Matrix Factorizations for Computer Networks

David Bindel

Department of Computer Science
Cornell University

30 Apr 2011

A fuzzy picture of the Internet

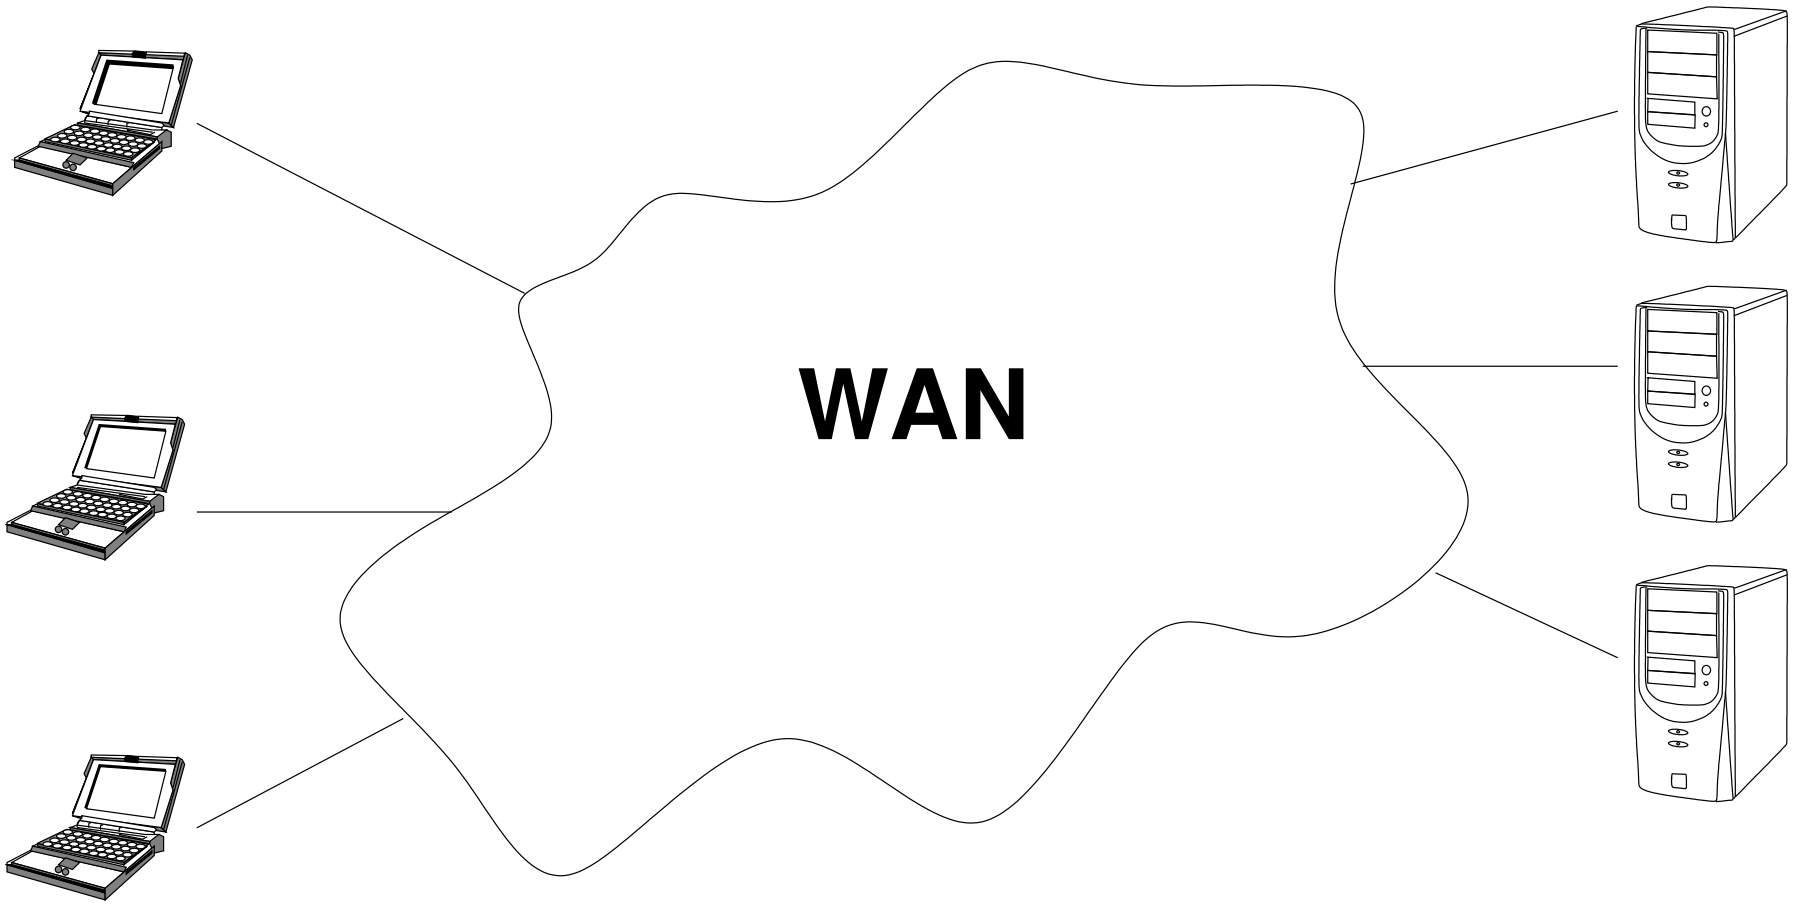

A slightly-less fuzzy picture

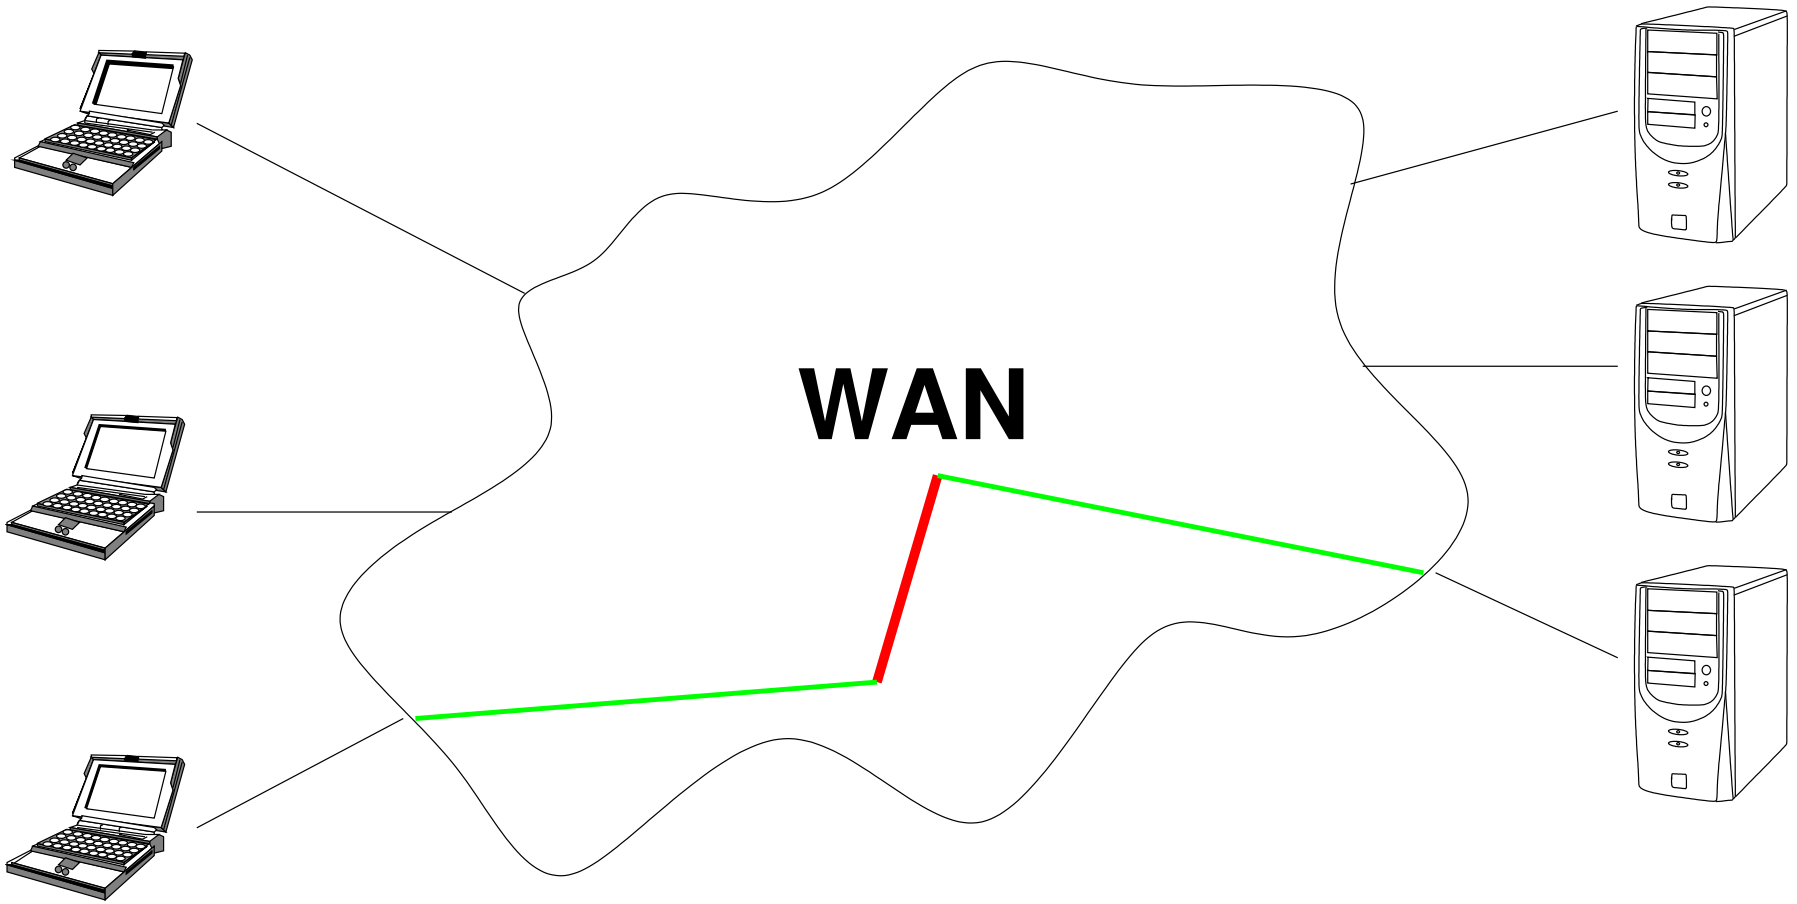

Overlays to the rescue?

