An Algebraic Approach to Internet Routing Lecture 12 Metric-neutral partitions

Timothy G. Griffin

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

Michaelmas Term
2010

Hot off the press ...

Hybrid Link-State, Path-Vector Routing
M. Abdul Alim and Timothy G. Griffin
AINTEC 2010, November 15–17, 2010
Bangkok, Thailand

Motivation

- Corporate networks routing objectives include
  - Fast convergence
  - Minimize resource usage (computation, memory, bandwidth)
  - Least-cost paths
  - Load balancing (by traffic engineering)
  - Fault isolation
  - Scalability
- Link-state routing offers fast convergence, but has very high resource requirements.
- OSPF and IS–IS use link-state with partitioning.
- They use complicated routing metrics (e.g., intra-area and inter-area and L1, L2, L1$\rightarrow$L2, and L2$\rightarrow$L1).
- Breaks global optimality and prevents traffic engineering for load balancing.

An example

- The following graph is labelled with shortest-path $(\mathbb{N}^+, \min, +, 0, \infty)$ algebra

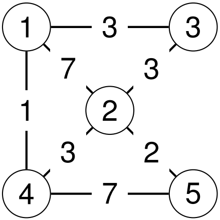

- The adjacency matrix and the all-pair shortest-paths matrix (solution to Equation (**??**)) are

$$\mathbf{A} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ \left[\begin{array}{ccccc} \infty & 7 & 3 & 1 & \infty \\ 7 & \infty & 3 & 3 & 2 \\ 3 & 3 & \infty & \infty & \infty \\ 1 & 3 & \infty & \infty & 7 \\ \infty & 2 & \infty & 7 & \infty \end{array} \right] \end{array} \Rightarrow \mathbf{A}^* = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ \left[\begin{array}{ccccc} 0 & 4 & 3 & 1 & 6 \\ 4 & 0 & 3 & 3 & 2 \\ 3 & 3 & 0 & 4 & 5 \\ 1 & 3 & 4 & 0 & 5 \\ 6 & 2 & 5 & 5 & 0 \end{array} \right] \end{array}.$$

Routing vs. Forwarding

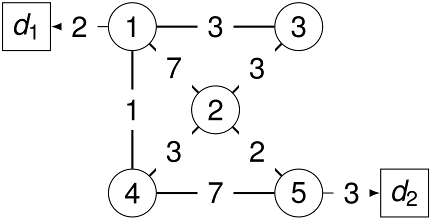

$$\mathbf{M} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cc} d_1 & d_2 \end{array} \\ \begin{bmatrix} 2 & \infty \\ \infty & \infty \\ \infty & \infty \\ \infty & \infty \\ \infty & 3 \end{bmatrix} \end{array} \qquad \mathbf{F} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cc} d_1 & d_2 \end{array} \\ \begin{bmatrix} 2 & 9 \\ 6 & 5 \\ 5 & 8 \\ 3 & 8 \\ 8 & 3 \end{bmatrix} \end{array}$$

- Routing solves $\mathbf{X} = \mathbf{AX} \oplus \mathbf{I}$ with router adjacencies $\mathbf{A}$.
- Forwarding solves $\mathbf{F} = \mathbf{AF} \oplus \mathbf{M}$, where $\mathbf{M}$ is the mapping matrix.
- Link-state: $\mathbf{F} = \mathbf{A}^*\mathbf{M}$.
- Path-vector: $\mathbf{F}^{[0]} = \mathbf{M}$, $\mathbf{F}^{[k+1]} = \mathbf{AF}^{[k]} \oplus \mathbf{M}$.

Partitioning

- $\pi(G) = \{V_1, V_2, \cdots, V_m\}$.
- $\pi(G) \Rightarrow E_{s,t} = \{(u, v) \in E \mid u \in V_s \wedge v \in V_t\}$.
- Region subgraph $G_r = (V_r, E_{r,r})$.
- Border node $u \in V_r$ if $\exists v \in V_{s \neq r} \wedge (u, v) \in E_{r,s}$ and $E_{r,s}$ is inter-region edge.
- Boundary subgraph $G_b = (V_b, E_{r,s})$, may not be connected.
- Adjacency matrix of a partitioned graph

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} & \cdots & \mathbf{A}_{1,m} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} & \cdots & \mathbf{A}_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{m,1} & \mathbf{A}_{m,2} & \cdots & \mathbf{A}_{m,m} \end{bmatrix}.$$

