Estimating Hop Distance Between Arbitrary Host Pairs

Brian Eriksson
University of Wisconsin - Madison

Paul Barford
Robert Nowak

Overview

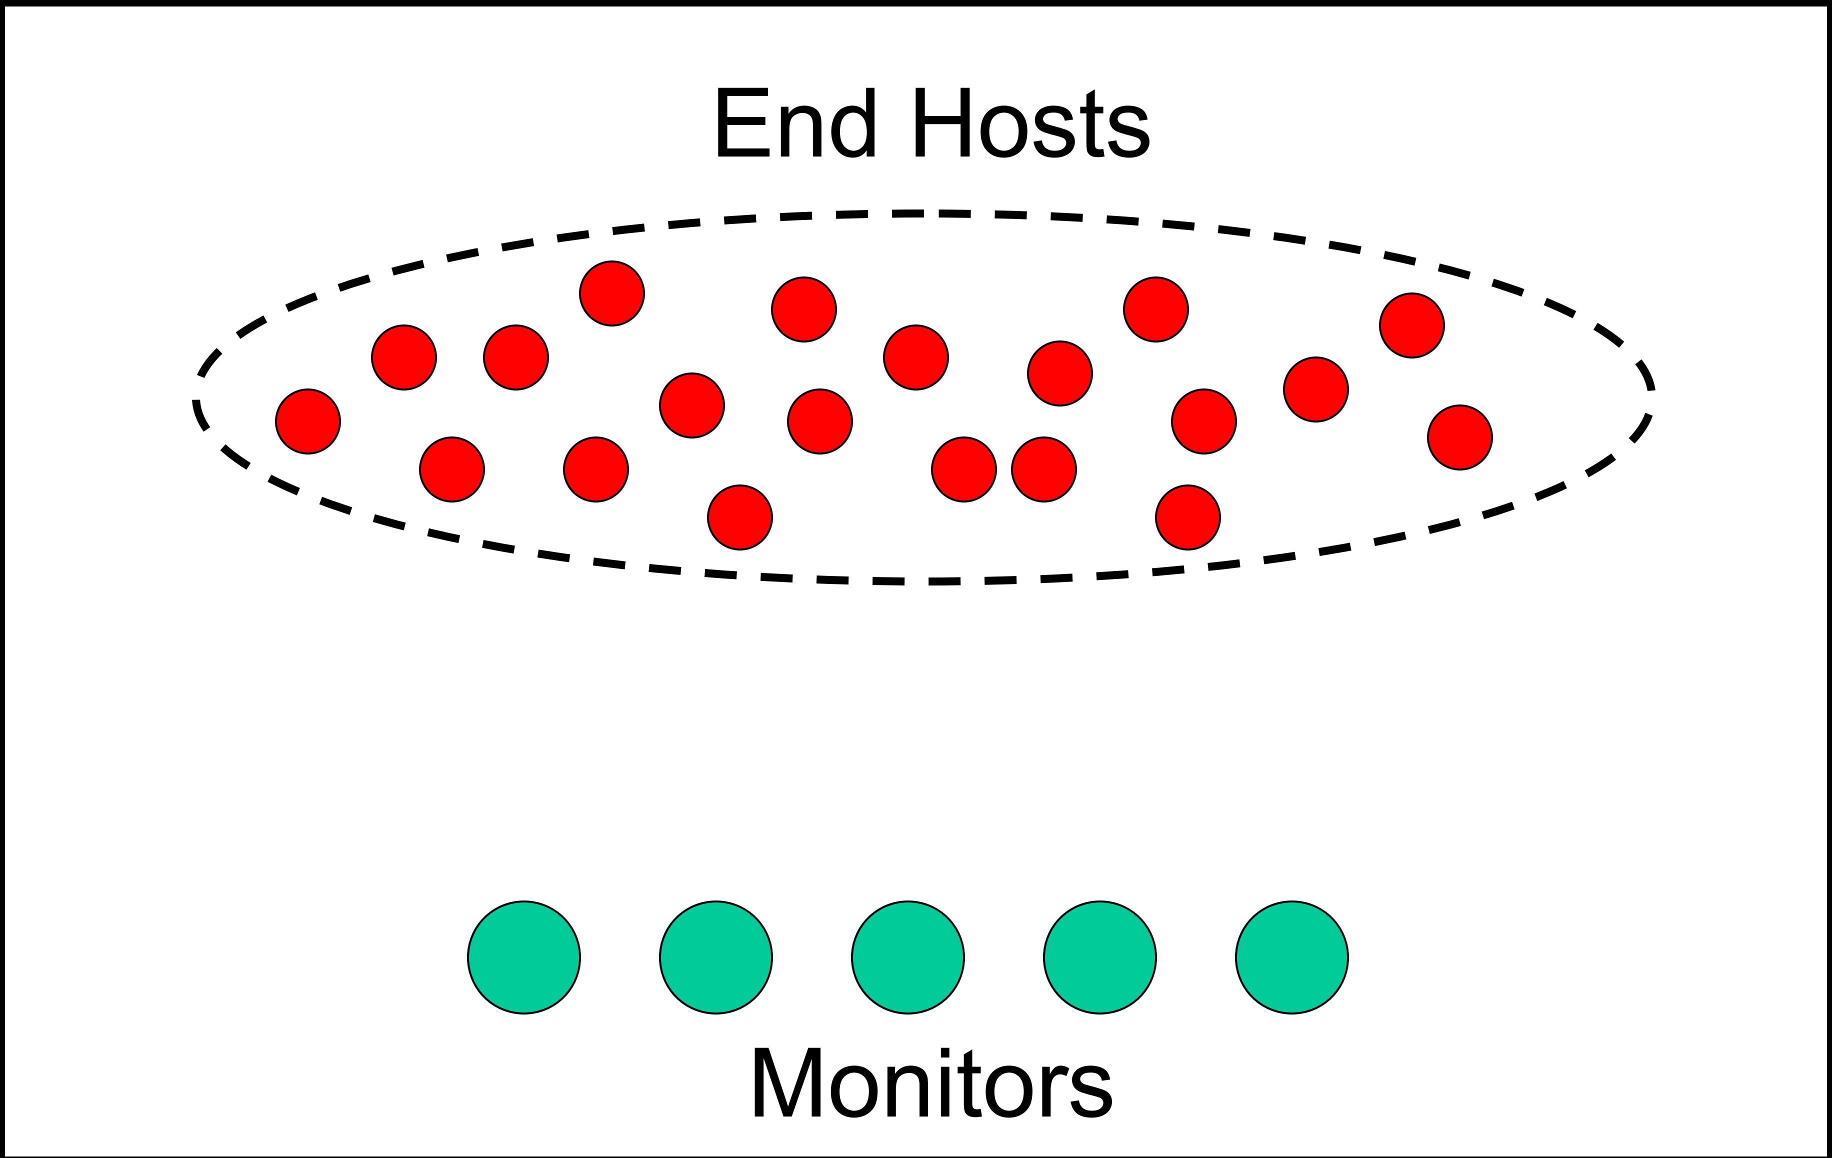

Overview

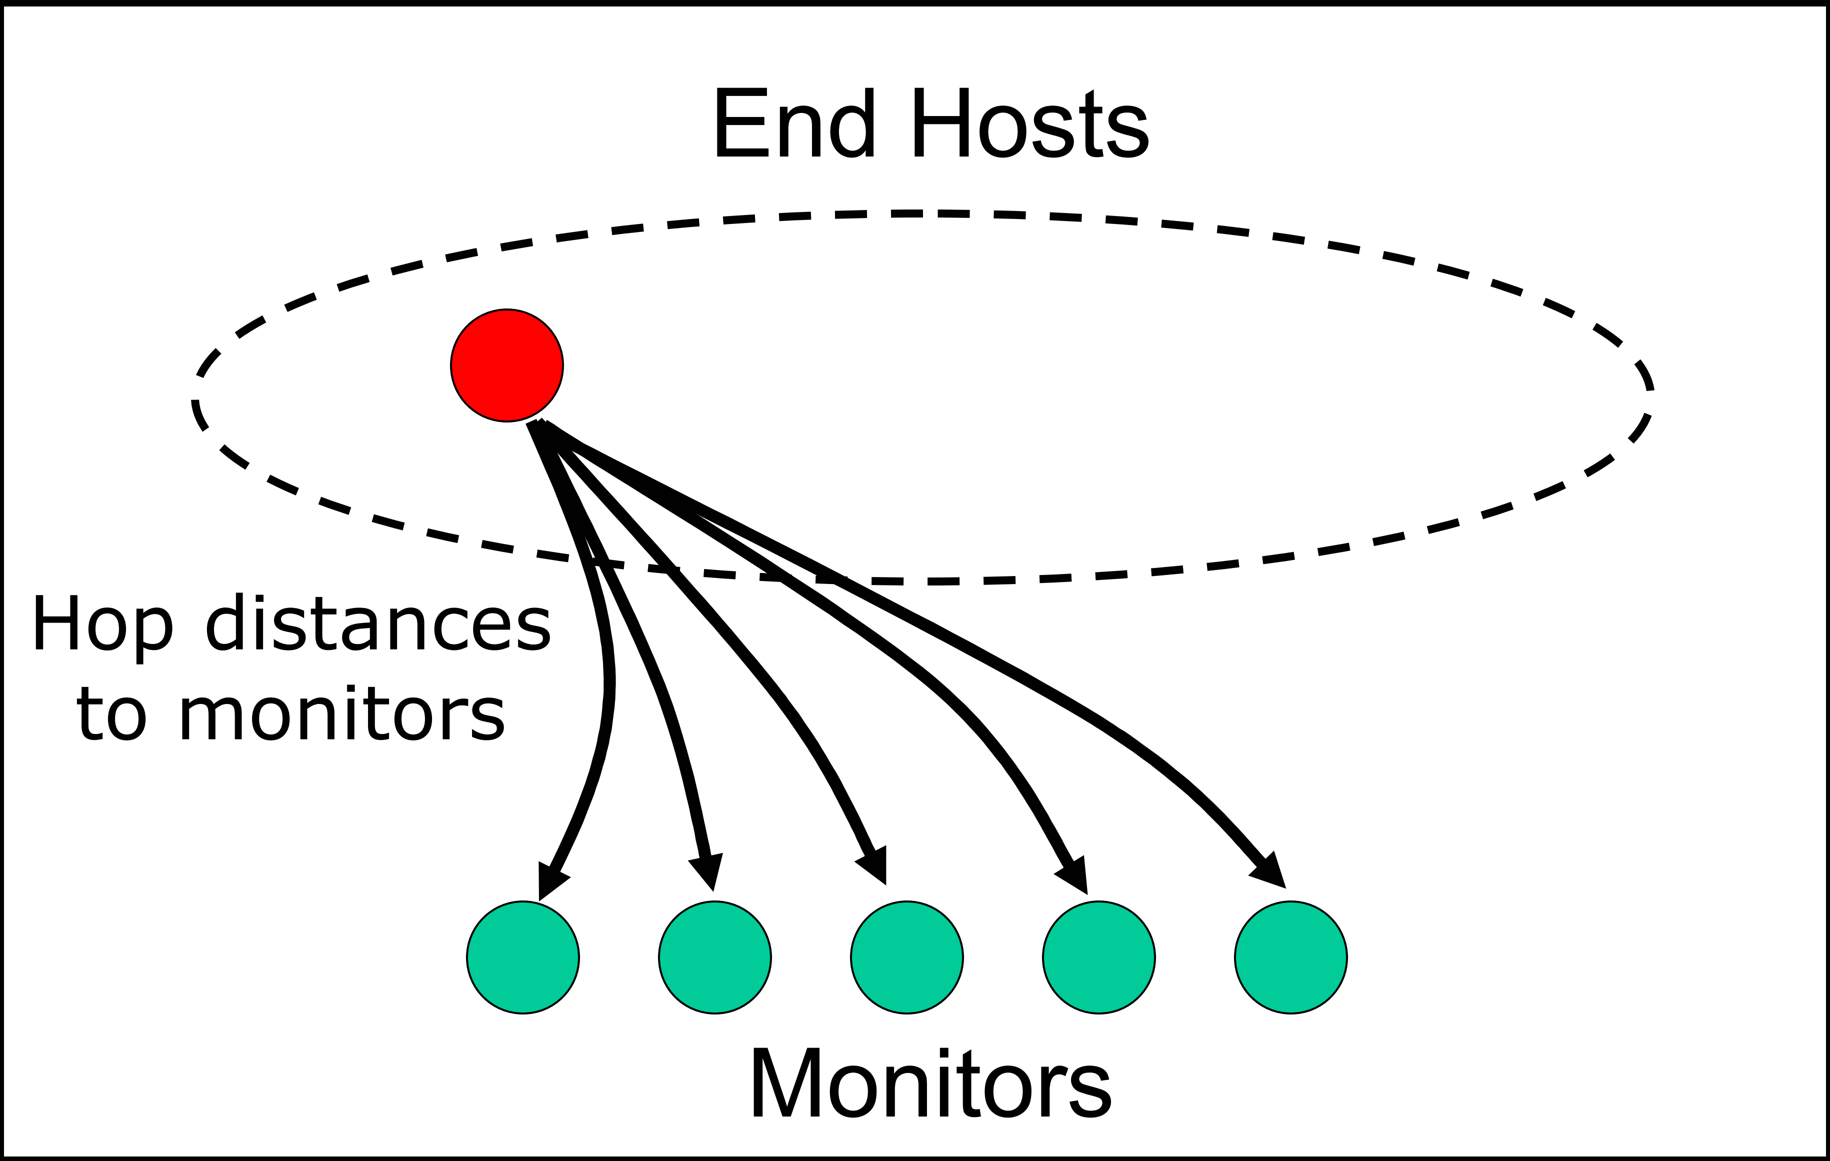

Overview

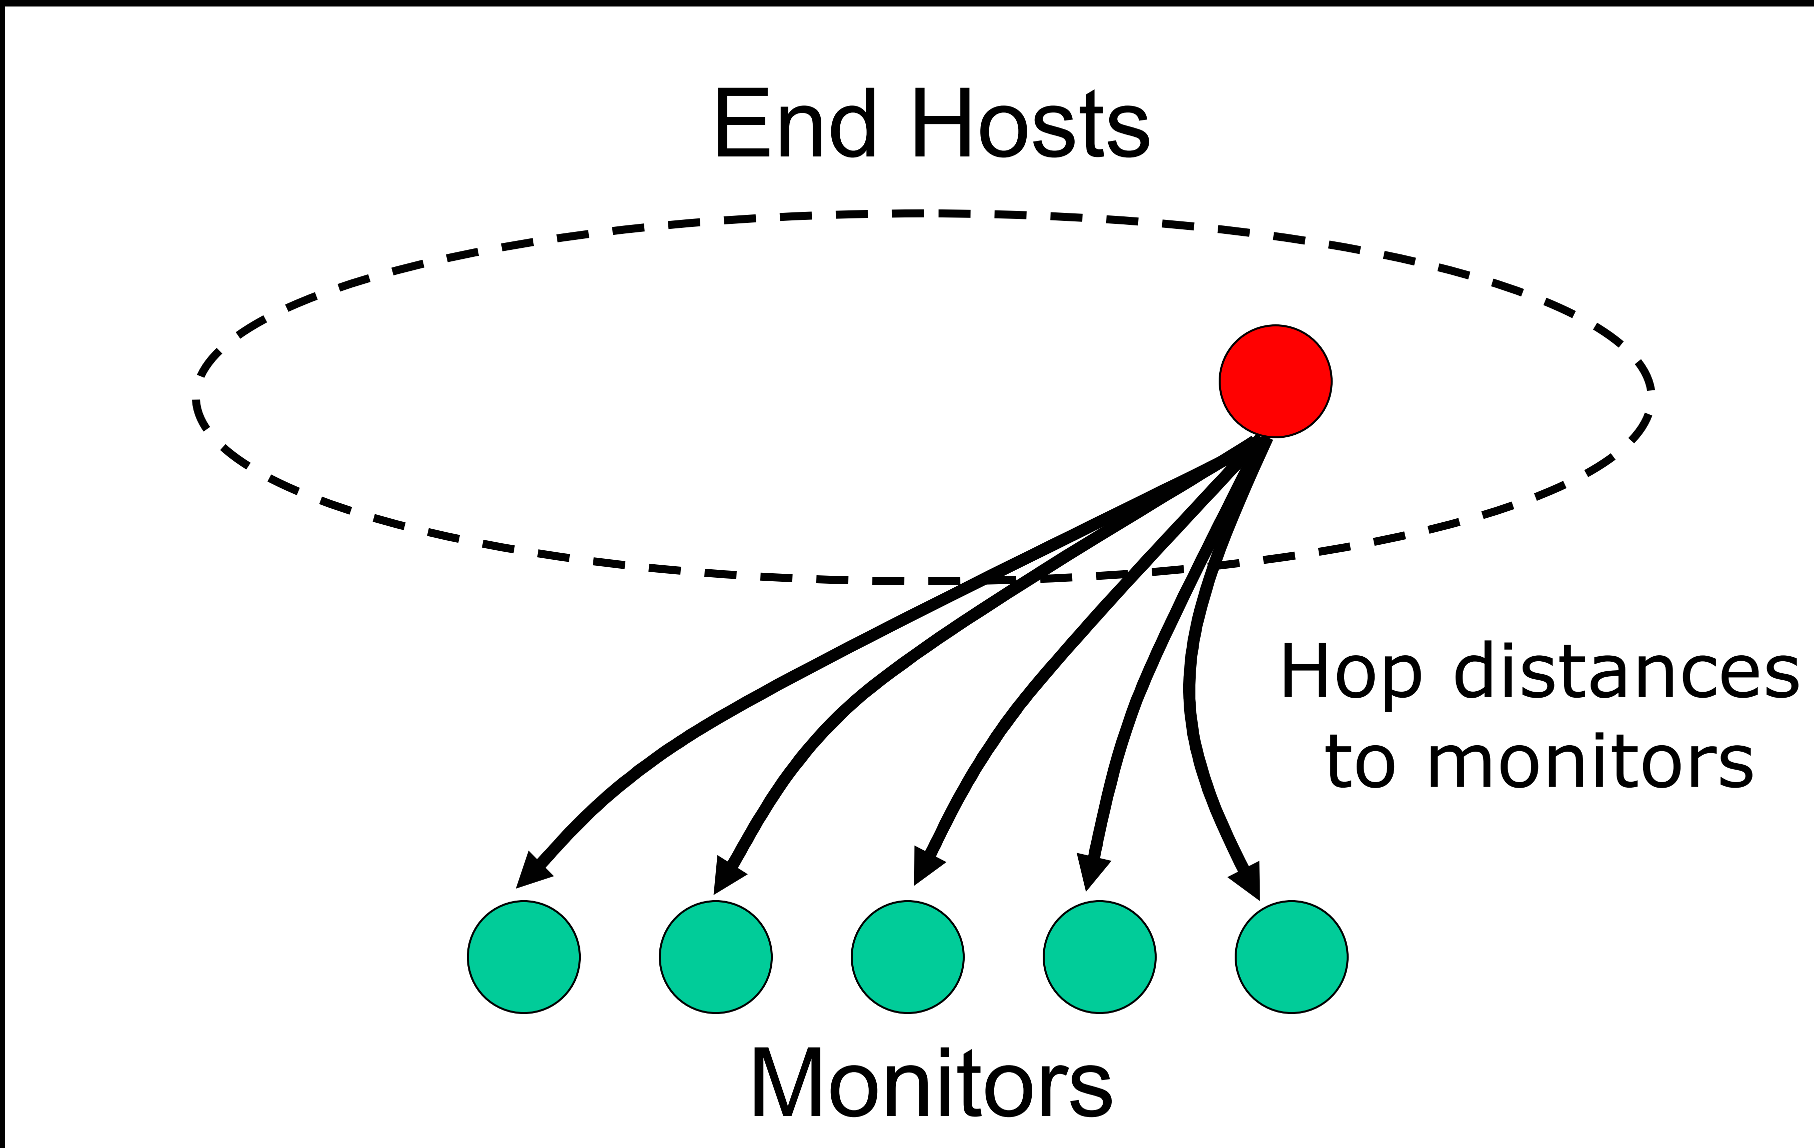

Overview

End Hosts

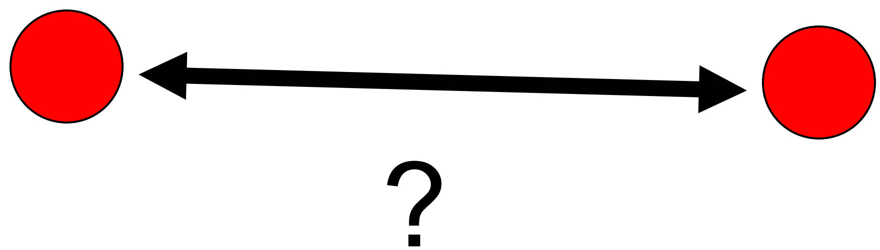

What is the hop distance between these two end hosts?