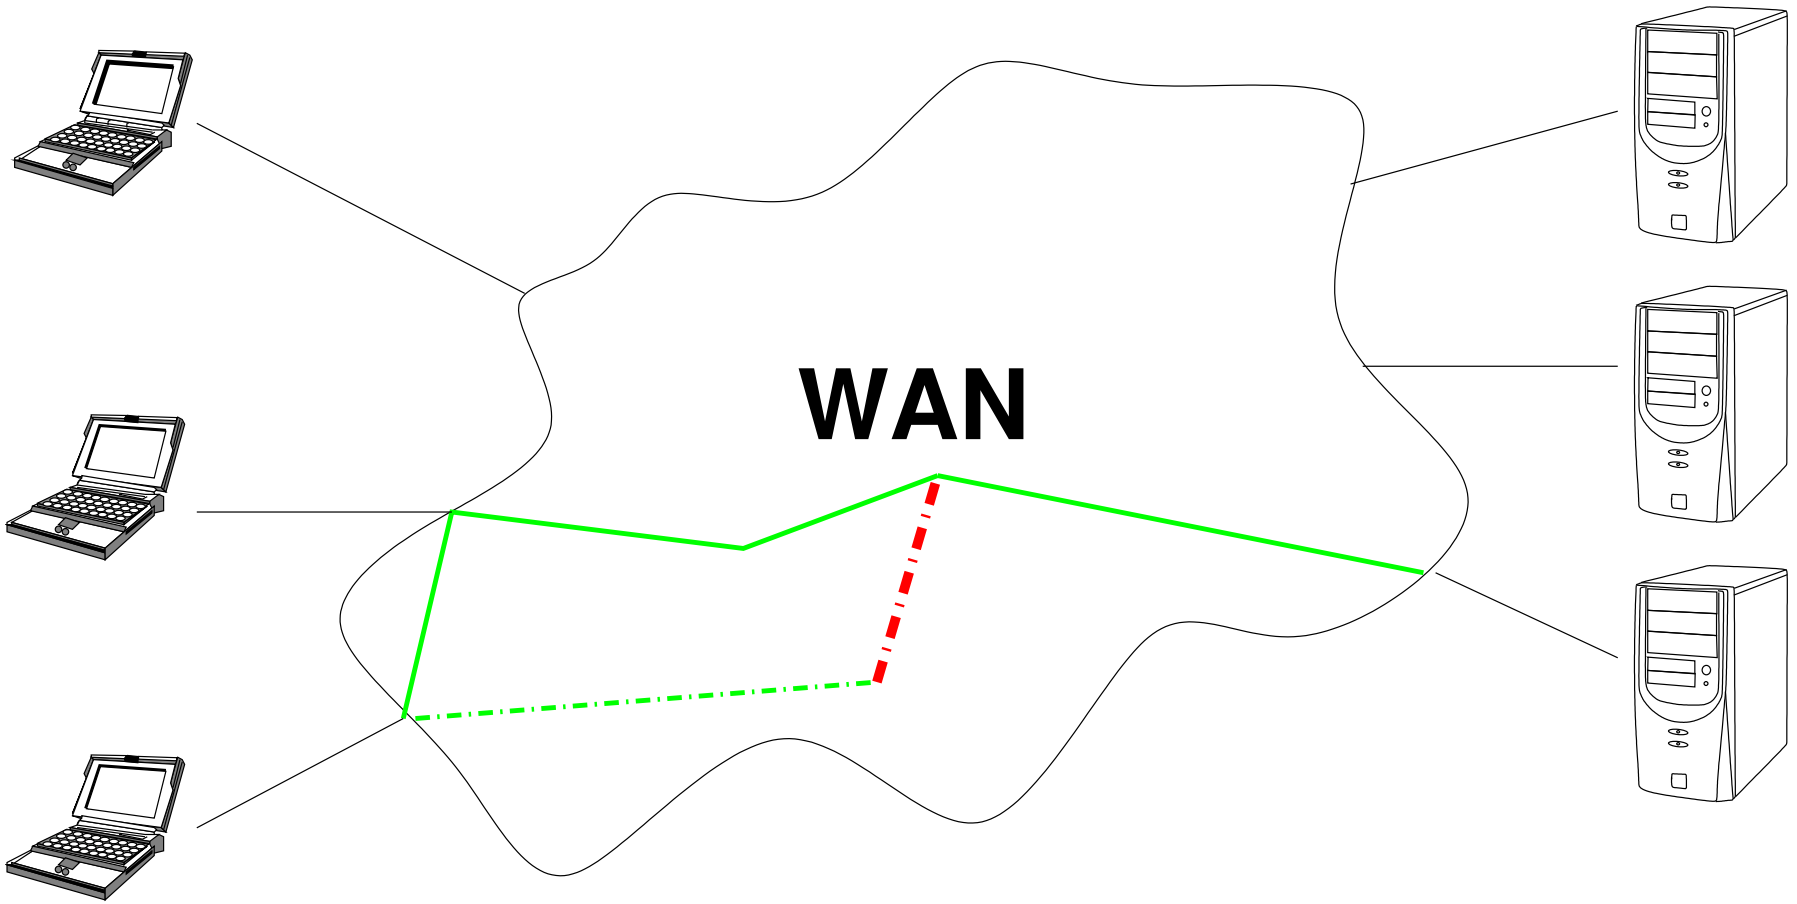

Overlays and measurement

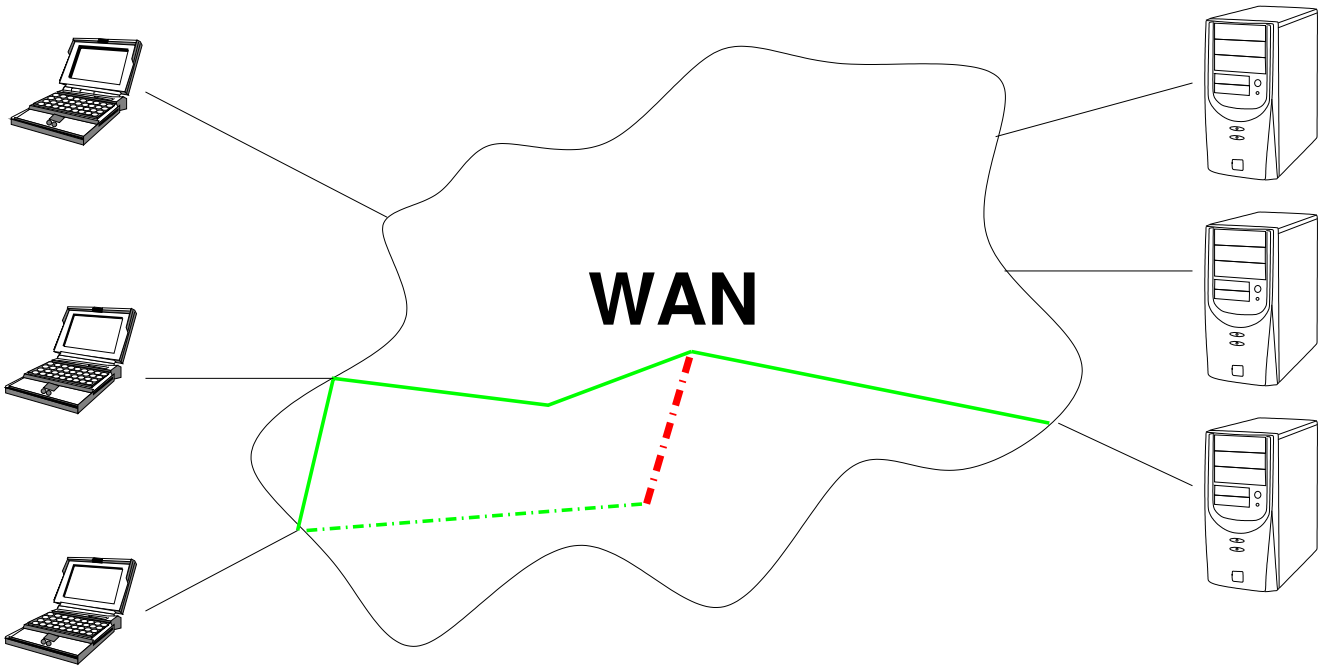

Measure a few paths to infer:

- *Path* properties (ToN 2007, SIGCOMM 2004, IMC 2003)