Algebraic Model of AS Partitioning

- The all-region adjacency matrix

$$\mathbf{R} = \begin{bmatrix} \mathbf{A}_{1,1} & \mathbf{O}_{1,2} & \cdots & \mathbf{O}_{1,m} \\ \mathbf{O}_{2,1} & \mathbf{A}_{2,2} & \cdots & \mathbf{O}_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{O}_{m,1} & \mathbf{O}_{m,2} & \cdots & \mathbf{A}_{m,m} \end{bmatrix},$$

  where $\mathbf{O}_{s,t}$ for $s \neq t$ is a sub-matrix of all $\overline{0}$s.

- The routing solution of all the regions

$$\mathbf{R}^* = \begin{bmatrix} \mathbf{A}^*_{1,1} & \mathbf{O}_{1,2} & \cdots & \mathbf{O}_{1,m} \\ \mathbf{O}_{2,1} & \mathbf{A}^*_{2,2} & \cdots & \mathbf{O}_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{O}_{m,1} & \mathbf{O}_{m,2} & \cdots & \mathbf{A}^*_{m,m} \end{bmatrix}.$$

Algebraic Model of AS Partitioning

- The adjacency matrix of of the boundary subgraph

$$\mathbf{B} = \begin{bmatrix} \mathbf{O}_{1,1} & \mathbf{A}_{1,2} & \cdots & \mathbf{A}_{1,m} \\ \mathbf{A}_{2,1} & \mathbf{O}_{2,2} & \cdots & \mathbf{A}_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{m,1} & \mathbf{A}_{m,2} & \cdots & \mathbf{O}_{m,m} \end{bmatrix}.$$

- We can then define the *transit matrix* **T** to capture the weights *virtual links* between border nodes u and v.

$$\mathbf{T}(u,\ v) = \begin{cases} \mathbf{A}^*_{r,r}(u,\ v) & \text{if } (u, v) \text{ is a transit arc in region } r \\ \overline{0} & \text{otherwise} \end{cases}$$

- Define (the adjacency matrix of) the core graph

$$\mathbf{C} = \mathbf{B} \oplus \mathbf{T}.$$

where **T** represents all the virtual links in the graph.

Hybrid Routing?

Main Claim

$$\mathbf{A}^* = \mathbf{R}^* \oplus \mathbf{R}^*\mathbf{C}^*\mathbf{R}^* \quad (1)$$

Interpretation

1. Compute $\mathbf{R}^*$ – solve the region routing problem.
2. Construct the the core graph, matrix $\mathbf{C} = \mathbf{B} \oplus \mathbf{T}$.
3. Compute $\mathbf{C}^*$ – solve the core routing problem.
4. Construct $\mathbf{A}^*$ by computing $\mathbf{R}^*(\mathbf{I} \oplus \mathbf{C}^*\mathbf{R}^*)$.

Note that *different algorithms* – link-state or path-vector – can be used in the 1st and the 3rd steps.

A Partitioning Example

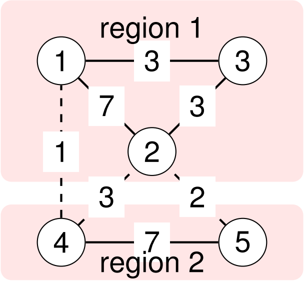

$$\mathbf{A} = \left[\begin{array}{c|c} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} \\ \hline \mathbf{A}_{2,1} & \mathbf{A}_{2,2} \end{array}\right] = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\ \left[\begin{array}{ccc|cc} \infty & 7 & 3 & 1 & \infty \\ 7 & \infty & 3 & 3 & 2 \\ 3 & 3 & \infty & \infty & \infty \\ \hline 1 & 3 & \infty & \infty & 7 \\ \infty & 2 & \infty & 7 & \infty \end{array}\right] \end{array}$$