Overview

- Goal: Estimate hop distances between arbitrary end hosts
  - Topology Estimation
  - Robust Overlay Network Design
  - Performance Analysis

Overview

- Goal: Estimate hop distances between arbitrary end hosts
  - Topology Estimation
  - Robust Overlay Network Design
  - Performance Analysis
- Using passive measurements, we can develop a much lighter-weight approach than previous methods

Overview

- Goal: Estimate hop distances between arbitrary end hosts
  - Topology Estimation
  - Robust Overlay Network Design
  - Performance Analysis
- Using passive measurements, we can develop a much lighter-weight approach than previous methods
- Use additional information (e.g. AS information) to improve estimation

Multidimensional Scaling

Multidimensional Scaling (MDS)

0	$D_{1,2}$	$D_{1,3}$
$D_{2,1}$	0	$D_{2,3}$
$D_{3,1}$	$D_{3,2}$	0

- Given the pairwise distances between a set of nodes

Multidimensional Scaling (MDS)

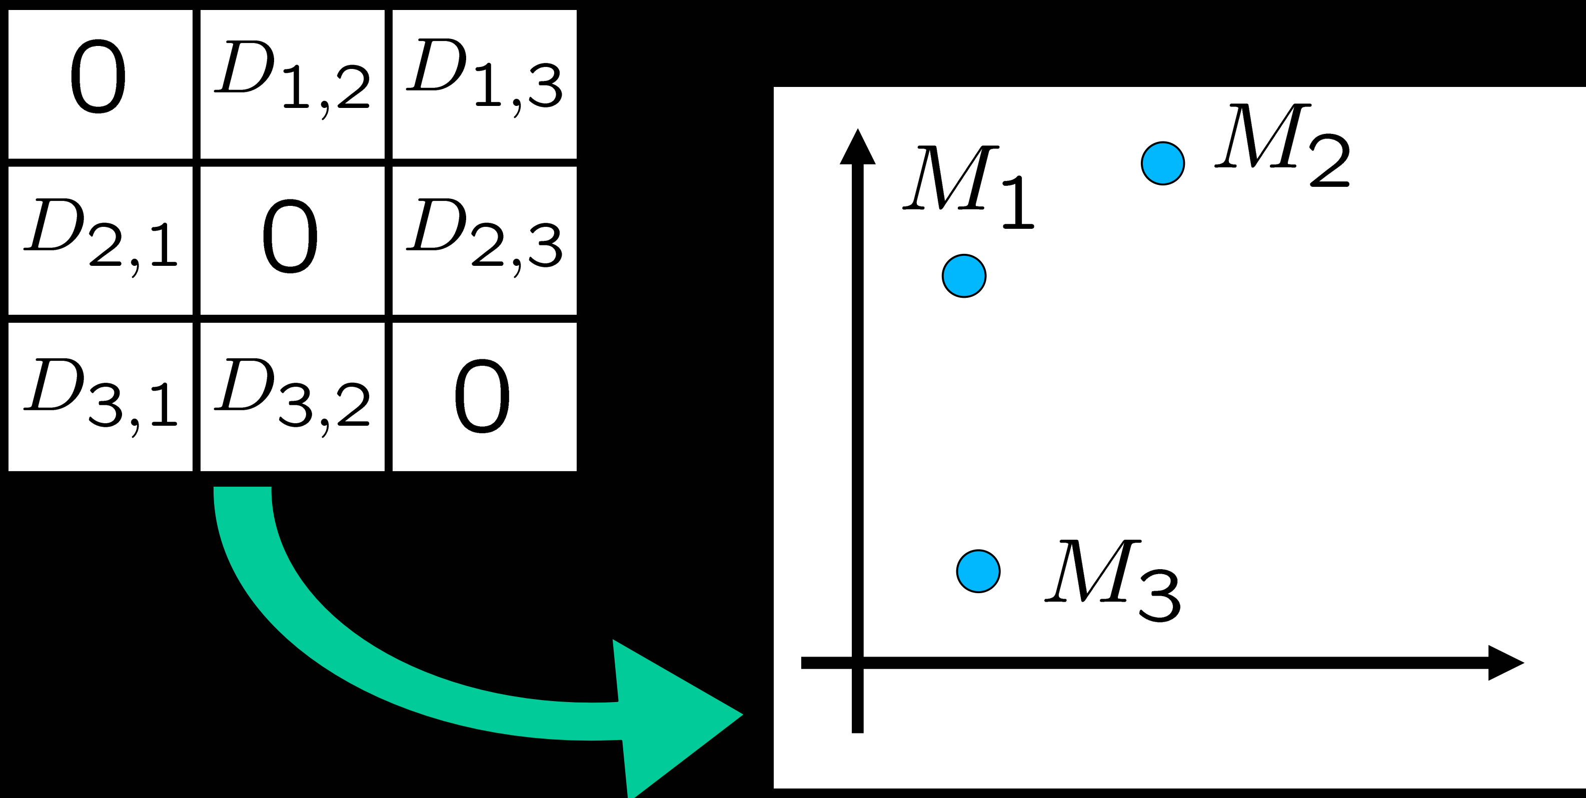

- Multidimensional Scaling (MDS) finds a low-dimensional coordinate mapping…

Multidimensional Scaling (MDS)

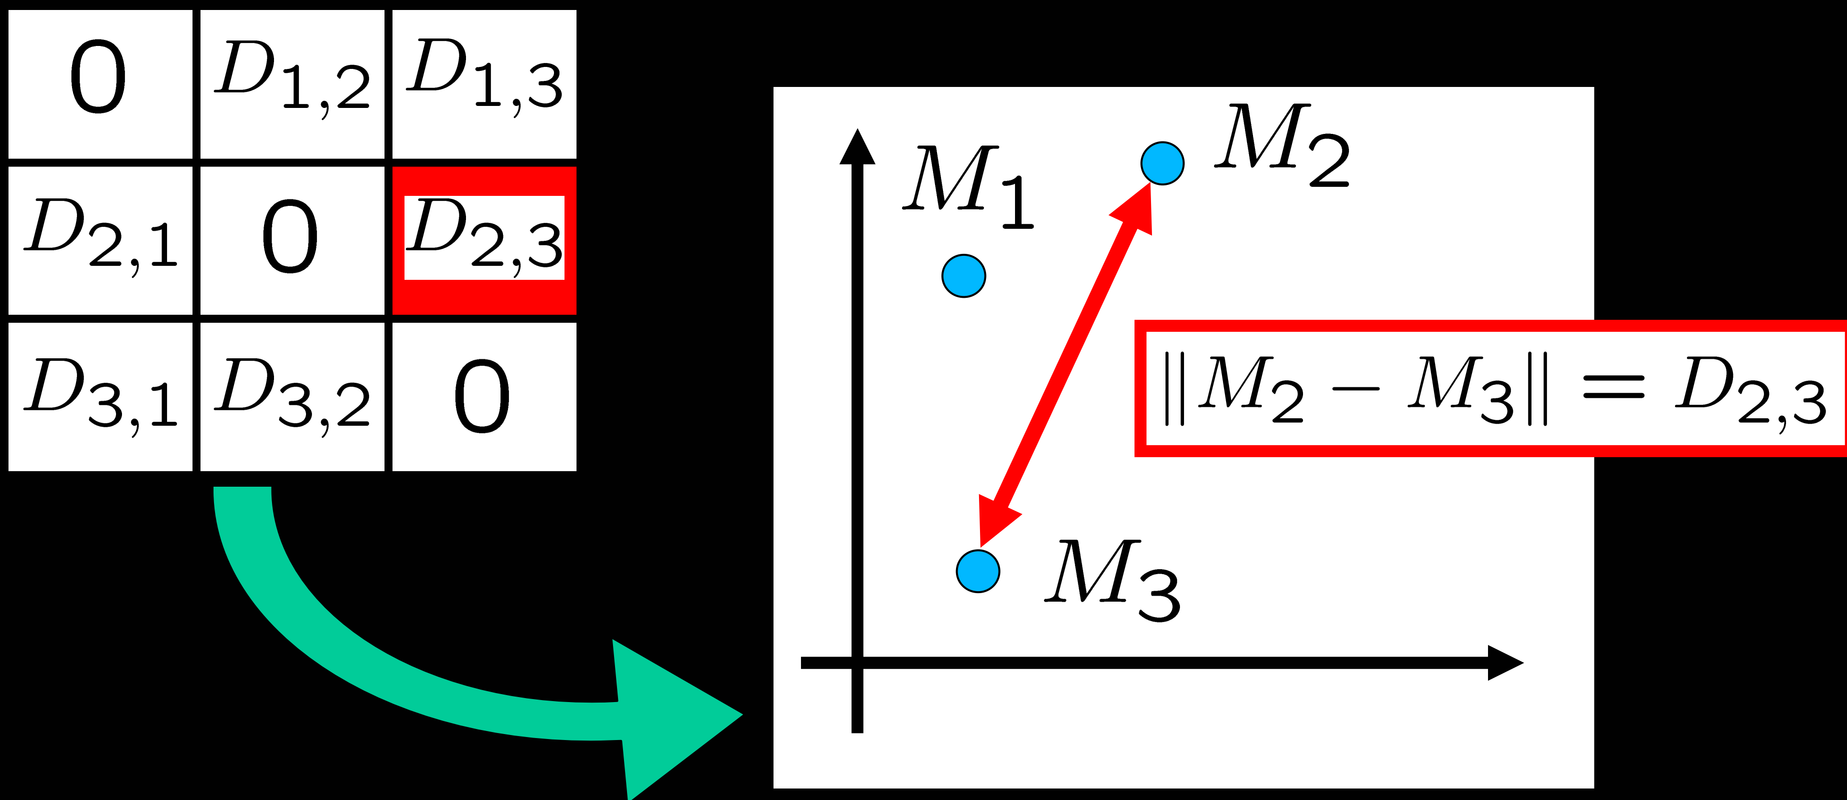

- Multidimensional Scaling (MDS) finds a low-dimensional coordinate mapping…such that, the <u>pairwise distances are preserved</u>

Why Multidimensional Scaling?

- Visualization of high dimensional data
  - Consider embedding in two dimensions

Why Multidimensional Scaling?

- Visualization of high dimensional data
  - Consider embedding in two dimensions
- Clustering

Why Multidimensional Scaling?

- Visualization of high dimensional data
  - Consider embedding in two dimensions
- Clustering
- Estimating pairwise distances using only distances to landmarks

Multidimensional Scaling

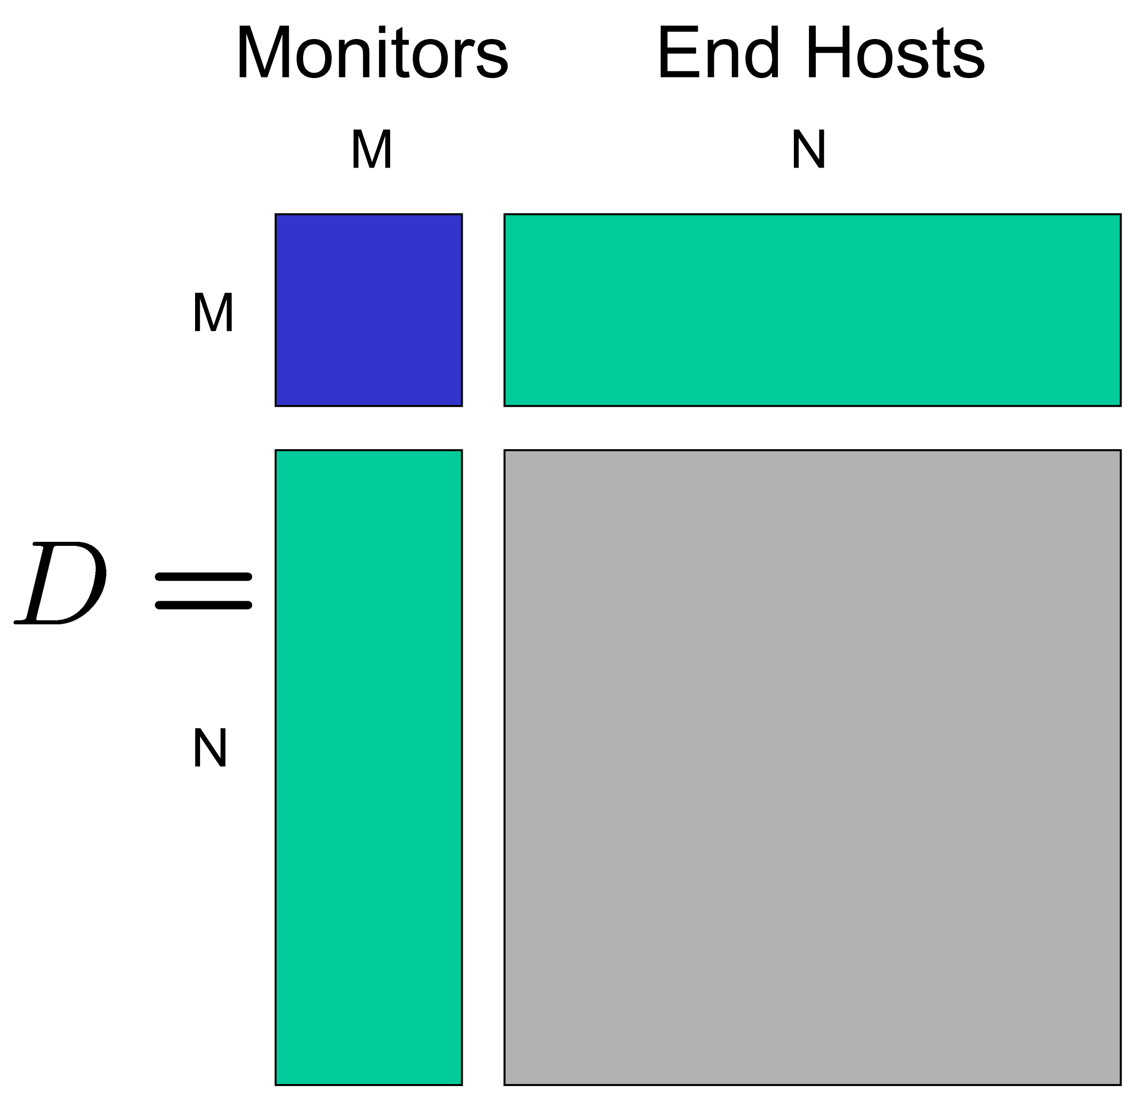

Multidimensional Scaling

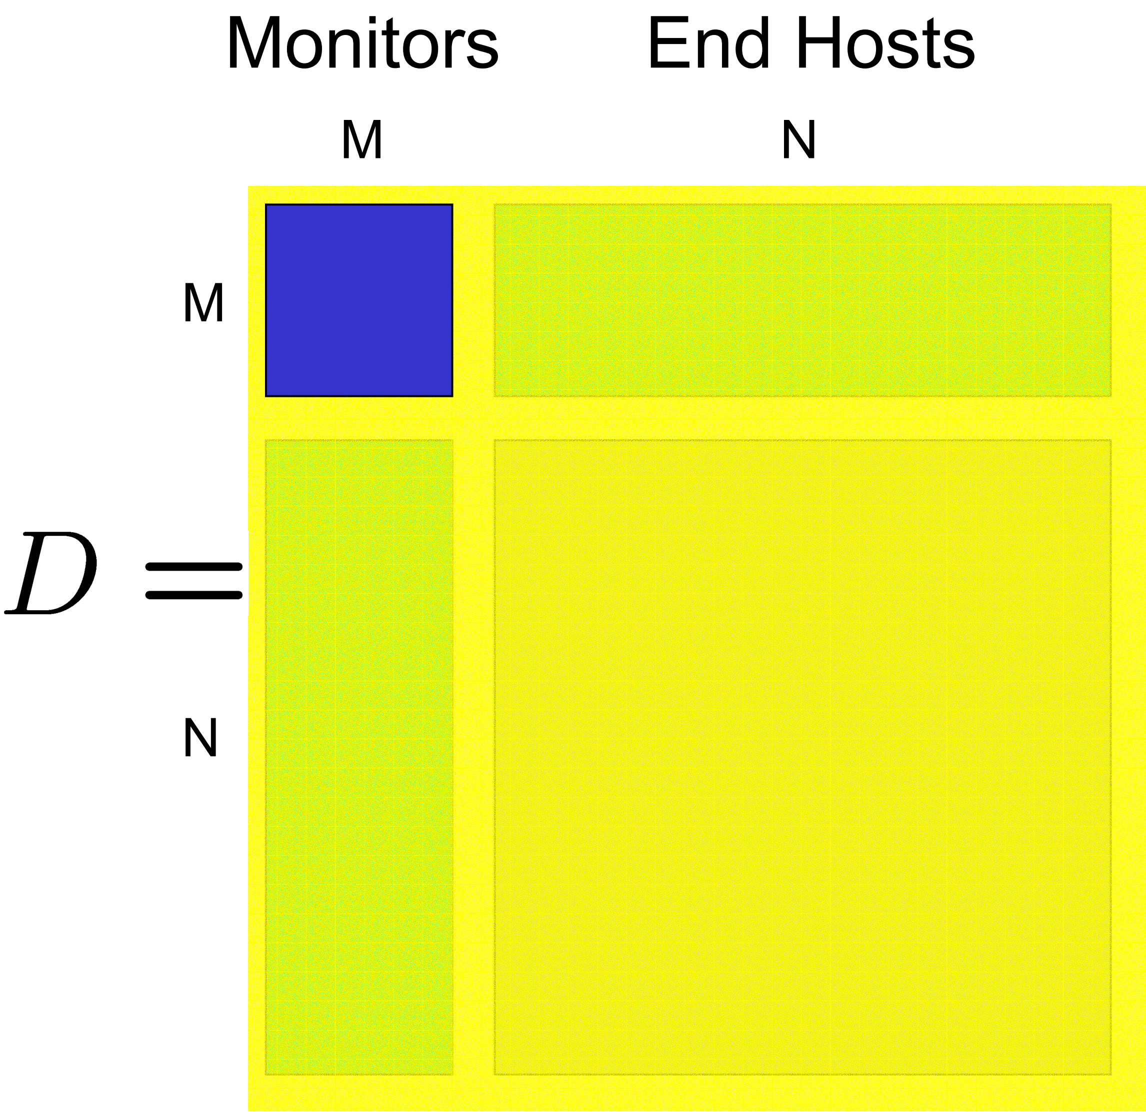

Monitor to Monitor Hop Distances – Known (active)