- *Link* properties (ToN 2009, SIGCOMM 2006)
- Routing topology? (underway)

Discrete Radon transform

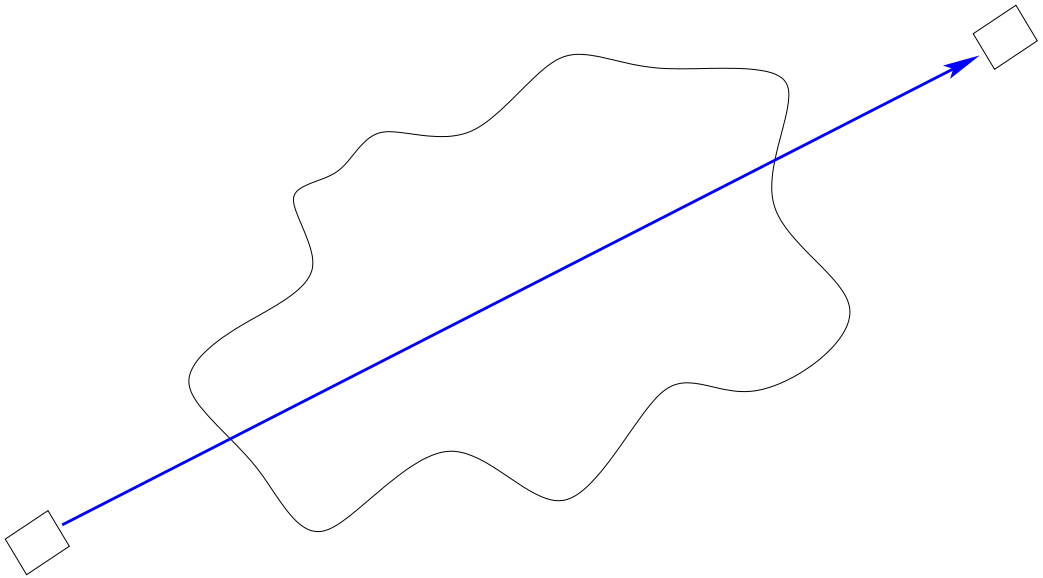

Radon transform:

$$(Ru)(L) = \int_L u(\mathbf{x})\,|d\mathbf{x}|$$

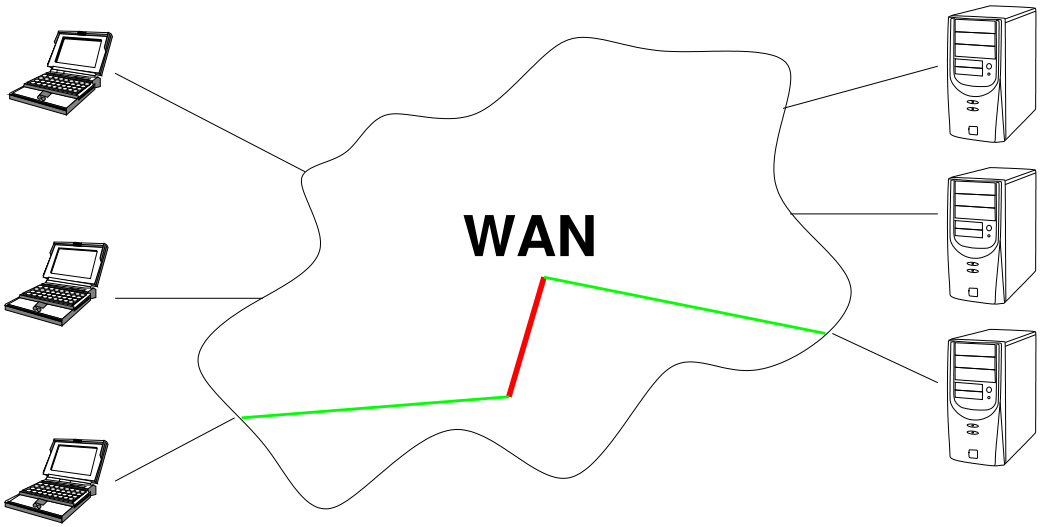

Discrete version:

$$(Gu)_i = \sum_{j \in \text{links}(i)} u_j$$

Path metrics and path matrices

Relate path and link properties via a linear map:

$$Gu = b$$

where

- b_i = property of ith end-to-end path
- u_j = property of link j
- $G_{ij} = \begin{cases} 1 & \text{if path } i \text{ uses link } j \\ 0 & \text{otherwise} \end{cases}$

Wanted: a sparse rank-revealing factorization of G

Network virtualization and column dependencies

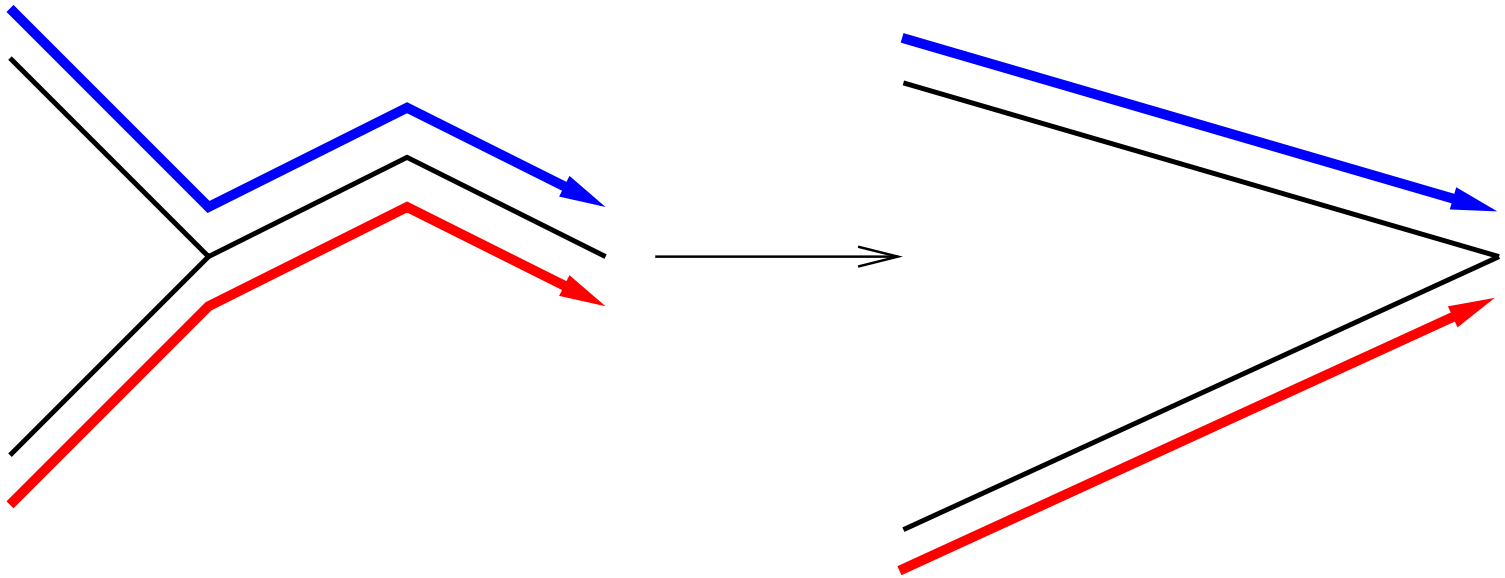

"Fans" == structured linear dependency among columns (links):

$$G(:, \text{fan}) = \begin{bmatrix} c_1 & c_2 & c_1 + c_2 & c_1 + c_2 \end{bmatrix} = \begin{bmatrix} c_1 & c_2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 \end{bmatrix}$$

Factor out a zero-one "virtualization matrix."