A Partitioning Example

- The all-region adjacency matrix and corresponding routing solution

$$
\mathbf{R} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array}
\begin{array}{c}
\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\
\left[\begin{array}{ccc|cc}
\infty & 7 & 3 & \infty & \infty \\
7 & \infty & 3 & \infty & \infty \\
3 & 3 & \infty & \infty & \infty \\ \hline
\infty & \infty & \infty & \infty & 7 \\
\infty & \infty & \infty & 7 & \infty
\end{array}\right]
\end{array}
\Rightarrow \mathbf{R}^* = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array}
\begin{array}{c}
\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\
\left[\begin{array}{ccc|cc}
0 & 6 & 3 & \infty & \infty \\
6 & 0 & 3 & \infty & \infty \\
3 & 3 & 0 & \infty & \infty \\ \hline
\infty & \infty & \infty & 0 & 7 \\
\infty & \infty & \infty & 7 & 0
\end{array}\right]
\end{array}
$$

- The boundary matrix and the transit matrix

$$
\mathbf{B} = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array}
\begin{array}{c}
\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \end{array} \\
\left[\begin{array}{ccc|cc}
\infty & \infty & \infty & 1 & \infty \\
\infty & \infty & \infty & 3 & 2 \\
\infty & \infty & \infty & \infty & \infty \\ \hline
1 & 3 & \infty & \infty & \infty \\
\infty & 2 & \infty & \infty & \infty
\end{array}\right]
\end{array}
, \ \mathbf{T} = \begin{array}{c} \\ 1 \\ 2 \\ 4 \\ 5 \end{array}
\begin{array}{c}
\begin{array}{cccc} 1 & 2 & 4 & 5 \end{array} \\
\left[\begin{array}{cccc}
\infty & 6 & \infty & \infty \\
6 & \infty & \infty & \infty \\
\infty & \infty & \infty & 7 \\
\infty & \infty & 7 & \infty
\end{array}\right]
\end{array}
$$

A Partitioning Example

- The core graph with virtual links

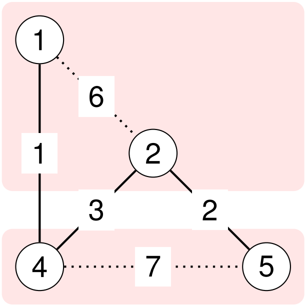

- Core graph adjacency matrix and the routing solution

$$\mathbf{C} = \begin{array}{c} \\ 1 \\ 2 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccc} 1 & 2 & 4 & 5 \end{array} \\ \left[\begin{array}{cccc} \infty & 6 & 1 & \infty \\ 6 & \infty & 3 & 2 \\ 1 & 3 & \infty & 7 \\ \infty & 2 & 7 & \infty \end{array}\right] \end{array} \Rightarrow \mathbf{C}^* = \begin{array}{c} \\ 1 \\ 2 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccc} 1 & 2 & 4 & 5 \end{array} \\ \left[\begin{array}{cccc} 0 & 4 & 1 & 6 \\ 4 & 0 & 3 & 2 \\ 1 & 3 & 0 & 5 \\ 6 & 2 & 5 & 0 \end{array}\right] \end{array}$$

Equations of Hybrid Model

Using the distinctions of routing and forwarding:

1. Solve $\mathbf{F}_1 = \mathbf{R}\mathbf{F}_1 \oplus \mathbf{I}$ for region routing $\mathbf{R}^*$.
2. Build the core graph $\mathbf{C} = \mathbf{B} \oplus \mathbf{T}$.
3. Solve $\mathbf{F}_2 = \mathbf{C}\mathbf{F}_2 \oplus \mathbf{F}_1$ for core routing $\mathbf{C}^*$ and exporting region internal routes to the core $\mathbf{C}^*\mathbf{R}^*$.
4. Solve $\mathbf{F} = \mathbf{F}_1 \oplus \mathbf{F}_1\mathbf{F}_2$ for importing inter-region routes $\mathbf{R}^*(\mathbf{I} \oplus \mathbf{C}^*\mathbf{R}^*)$.

Combinations of link-state and path-vector mechanisms

Hybrid	Region	Core
D-over-D	link-state	link-state
B-over-D	link-state	path-vector
D-over-B	path-vector	link-state
B-over-B	path-vector	path-vector

Problem Set III (Due 1 December)

1. Construct an interesting weighted graph using the *scoped product* $(S \Theta T)$. Show adjacency and routing matrix. A picture might help.
2. Construct an interesting weighted graph using the *metric-neutral partitions* (this lecture). Show adjacency and routing matrix.
3. Prove this : $(()^{*}\mathbf{X} \oplus \mathbf{Y}) = \mathbf{X}^{*}(\mathbf{Y}\mathbf{X}^{*})^{*}$.
4. Prove the **Main Claim** above.