Multidimensional Scaling

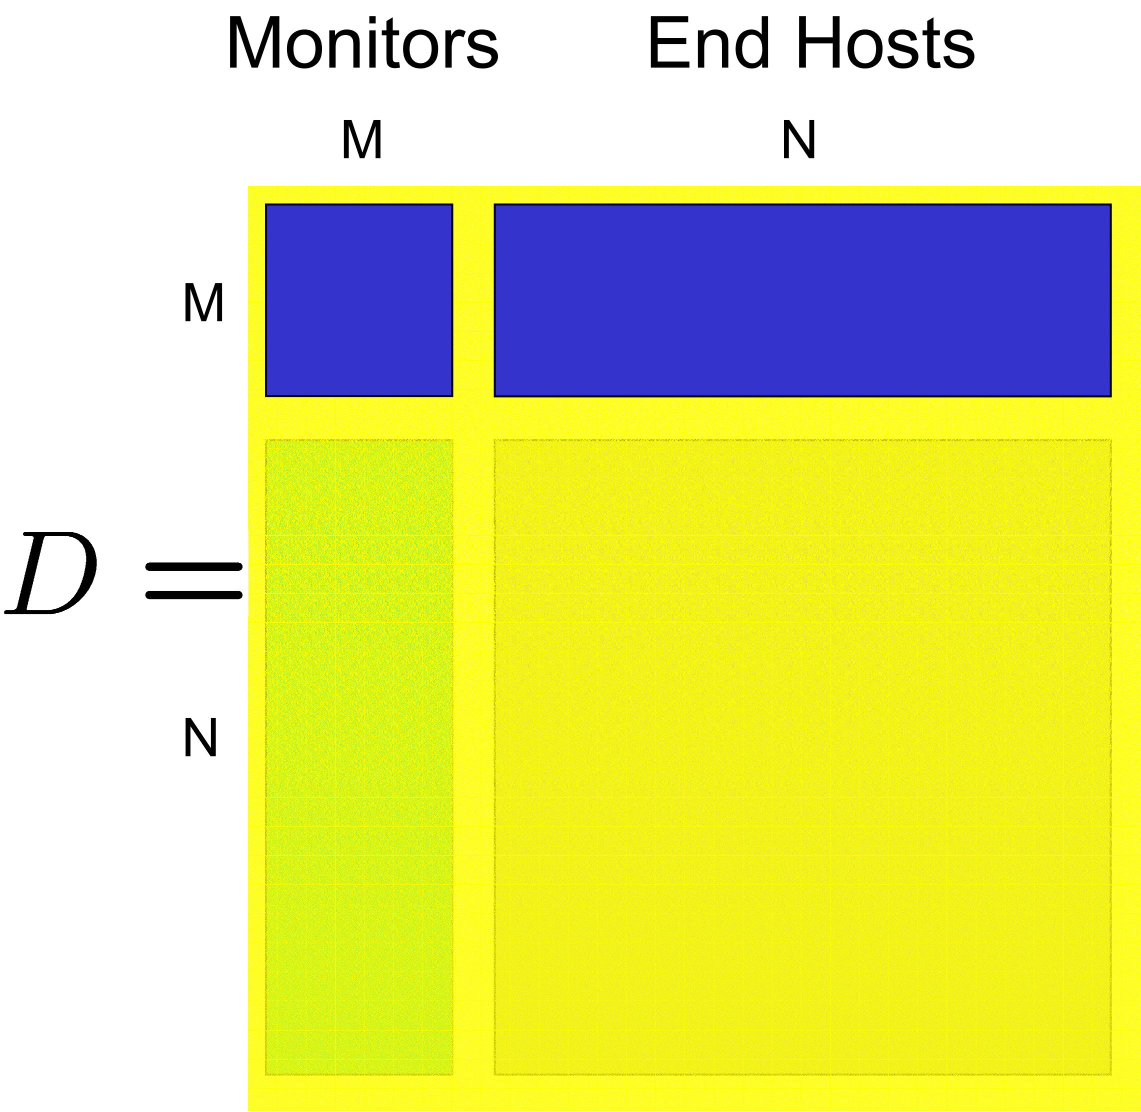

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Known (active)

Multidimensional Scaling

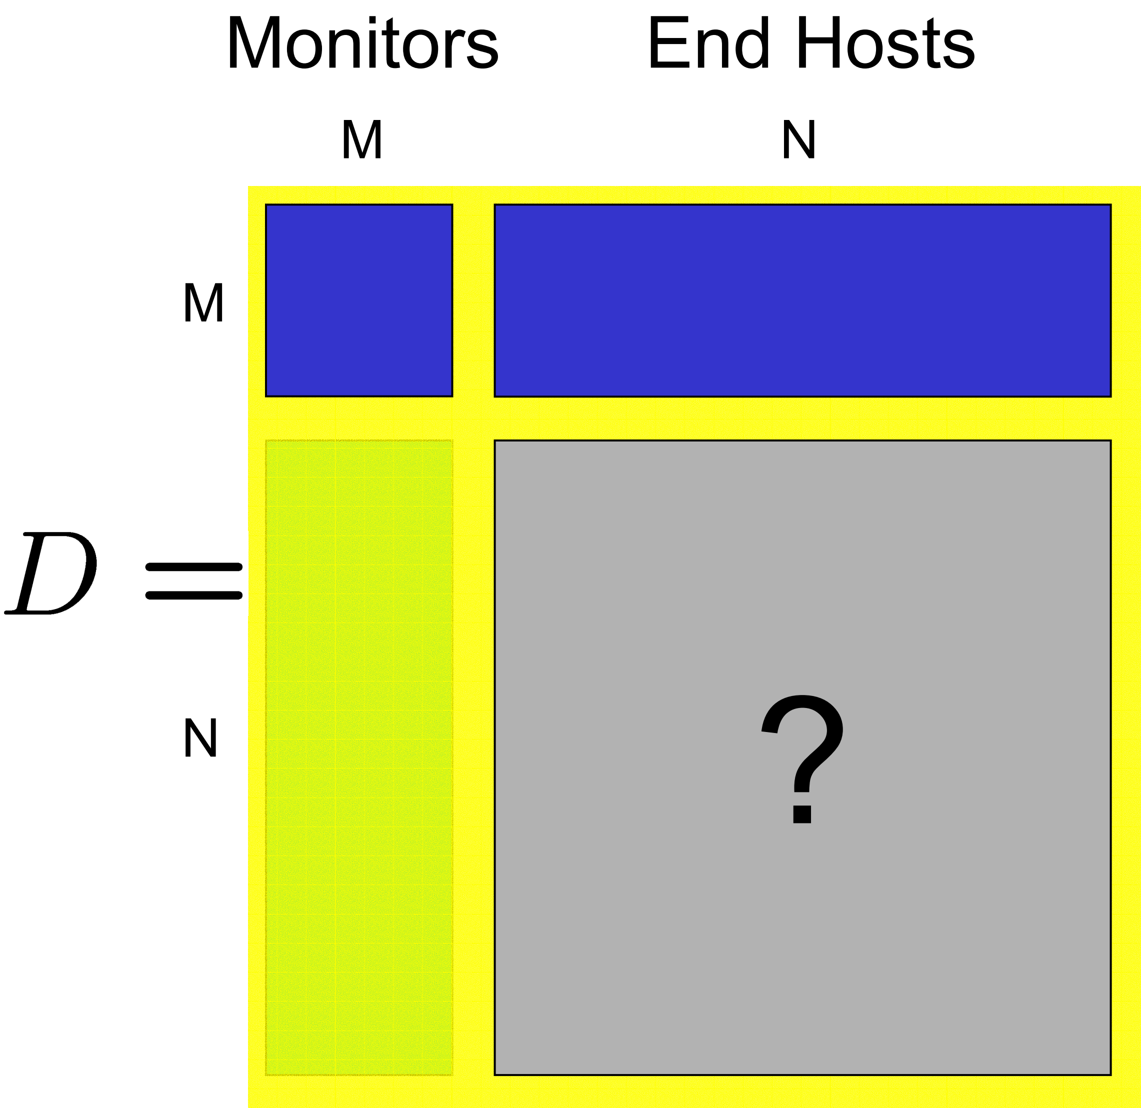

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Known (active)