Row dependencies and dependent paths

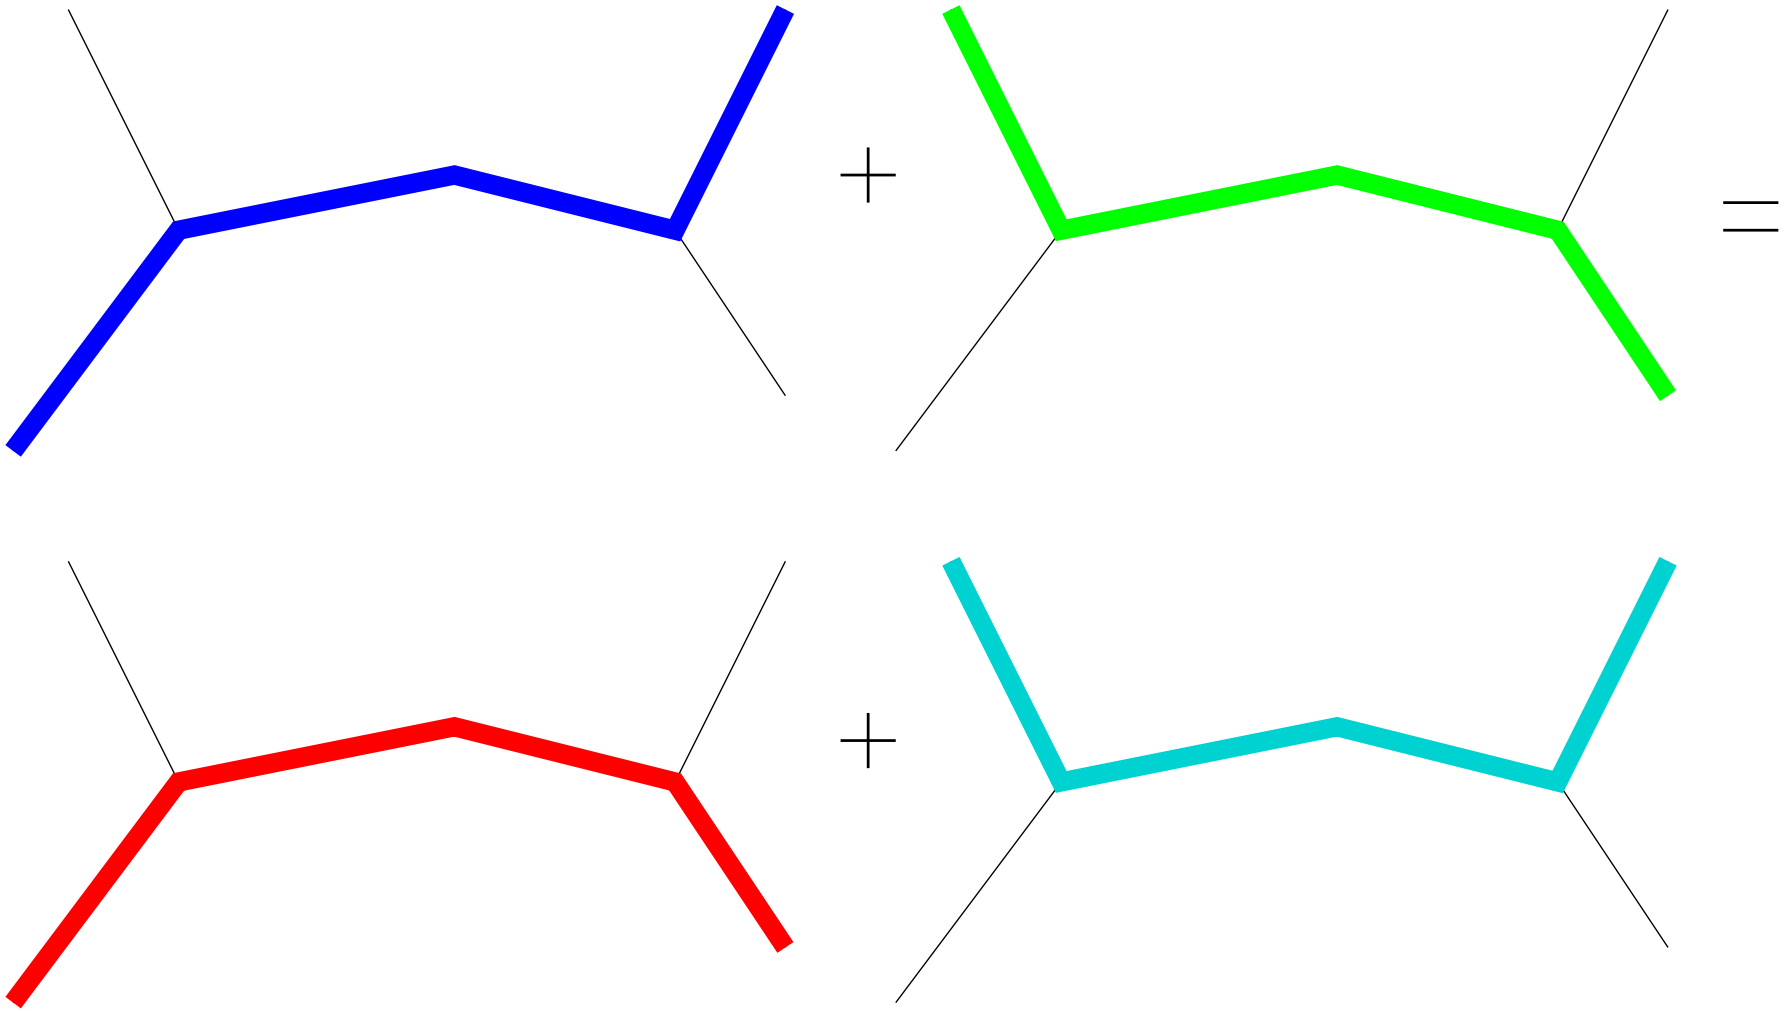

More complicated linear dependencies?

