 377 \\ 1 & 1 & 0 & 1 & 377 \\ 2 & 1 & 1 & 0 & 377 \\ 3 & \infty & \infty & \infty & 0 \end{array}$$

$\vdots \quad \vdots \quad \vdots$

$$\mathbf{N}_{998} = \begin{array}{c|cccc} & 0 & 1 & 2 & 3 \\ \hline 0 & 0 & 1 & 1 & 999 \\ 1 & 1 & 0 & 1 & 999 \\ 2 & 1 & 1 & 0 & 999 \\ 3 & \infty & \infty & \infty & 0 \end{array}$$

$\vdots \quad \vdots \quad \vdots$

RIP-like example — What's going on?

Recall

$$\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}} = (\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}}) \oplus (\mathbf{A}_{\mathrm{new}} \rhd^{*} \mathbf{B}_{\mathrm{new}})$$

For some i and d it may be that ...

- $(\mathbf{A}_{\mathrm{new}} \rhd^{*} \mathbf{B}_{\mathrm{new}})(i,\ d)$ is arrived at very quickly
- but $(\mathbf{A}_{\mathrm{new}} \rhd^{\langle k \rangle}_{\mathrm{new}} \mathbf{B}_{\mathrm{new}})(i,\ d)$ may be better until a very large value of k is reached (counting to convergence)
- or it may always be better (counting to infinity).

Distance vector solution : define a very small ∞

- RIP: $\infty = 16$

The path vector solution

The Border Gateway Protocol (BGP)

BGP exchanges metrics **and** paths. It avoids counting to infinity by throwing away routes that have a loop in the path.

The plan ...

Starting from $(\mathbb{N}, \ \mathit{min}, \ +)$ we will attempt to construct a semiring or IAME (using our lexicographic operators) that has elements of the form $(d, \ X)$, where d is a shortest-path metric and X is a set of paths. Then, by succsessive refinements, we will arrive at a BGP-like solution.