End host to end host distances – <u>Completely unknown</u>

Multidimensional Scaling

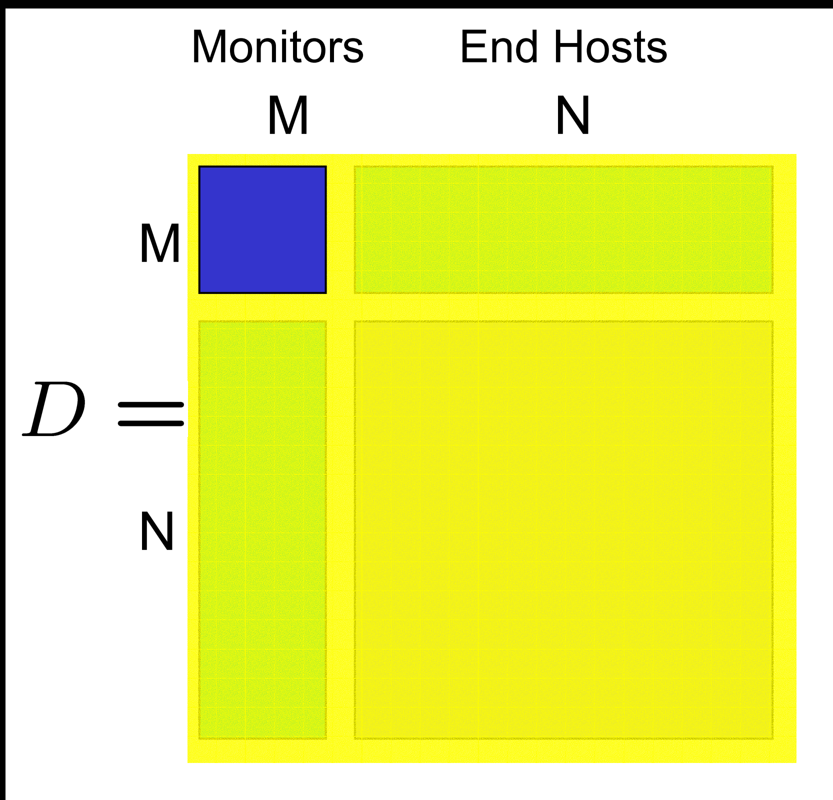

Multidimensional Scaling

- MDS finds monitor embedding

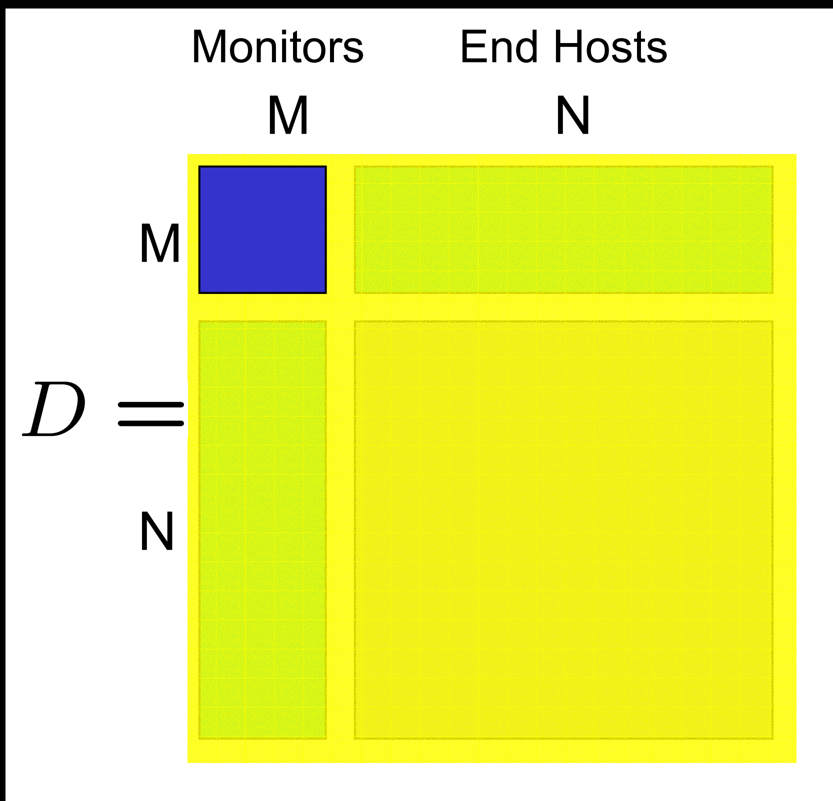

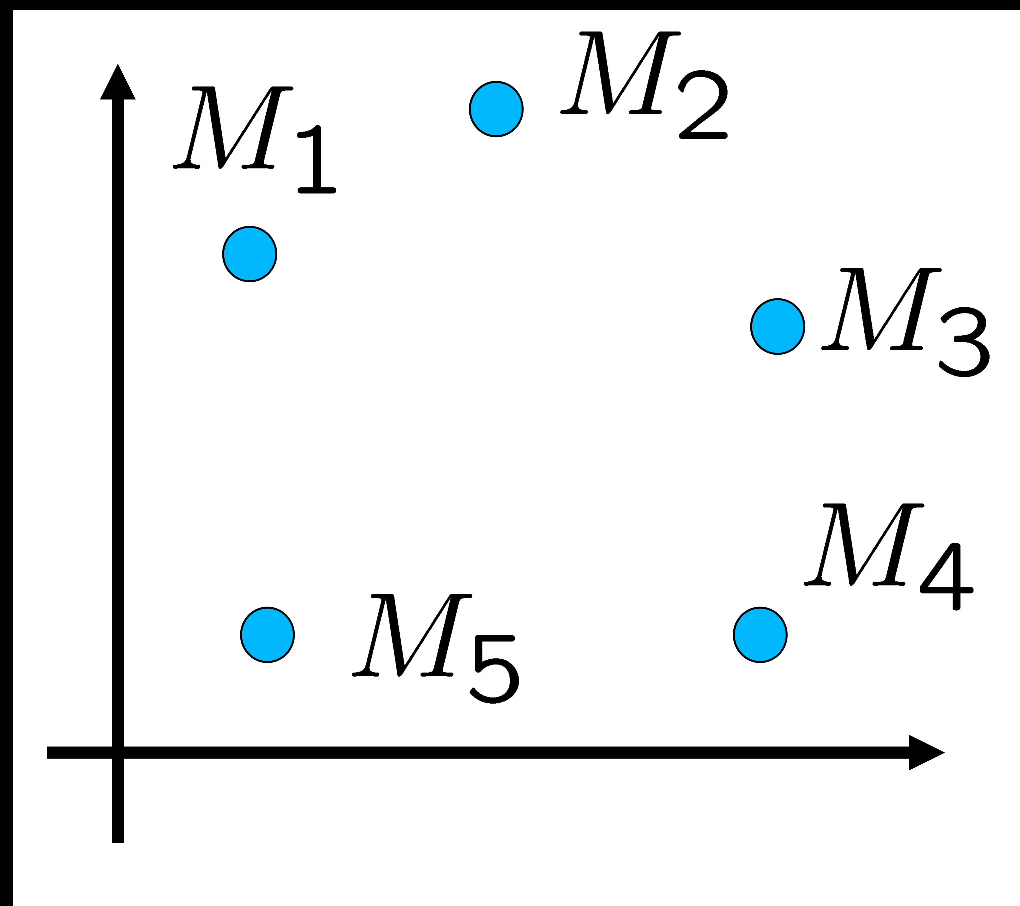

Multidimensional Scaling

- MDS finds monitor embedding

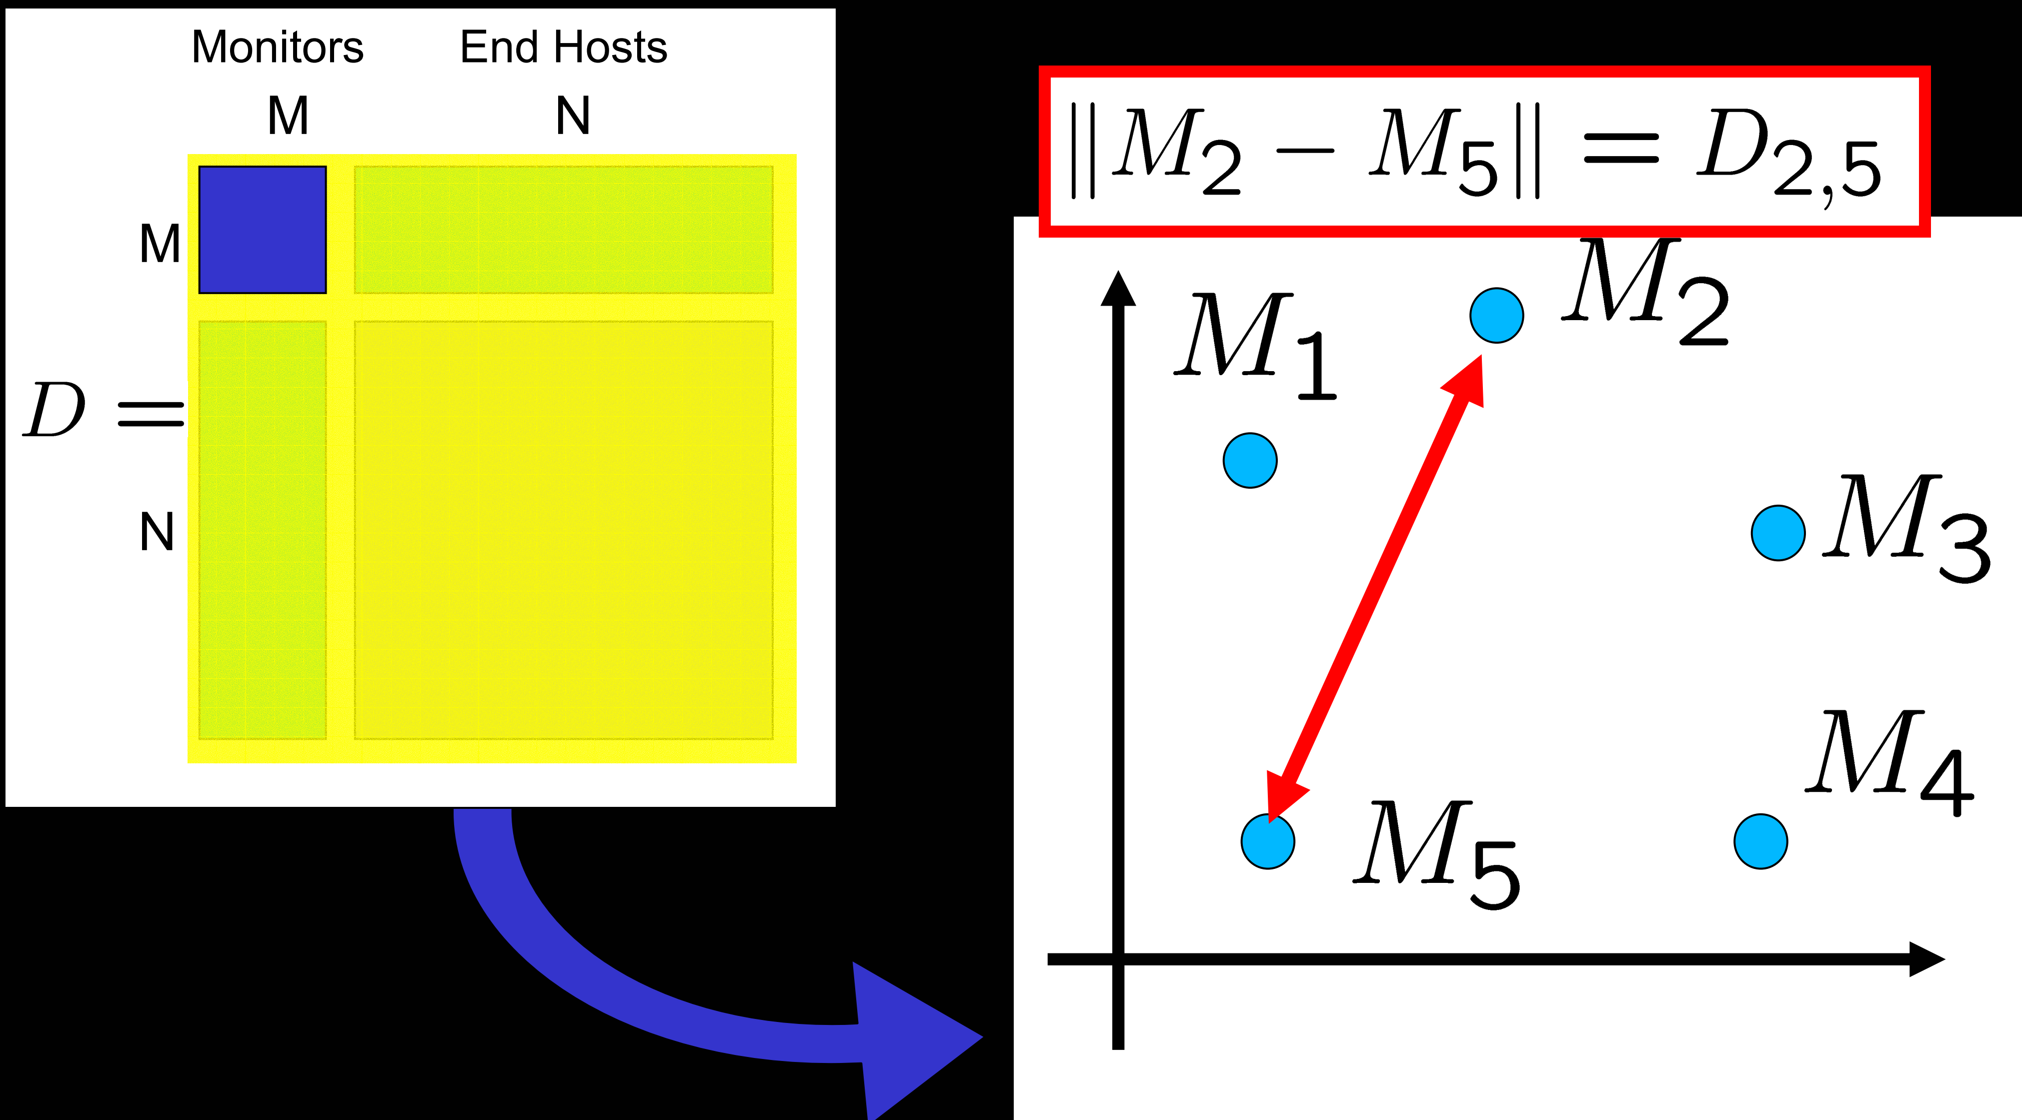

Landmark MDS (LMDS)

- Where should we place the end host?

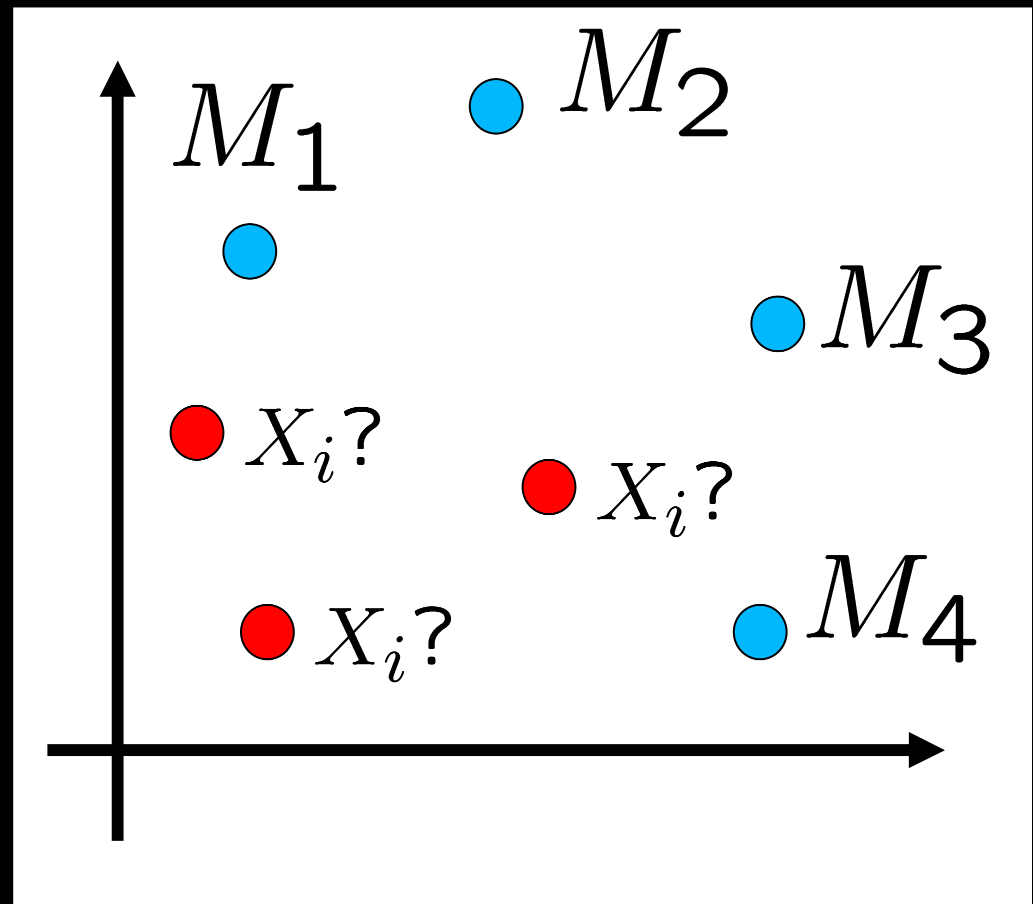

Monitor Embedding

Landmark MDS (LMDS)

- Where should we place the end host?
- Using distances to monitors
  - "Triangulate"

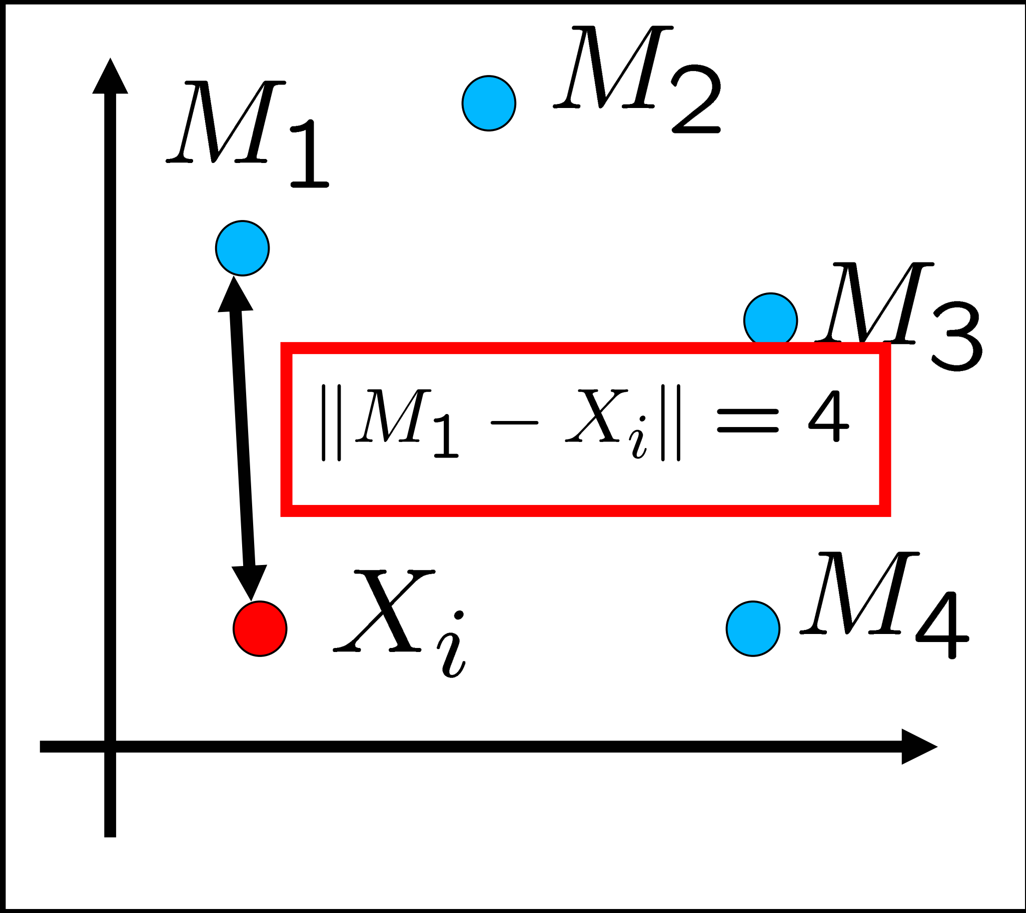

h_i

4	6	3	8	9	5	2	3

Landmark MDS (LMDS)

- End host to End Host distance can be found by embedding distance.

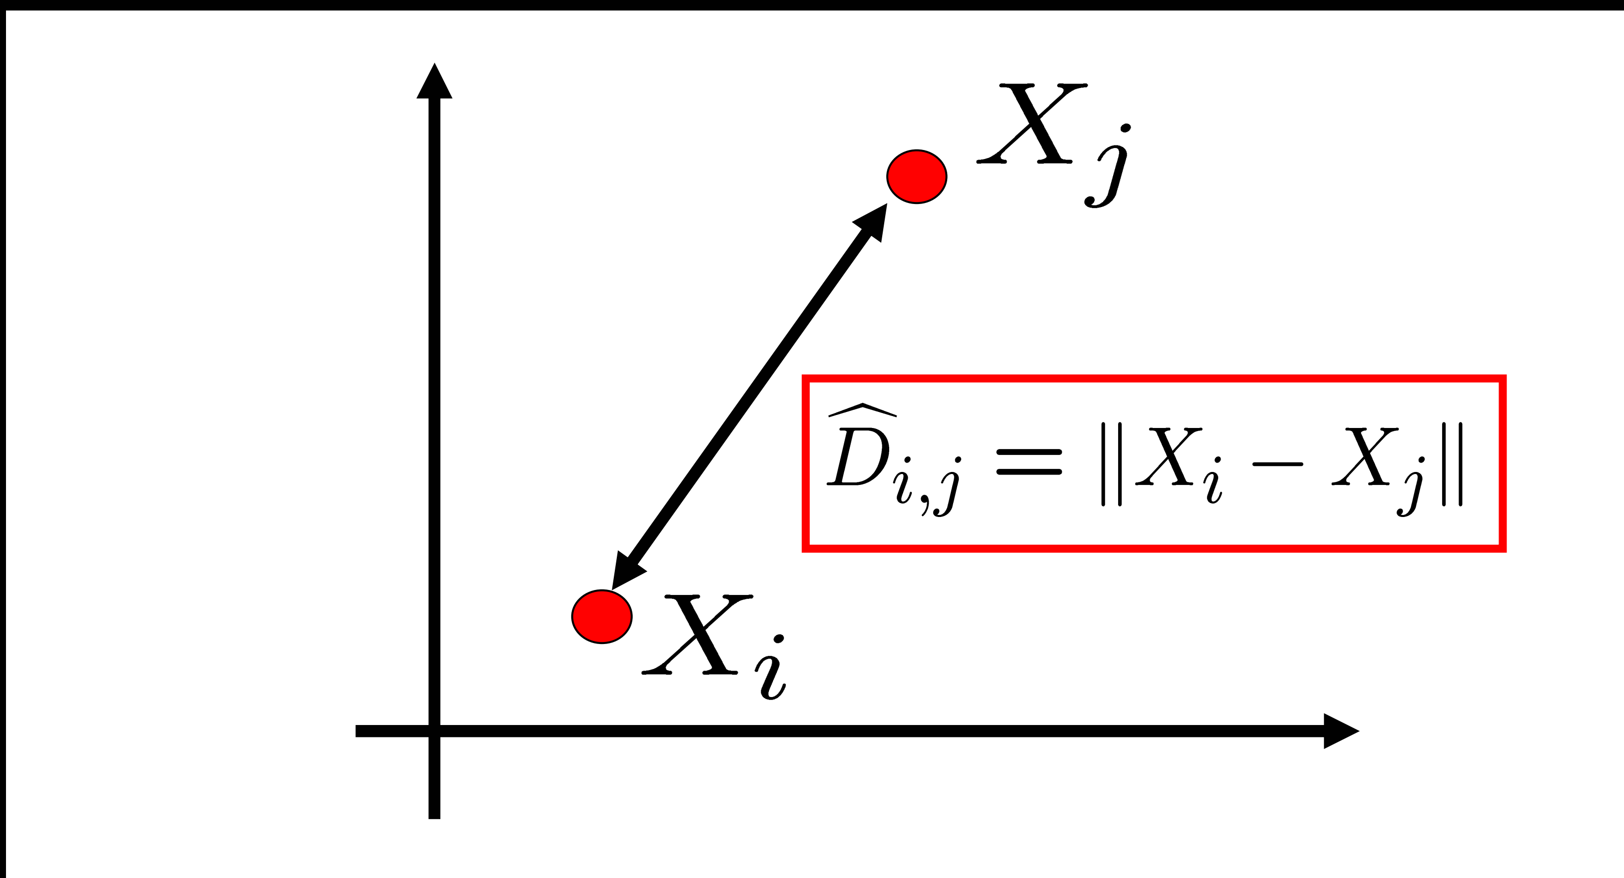

End to End Estimation

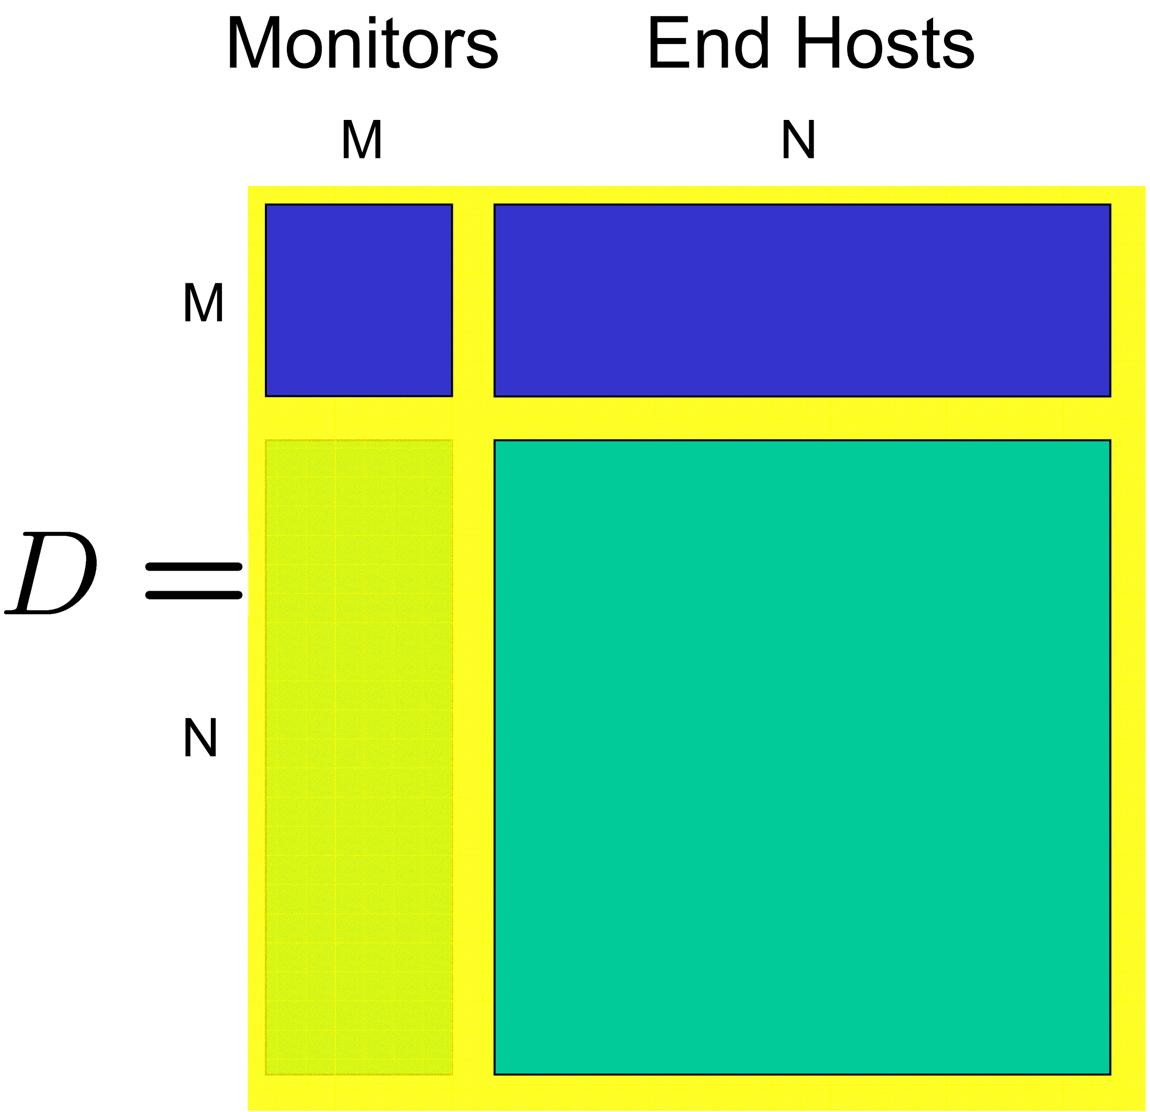

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Known (active)