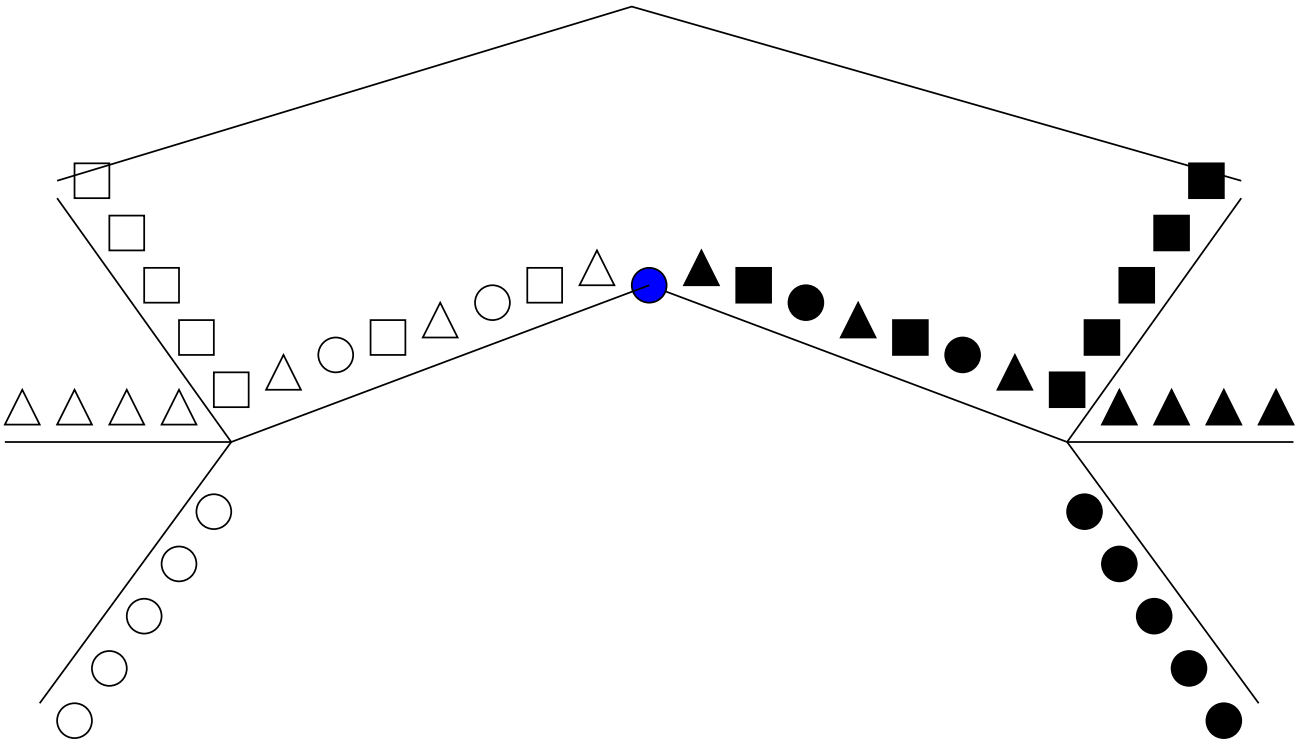

Linear dependencies

$$\square + \blacktriangle$$
$$- \quad \triangle + \blacktriangle$$
$$+ \quad \triangle + \bullet$$
$$= \quad \square + \bullet$$