End host to end host data – <u>Estimated by MDS</u>

End to End Estimation

- Prior work has examined embedding using hop counts

M. Costa, M. Castro, A. Rowstron, and P. Key, "PIC: Practical Internet coordinates for distance estimation," in *In International Conference on Distributed Systems*, March 2004.
Y. Shavitt and T. Tankel, "Hyperbolic Embedding of Internet Graphs for Distance Estimation and Overlay Construction," *IEEE/ACM Transactions on Networking*,

End to End Estimation

- Prior work has examined embedding using hop counts
- For prior techniques:

$$\text{Number of Active Probes} \sim O\left(M^2 + NM\right)$$

M. Costa, M. Castro, A. Rowstron, and P. Key, "PIC: Practical Internet coordinates for distance estimation," in *In International Conference on Distributed Systems*, March 2004.
Y. Shavitt and T. Tankel, "Hyperbolic Embedding of Internet Graphs for Distance Estimation and Overlay Construction," *IEEE/ACM Transactions on Networking*,

End to End Estimation

- Prior work has examined embedding using hop counts
- For prior techniques:

$$\text{Number of Active Probes} \sim O\left(M^2 + NM\right)$$

- What if N is very large?

M. Costa, M. Castro, A. Rowstron, and P. Key, "PIC: Practical Internet coordinates for distance estimation," in *In International Conference on Distributed Systems*, March 2004.
Y. Shavitt and T. Tankel, "Hyperbolic Embedding of Internet Graphs for Distance Estimation and Overlay Construction," *IEEE/ACM Transactions on Networking*,

End to End Estimation

- Prior work has examined embedding using hop counts
- For prior techniques:

$$\text{Number of Active Probes} \sim O\left(M^2 + NM\right)$$

- What if N is very large?
- What if some ISPs block probes?

M. Costa, M. Castro, A. Rowstron, and P. Key, "PIC: Practical Internet coordinates for distance estimation," in *In International Conference on Distributed Systems*, March 2004.
Y. Shavitt and T. Tankel, "Hyperbolic Embedding of Internet Graphs for Distance Estimation and Overlay Construction," *IEEE/ACM Transactions on Networking*,