Junctions and bipartite graphs

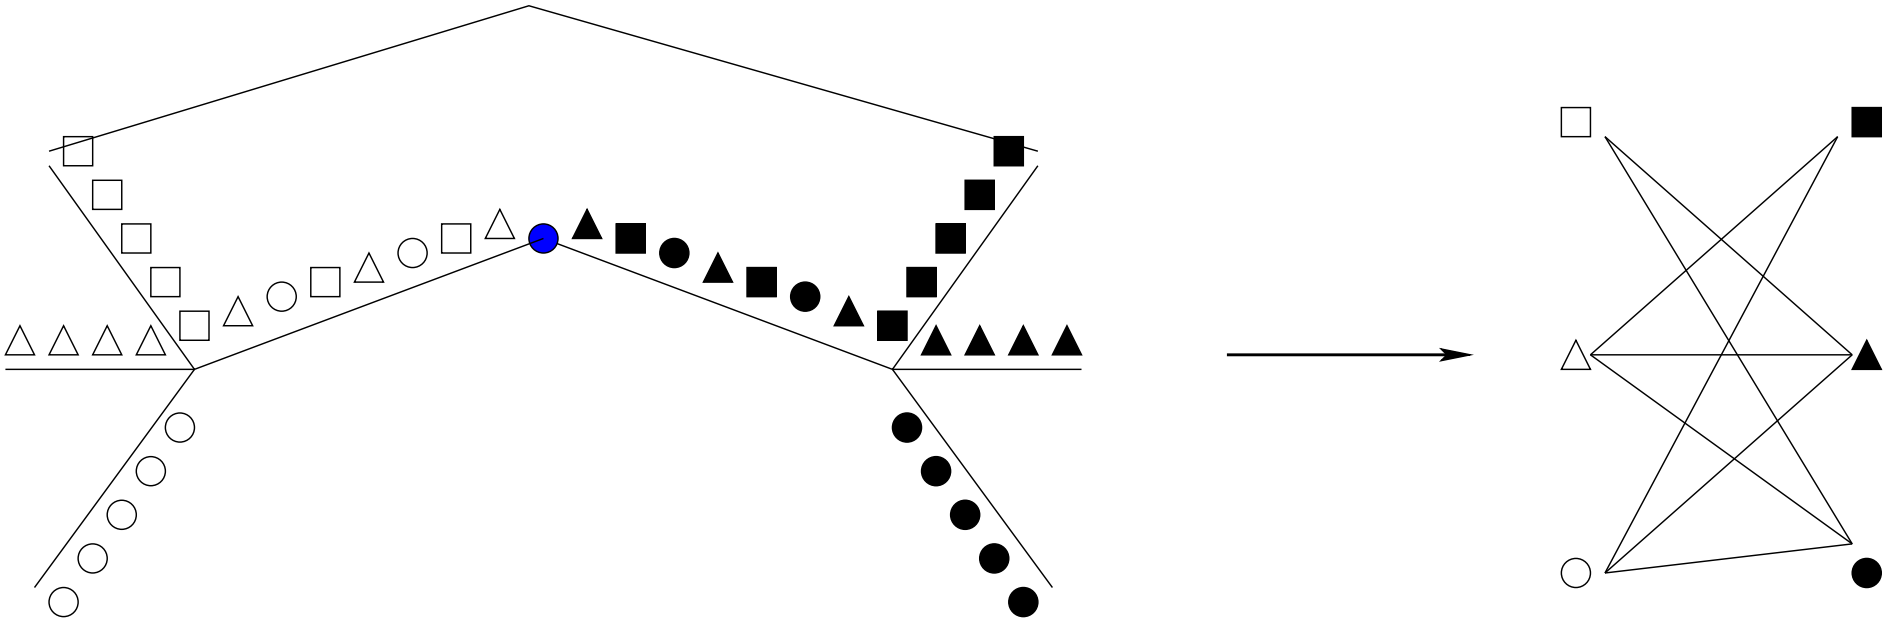

Bipartite router graph at r:

Left node	= segment from source $\rightarrow r$
Right node	= segments $r \rightarrow$ destination
Edge	= path traversing r

Junctions and bipartite graphs

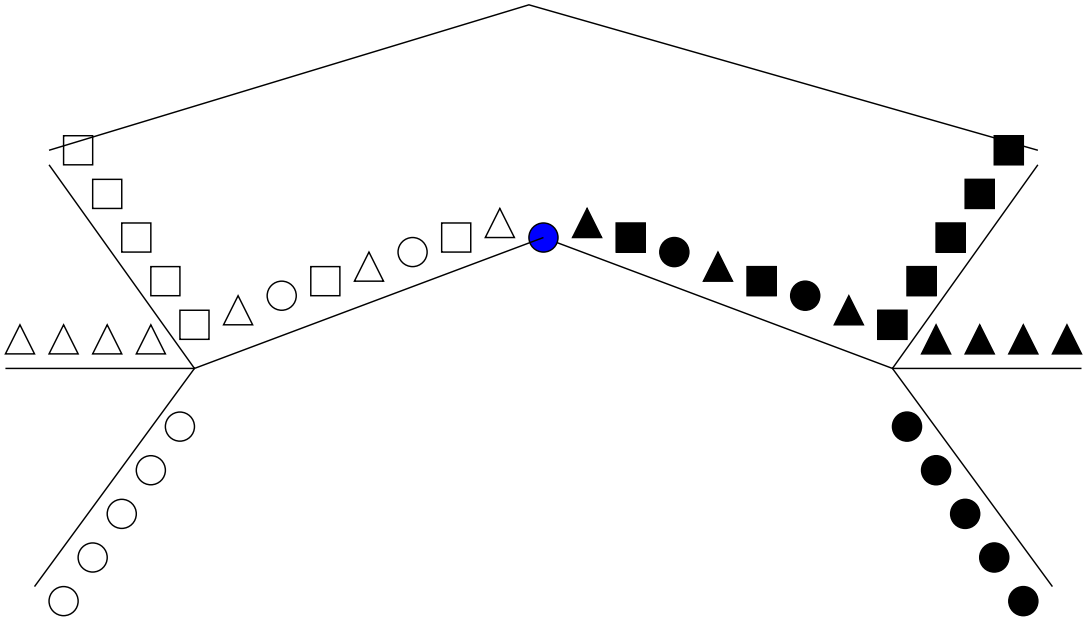

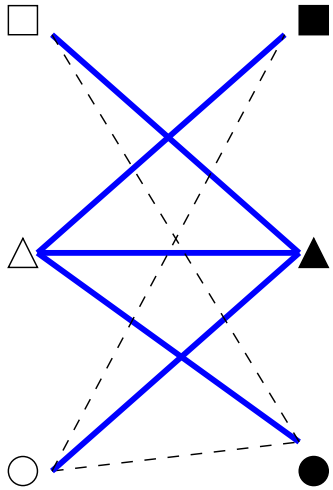

Spanning trees in the router graph

$\Longrightarrow$

spanning sets among path vectors

For more

On network tomography:

- Chen, B., Song, Chavez, Katz. *Algebra-based scalable overlay network monitoring: Algorithms, evaluation, and applications*. ACM ToN, 17(6), 2009.
- Zhao, Chen, B. *Towards unbiased end-to-end network diagnosis*. ACM ToN, 15(5), 2007
- `http://www.cs.cornell.edu/~bindel`
- ... or come talk to me!

On completely different things:

- Micro-machined wineglass gyroscopes (Erdal Yilmaz)
- Resonances and nonlinear eigenvalue problems (Amanda Hood)