Passive Measurements

Time To Live

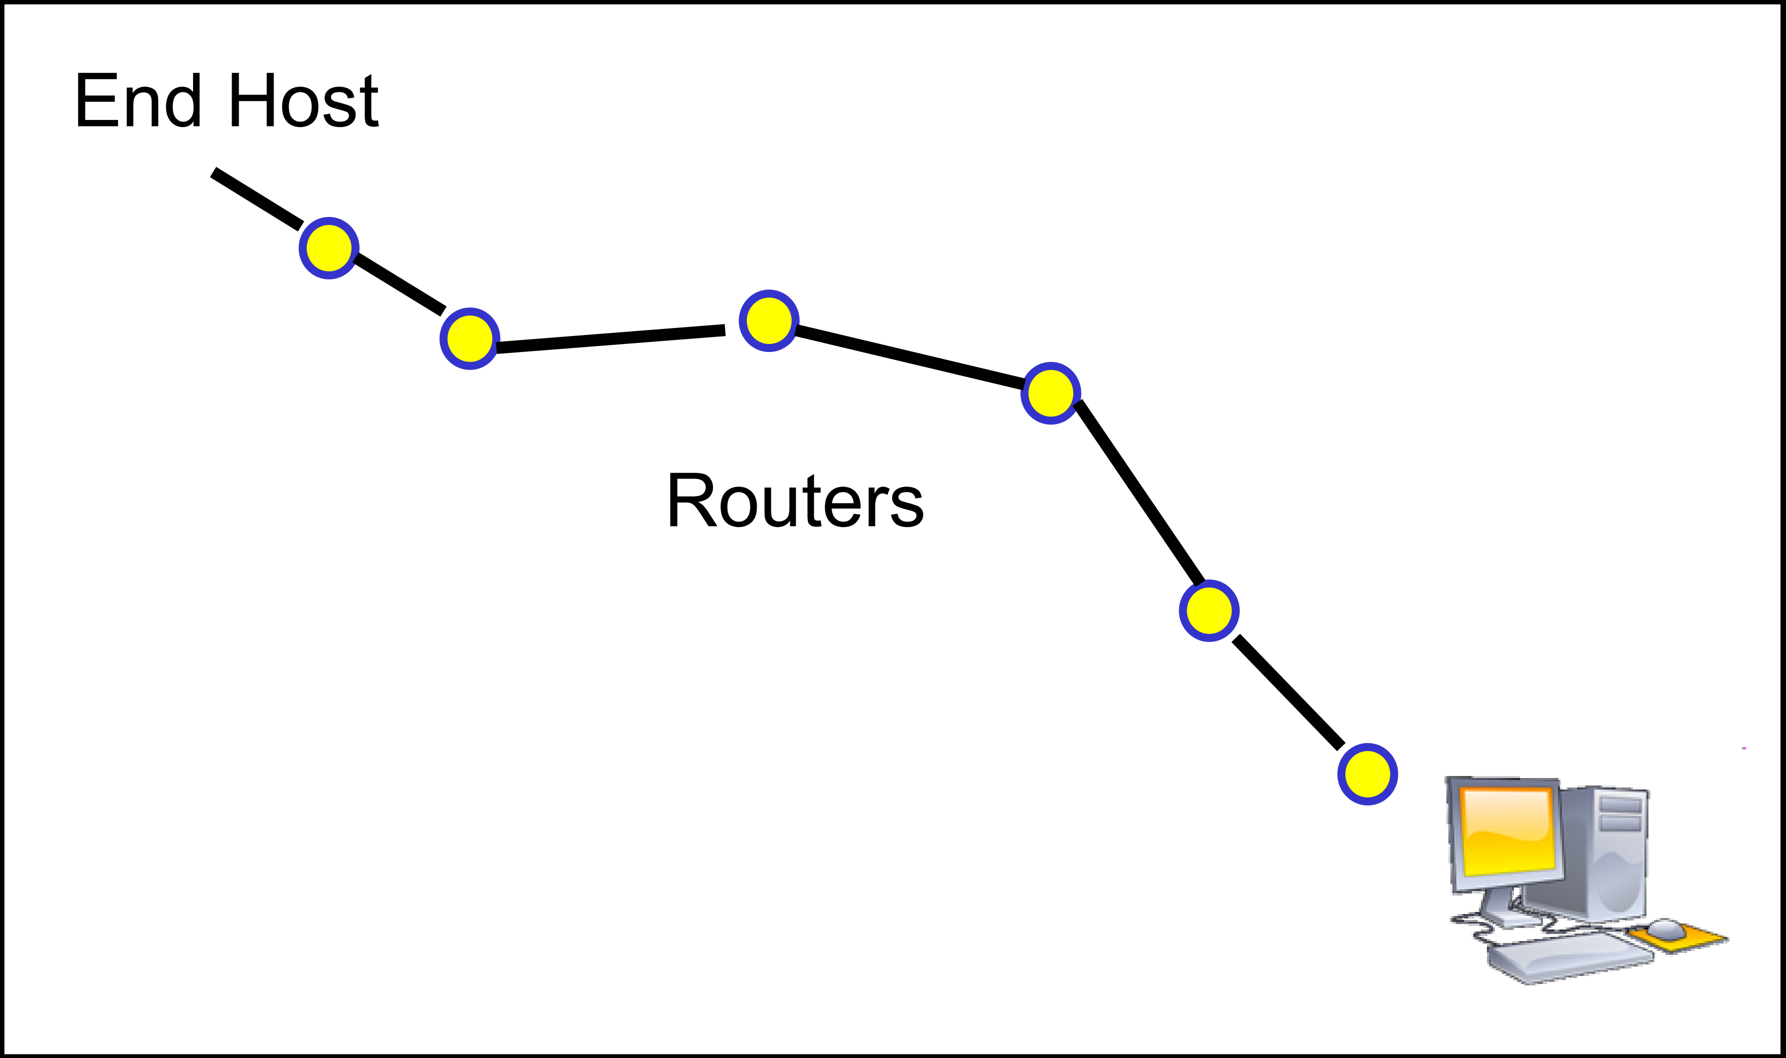

Time To Live

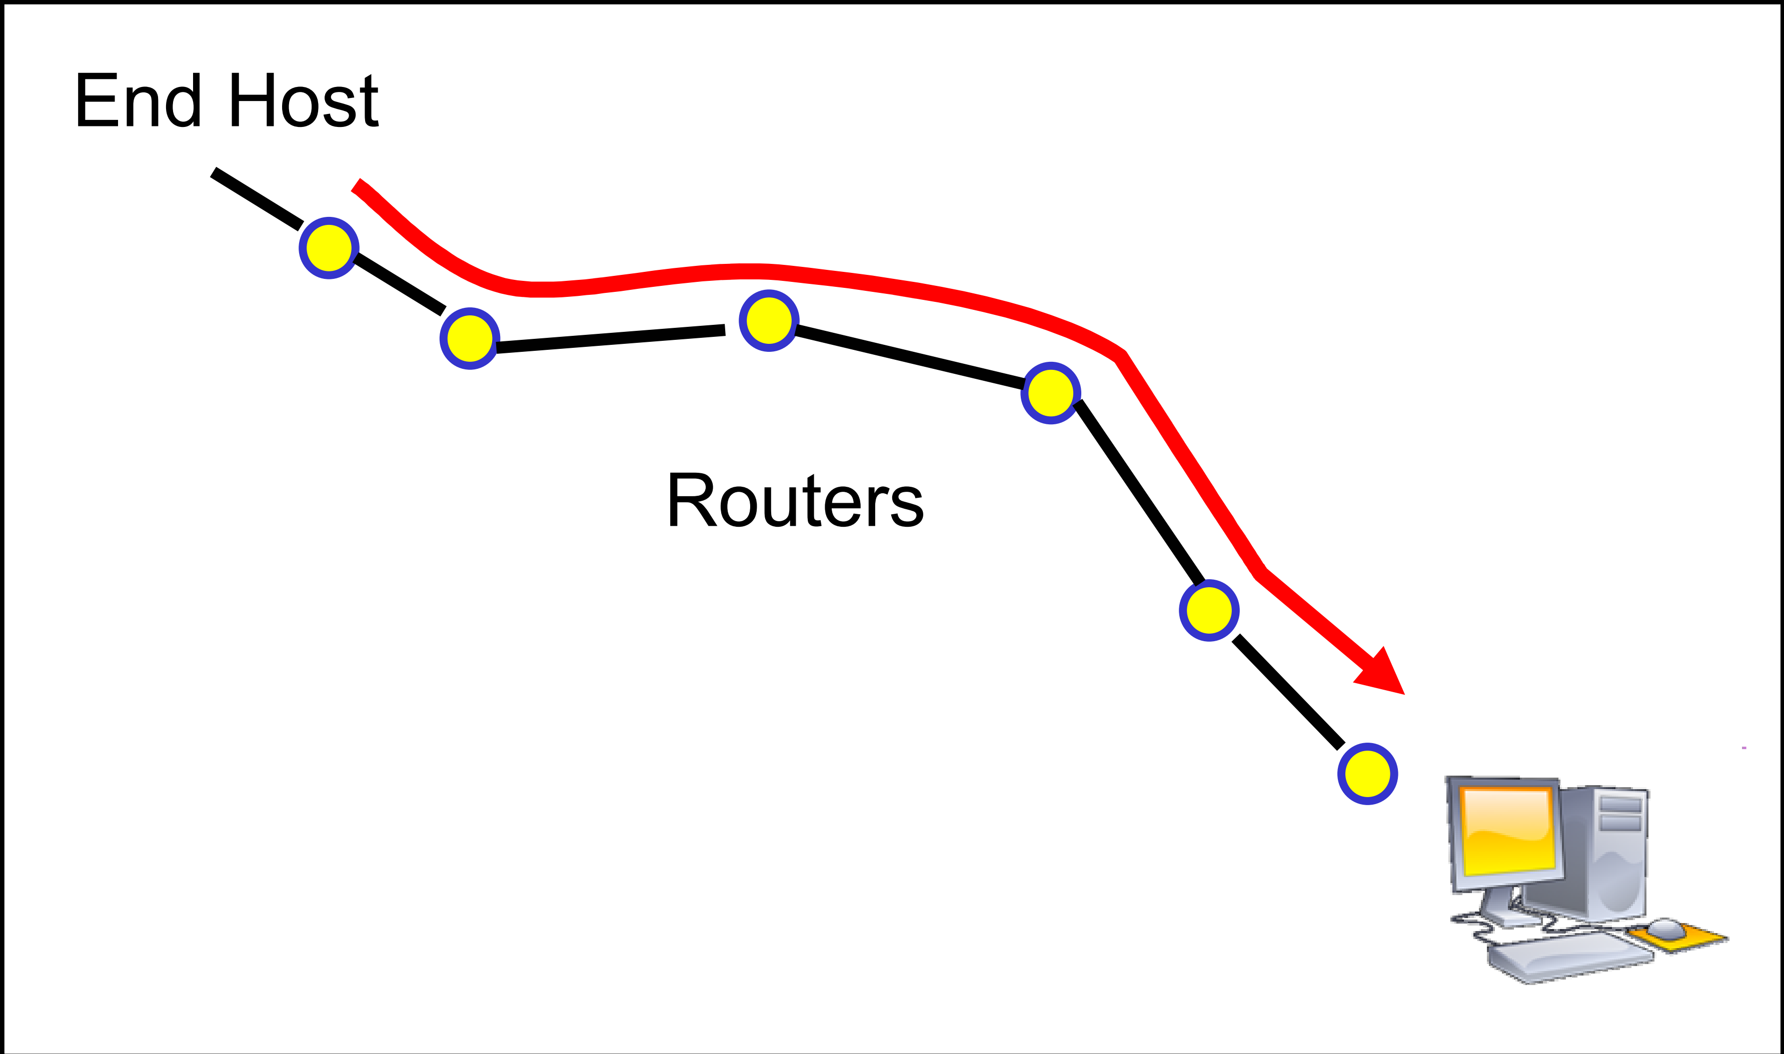

Time To Live

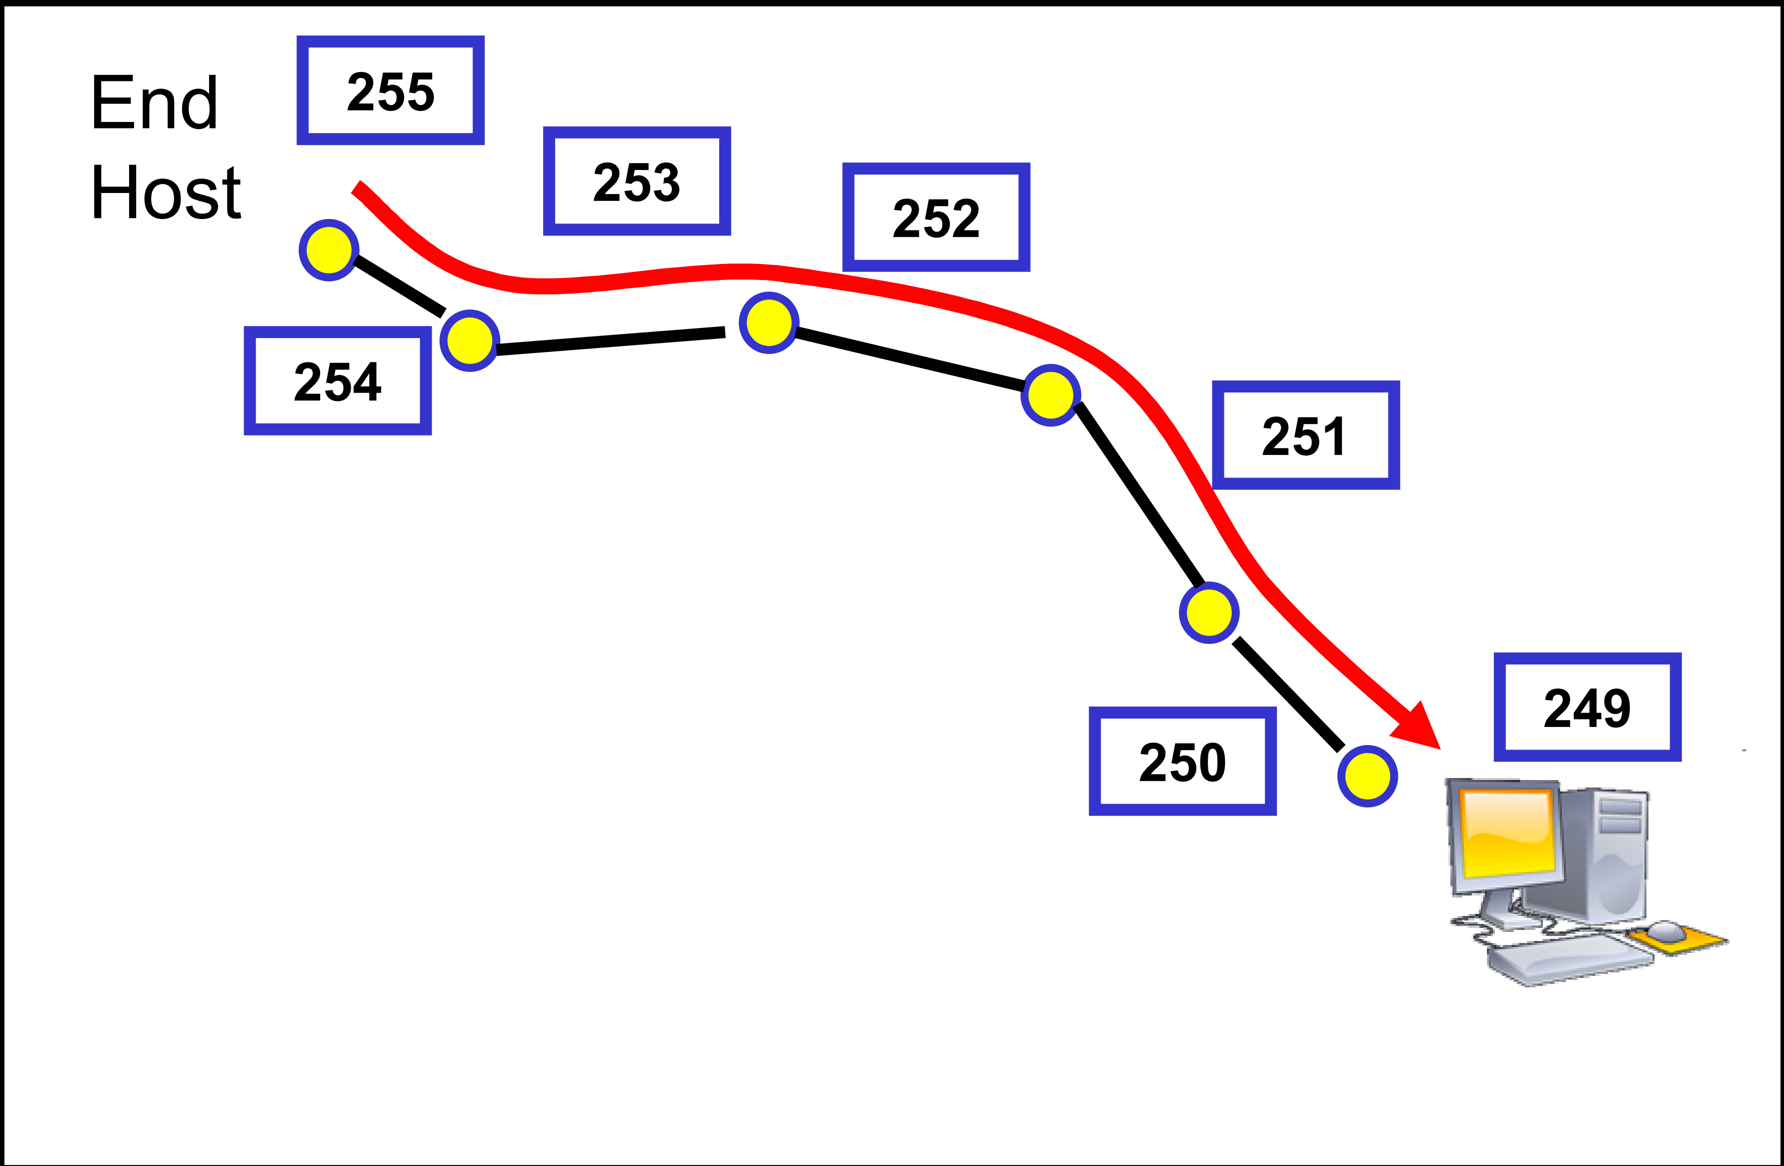

Time To Live

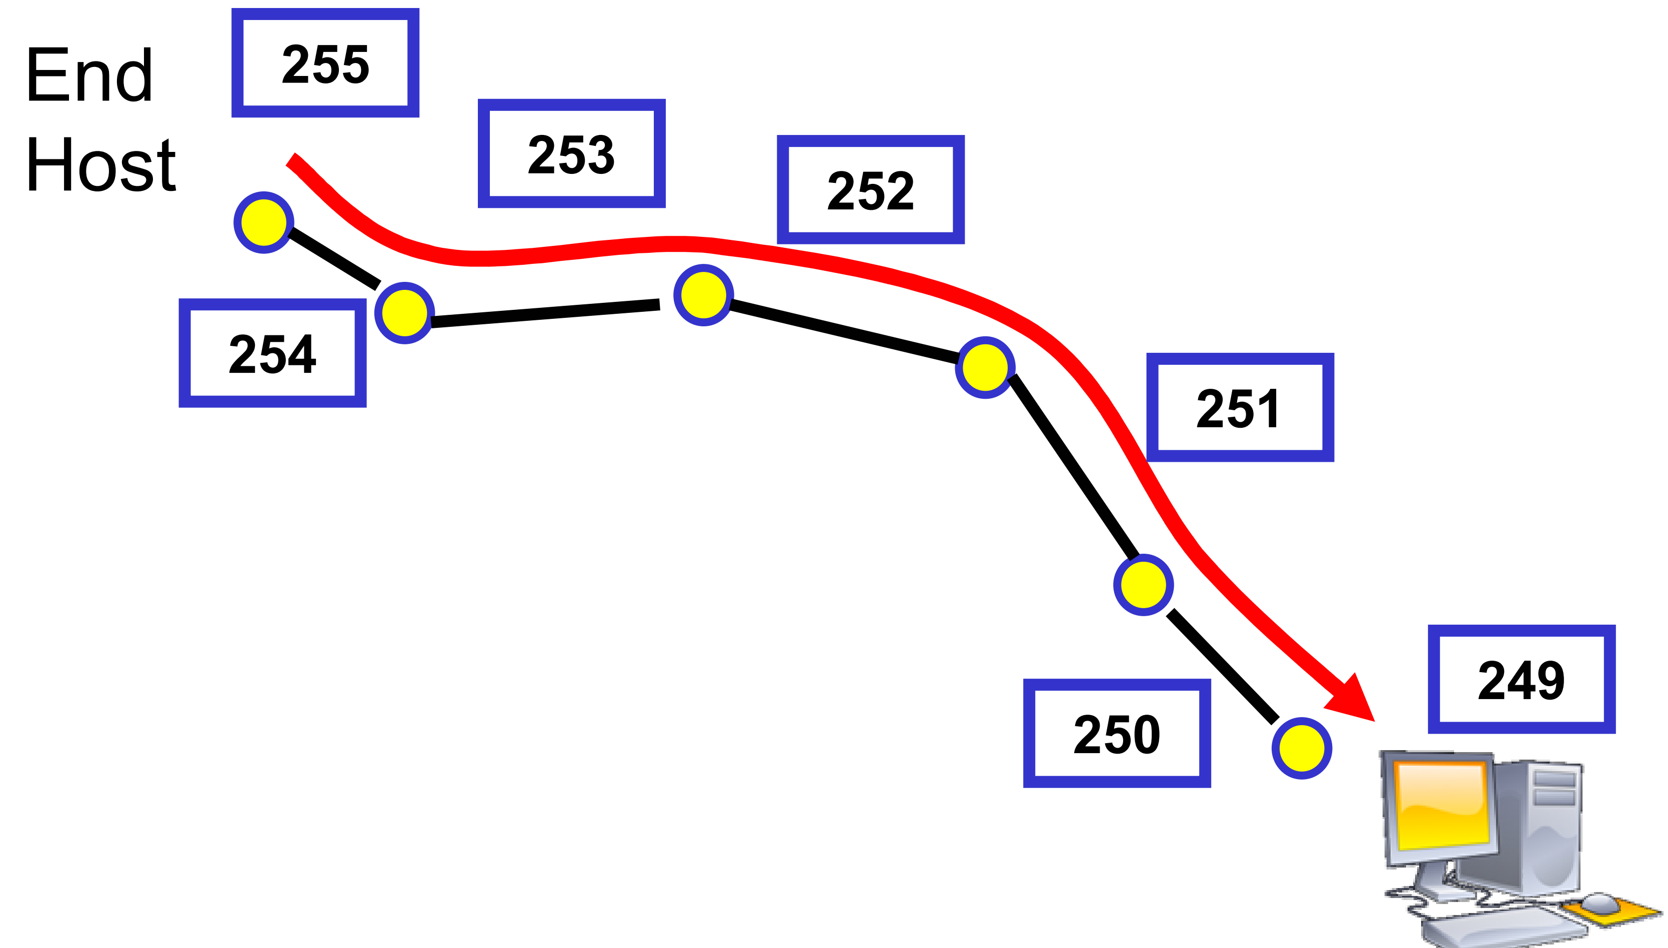

$255 - 249 = 6$ routers

Passive Measurements

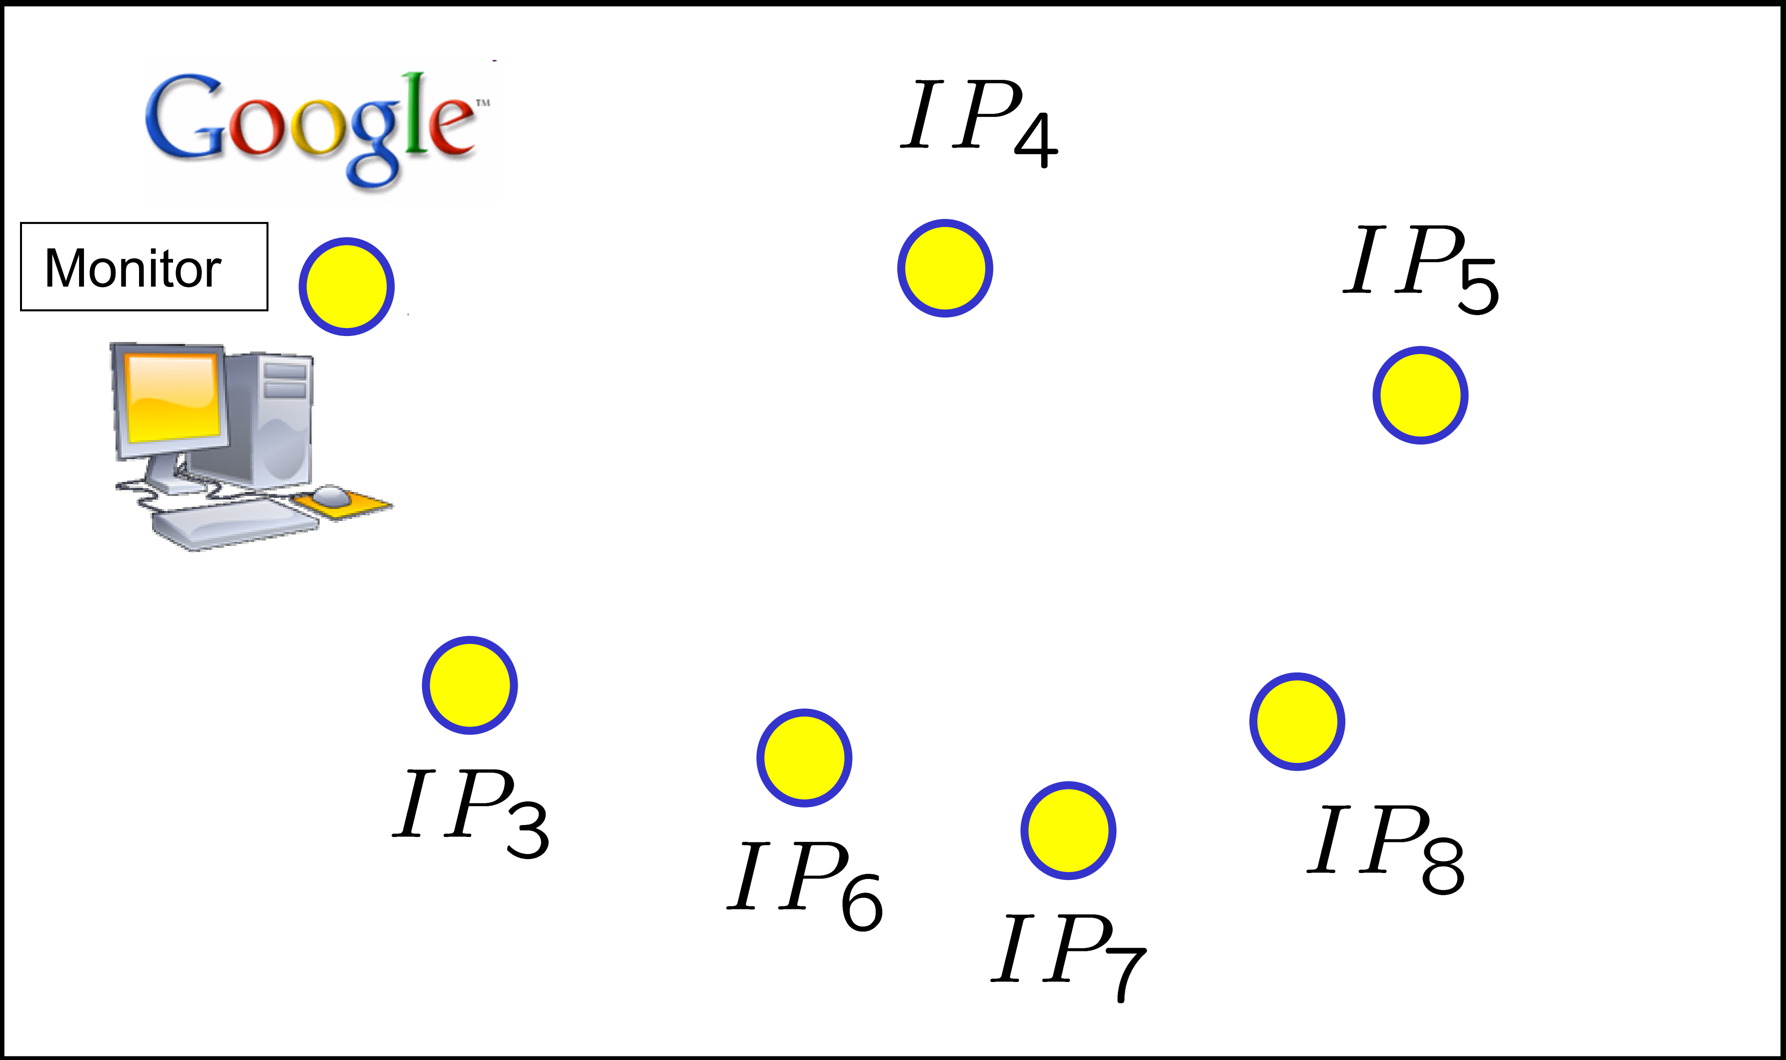

Passive Measurements

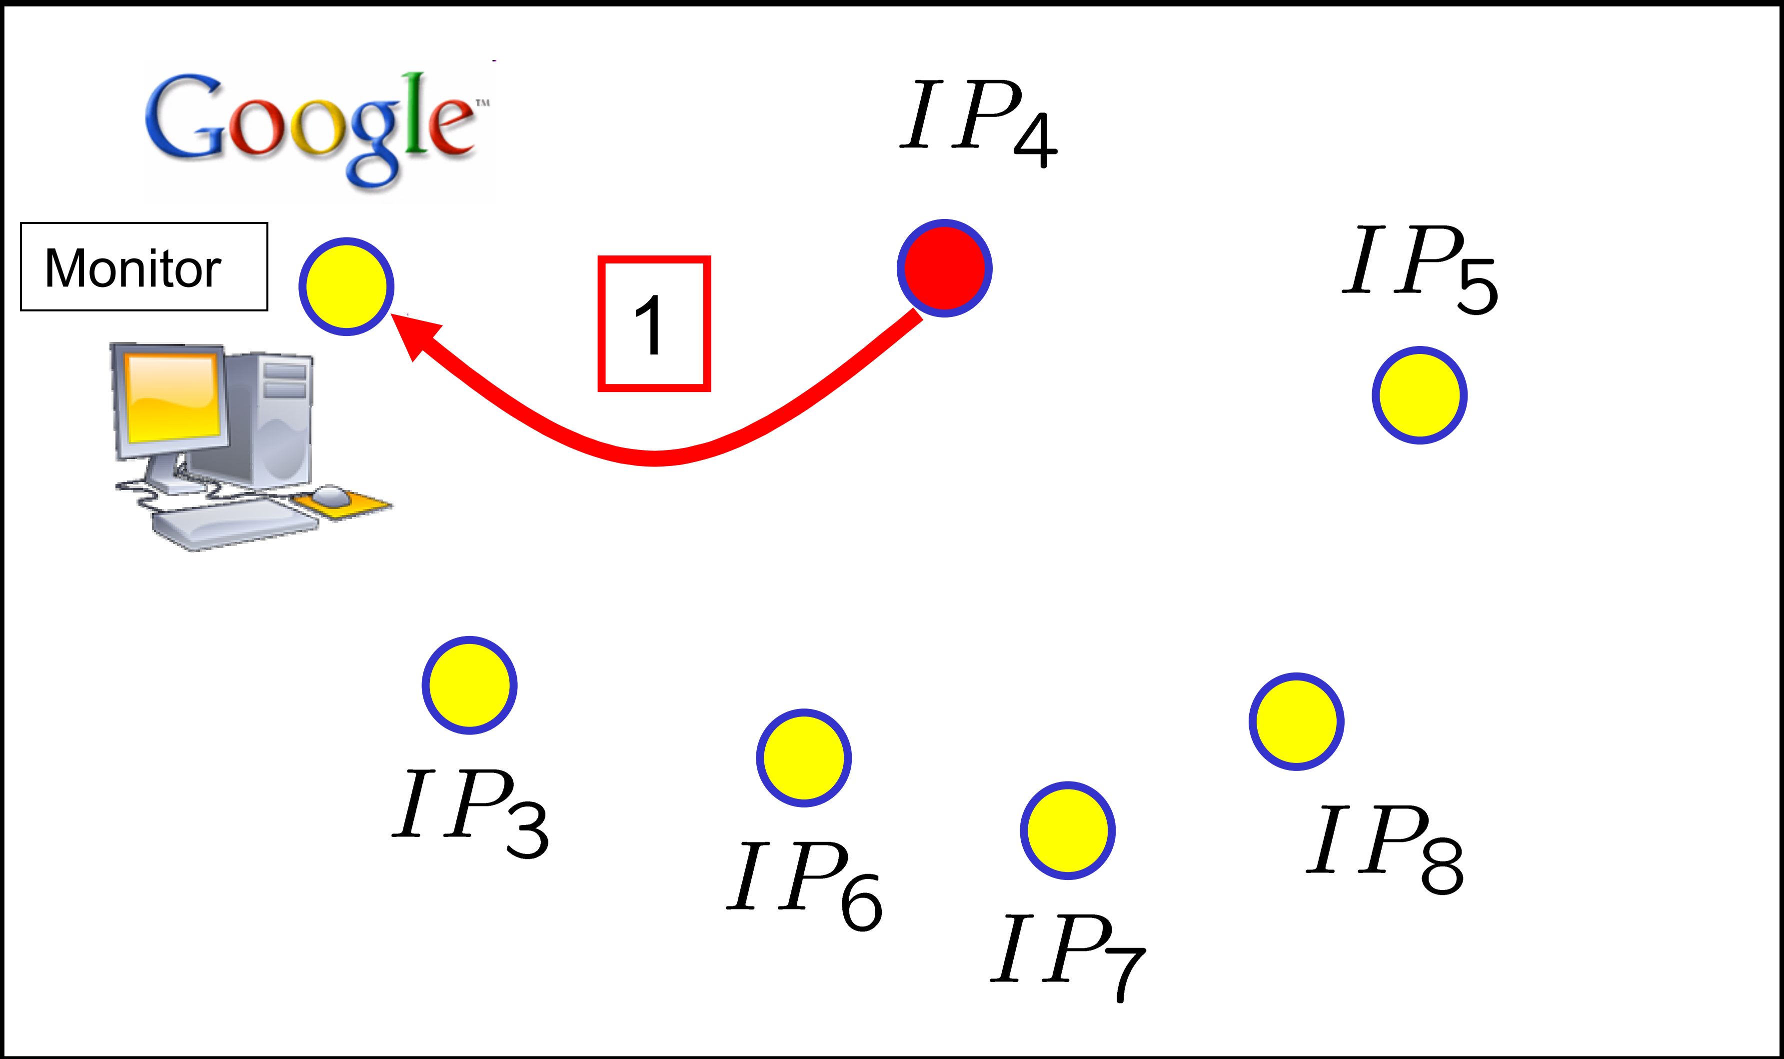

Passive Measurements

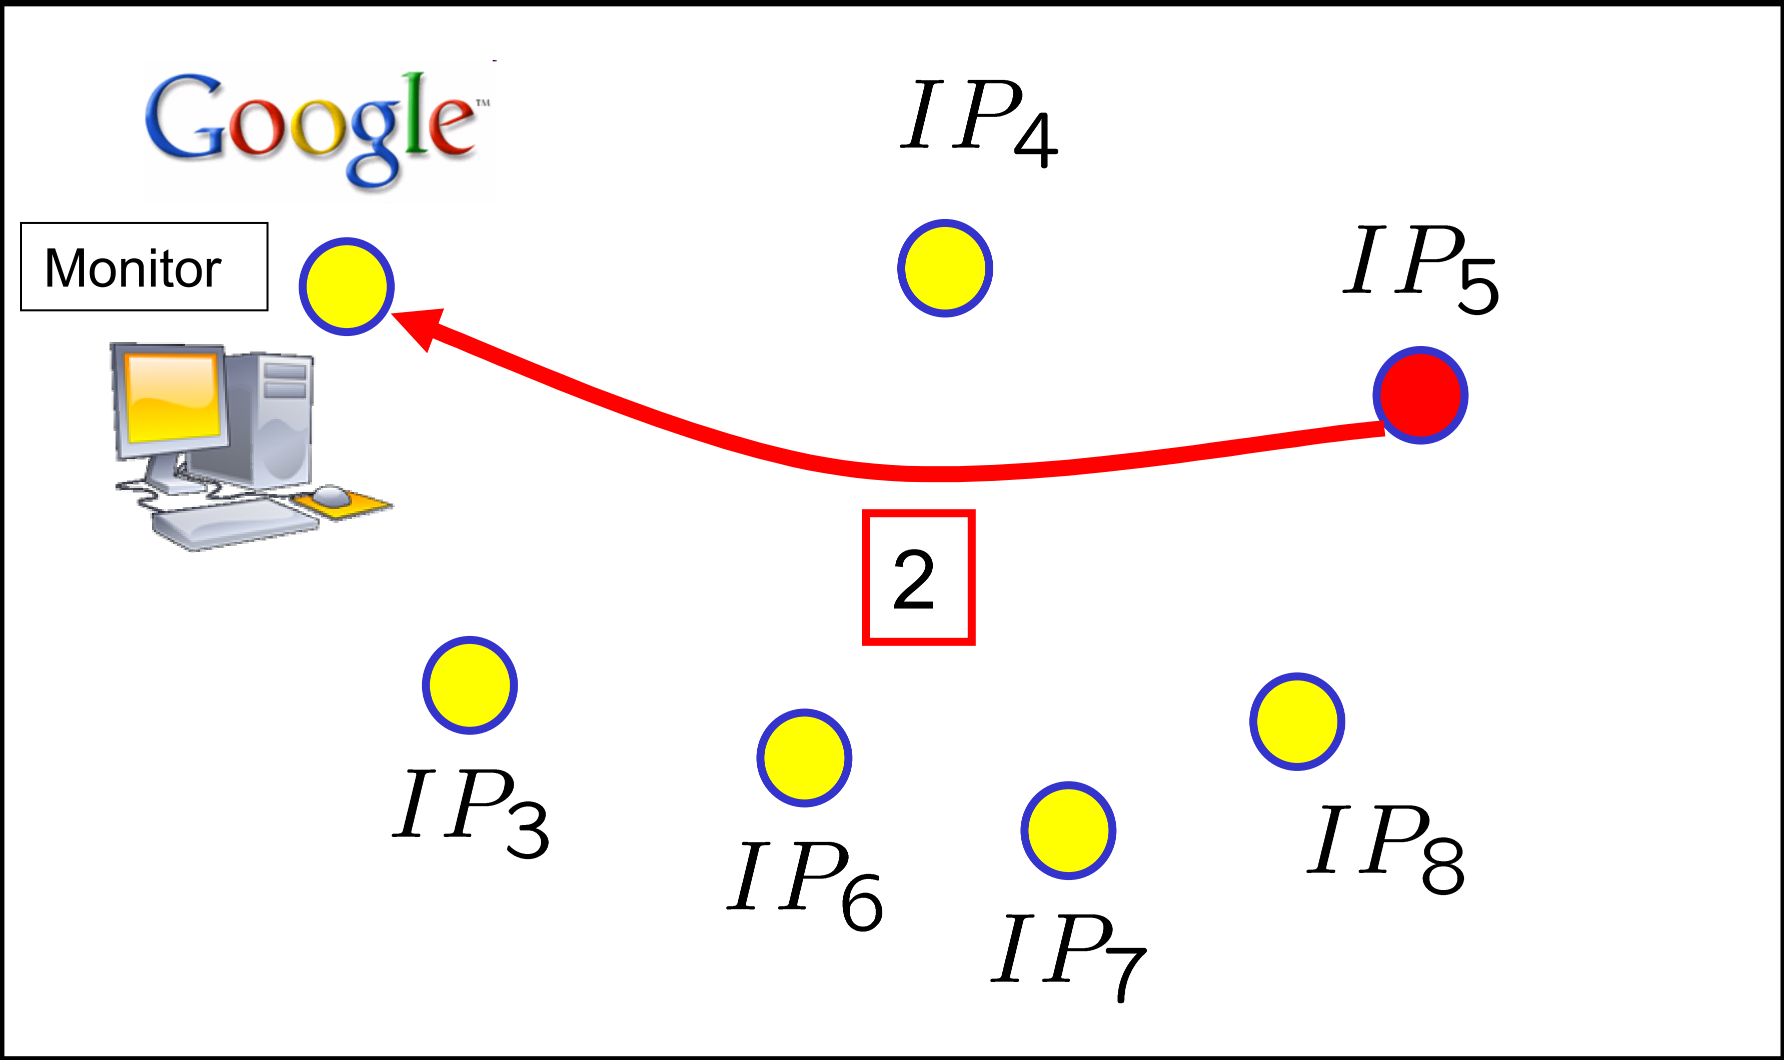

Passive Measurements

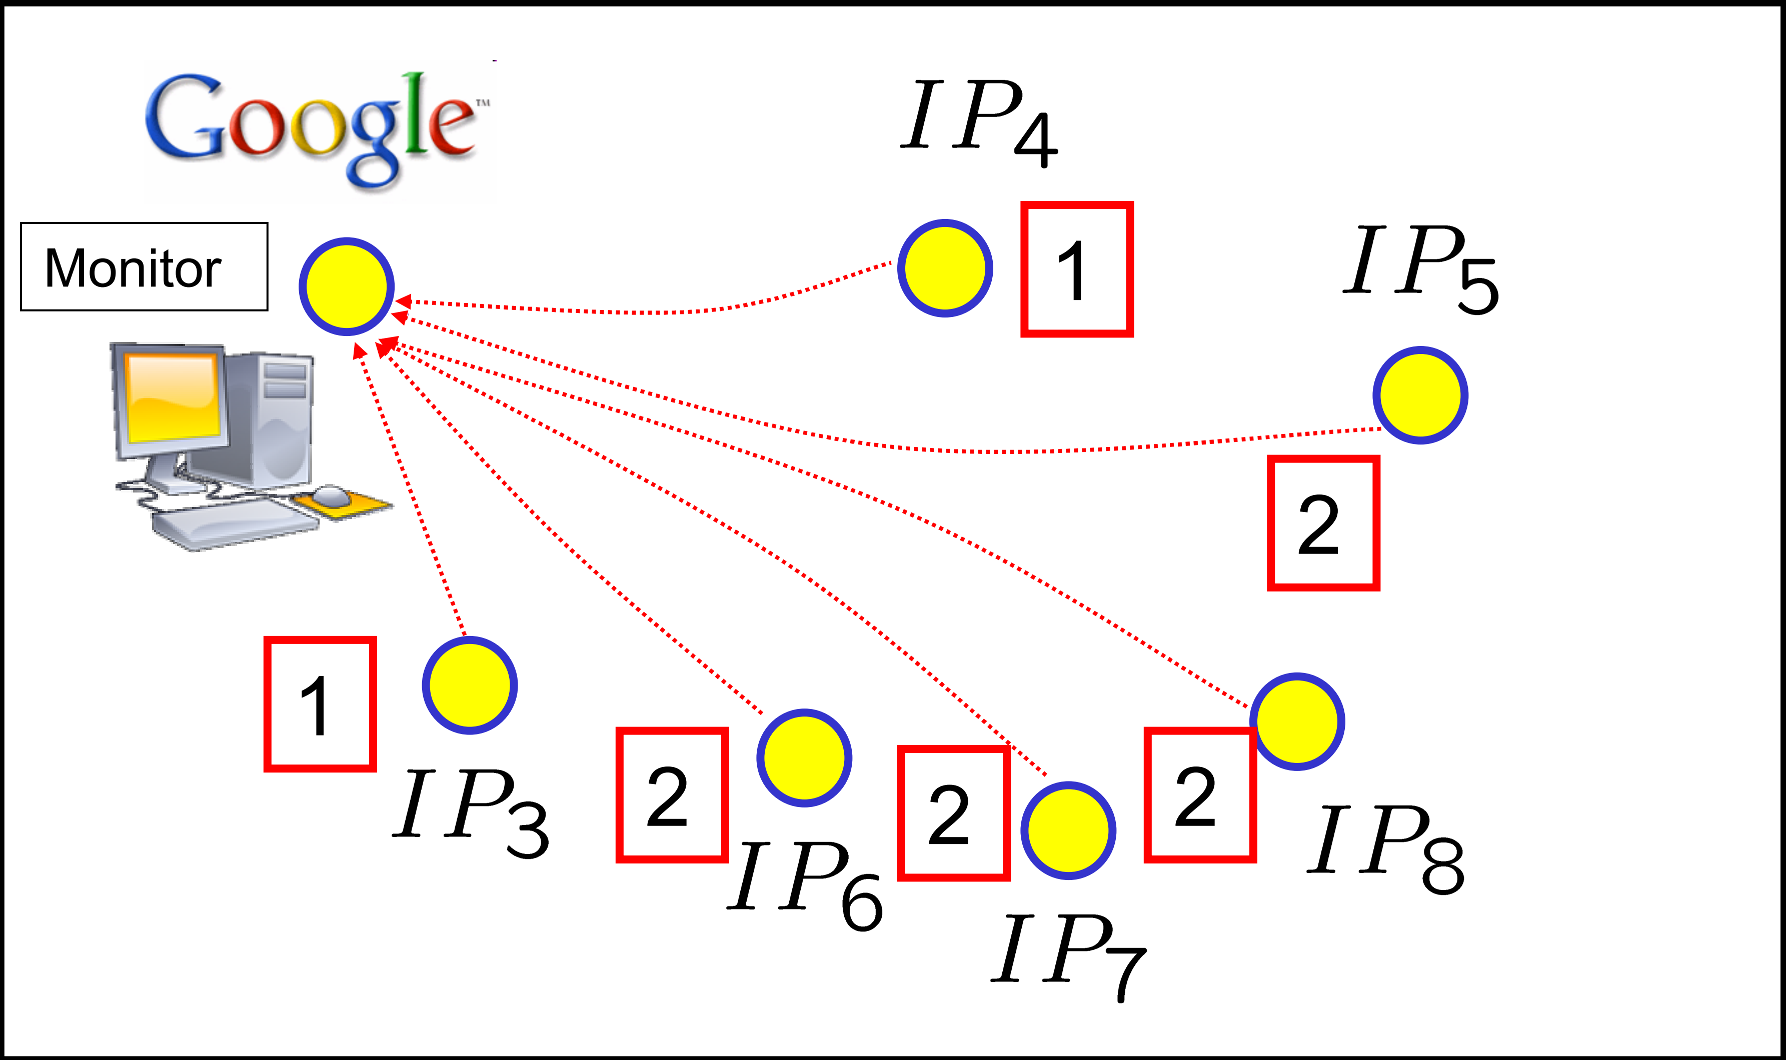

Passive Measurements

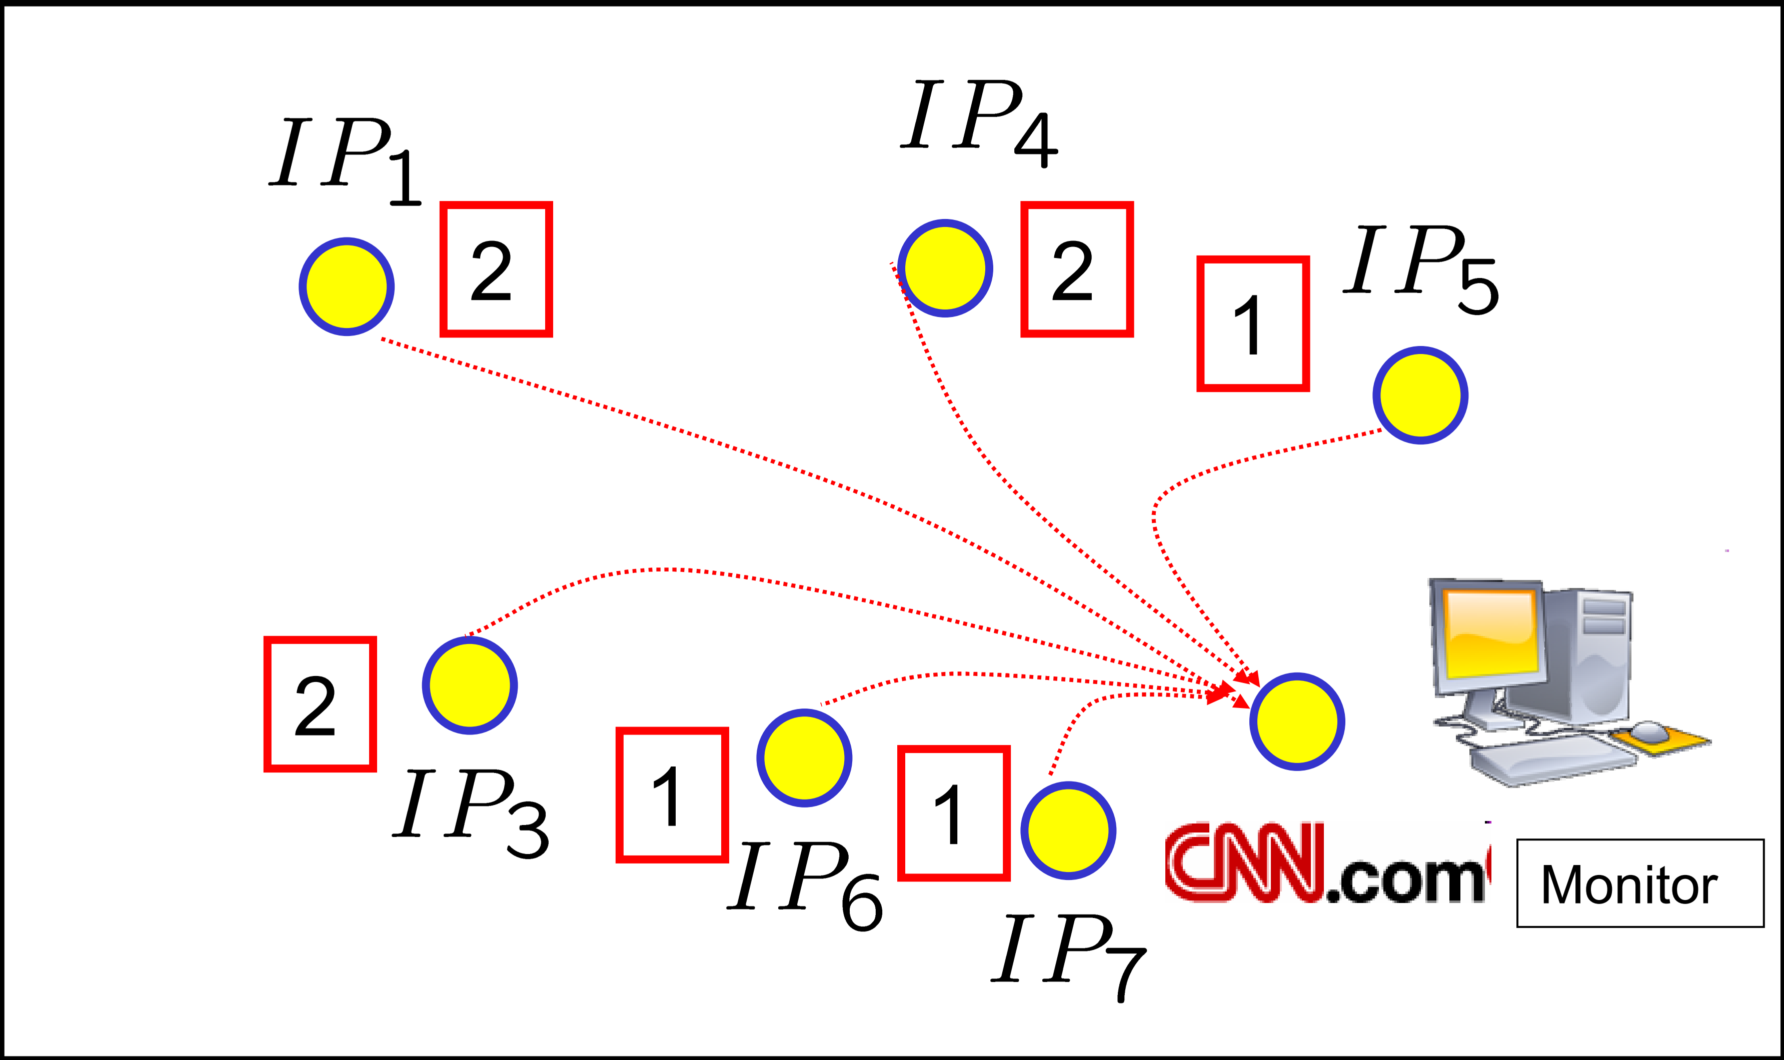

Passive Measurements

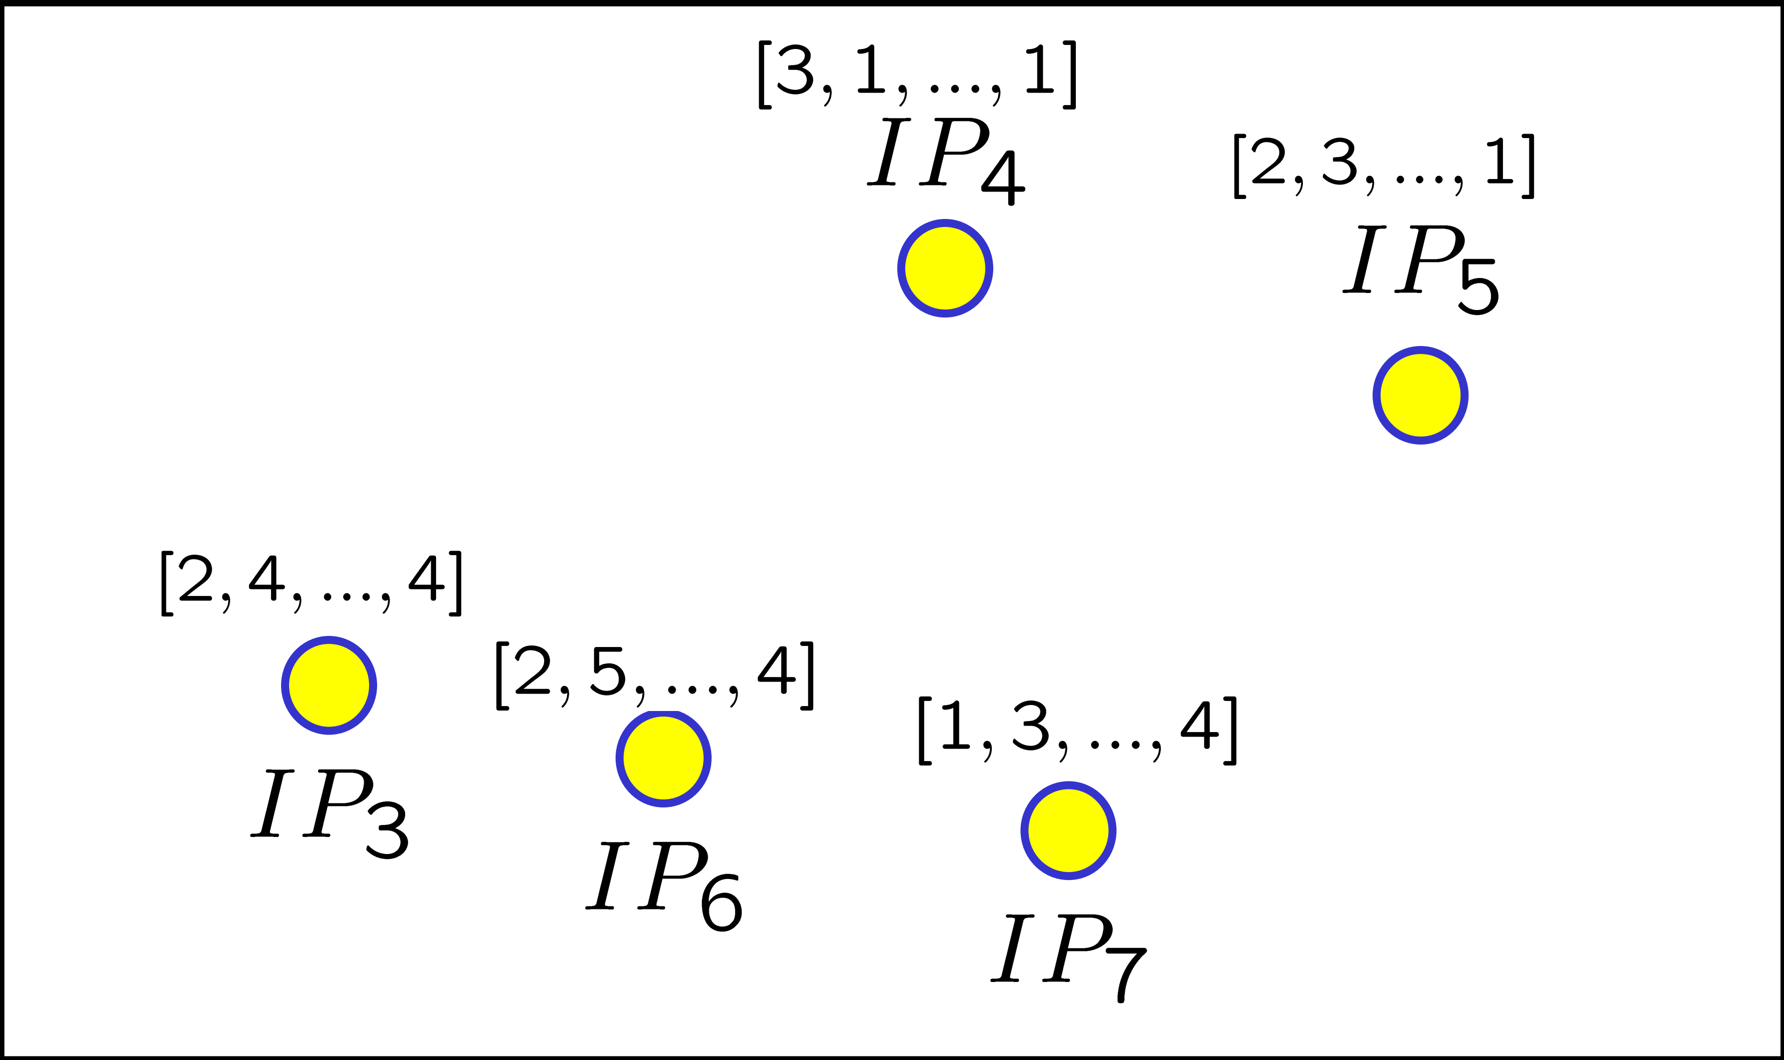

Incomplete Passive Measurements

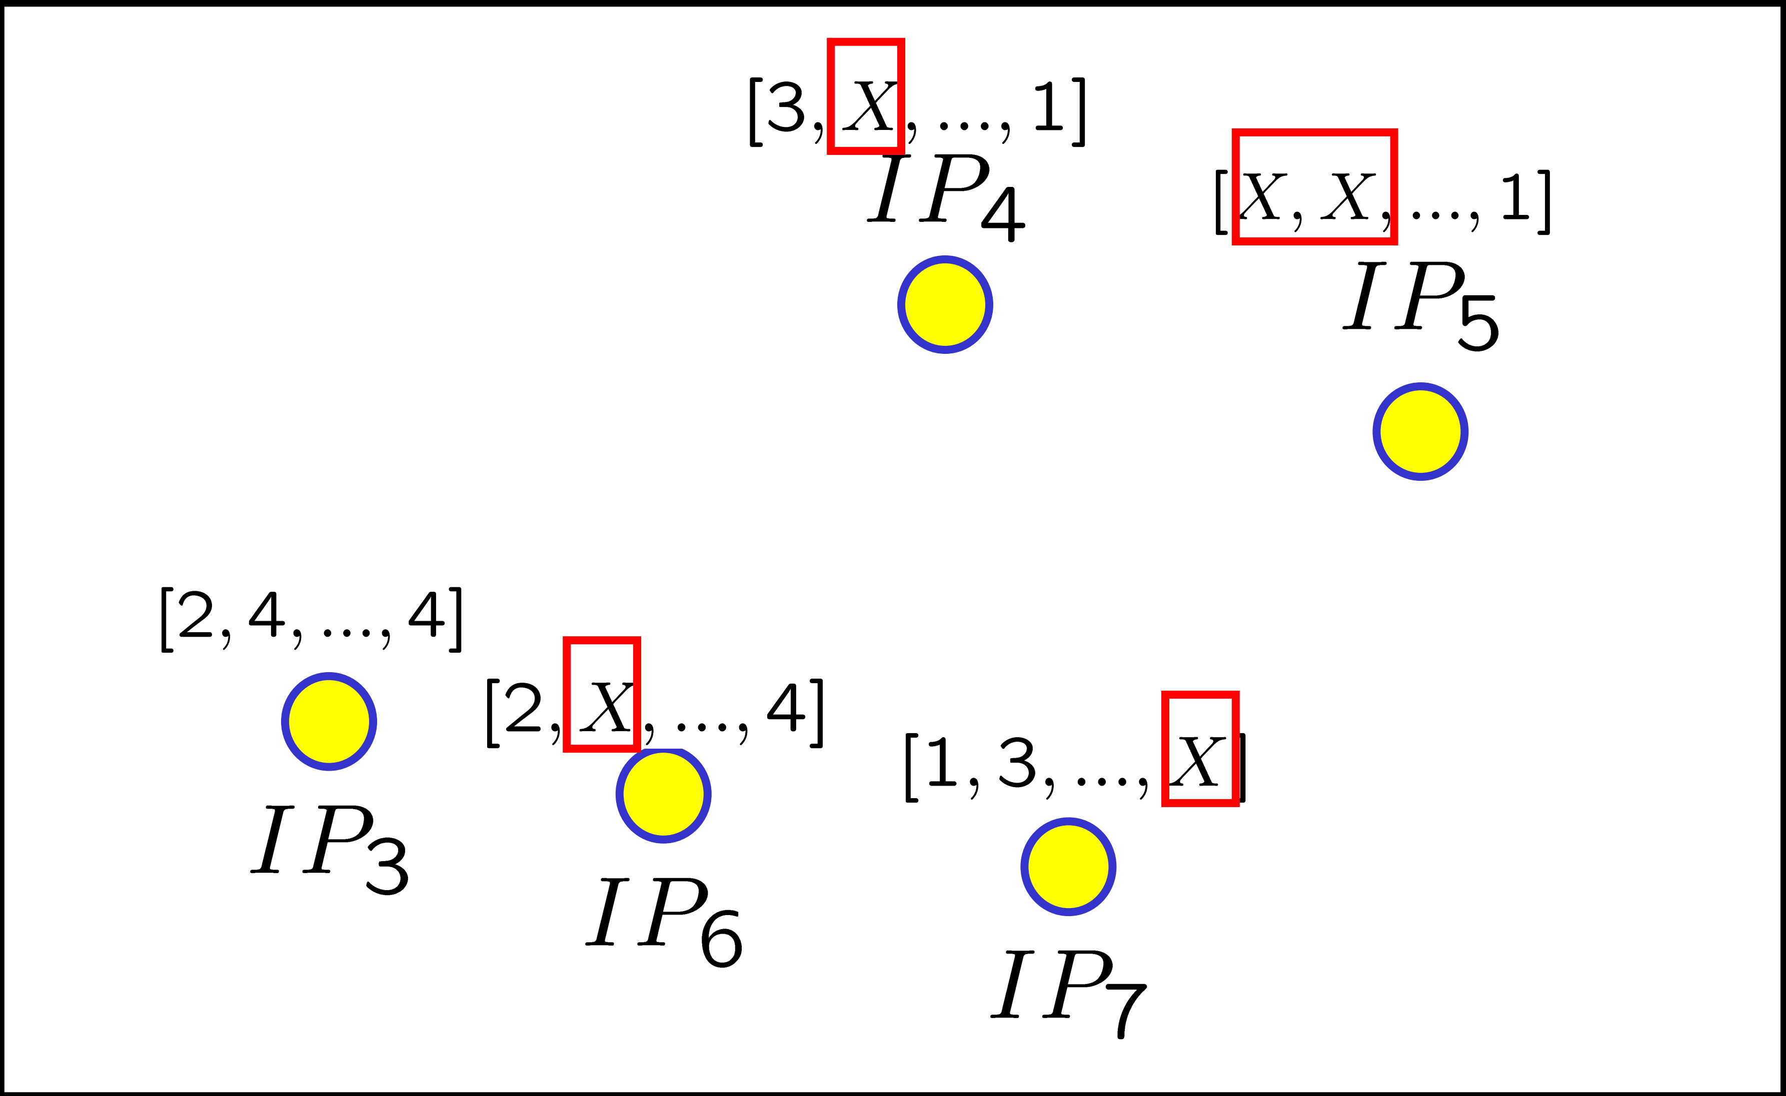

End to End Estimation

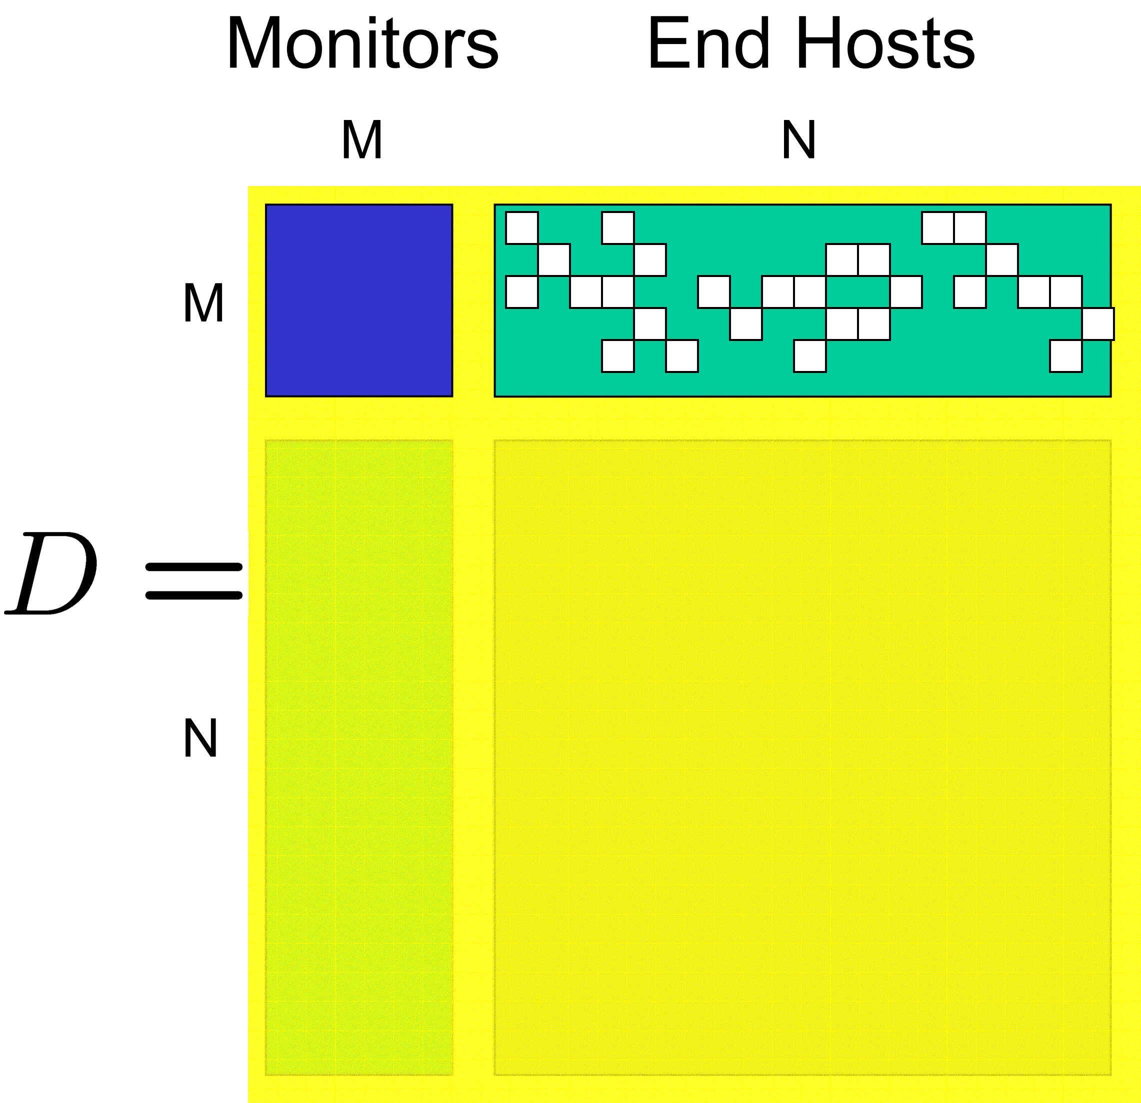

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Incomplete (passive)

End to End Estimation

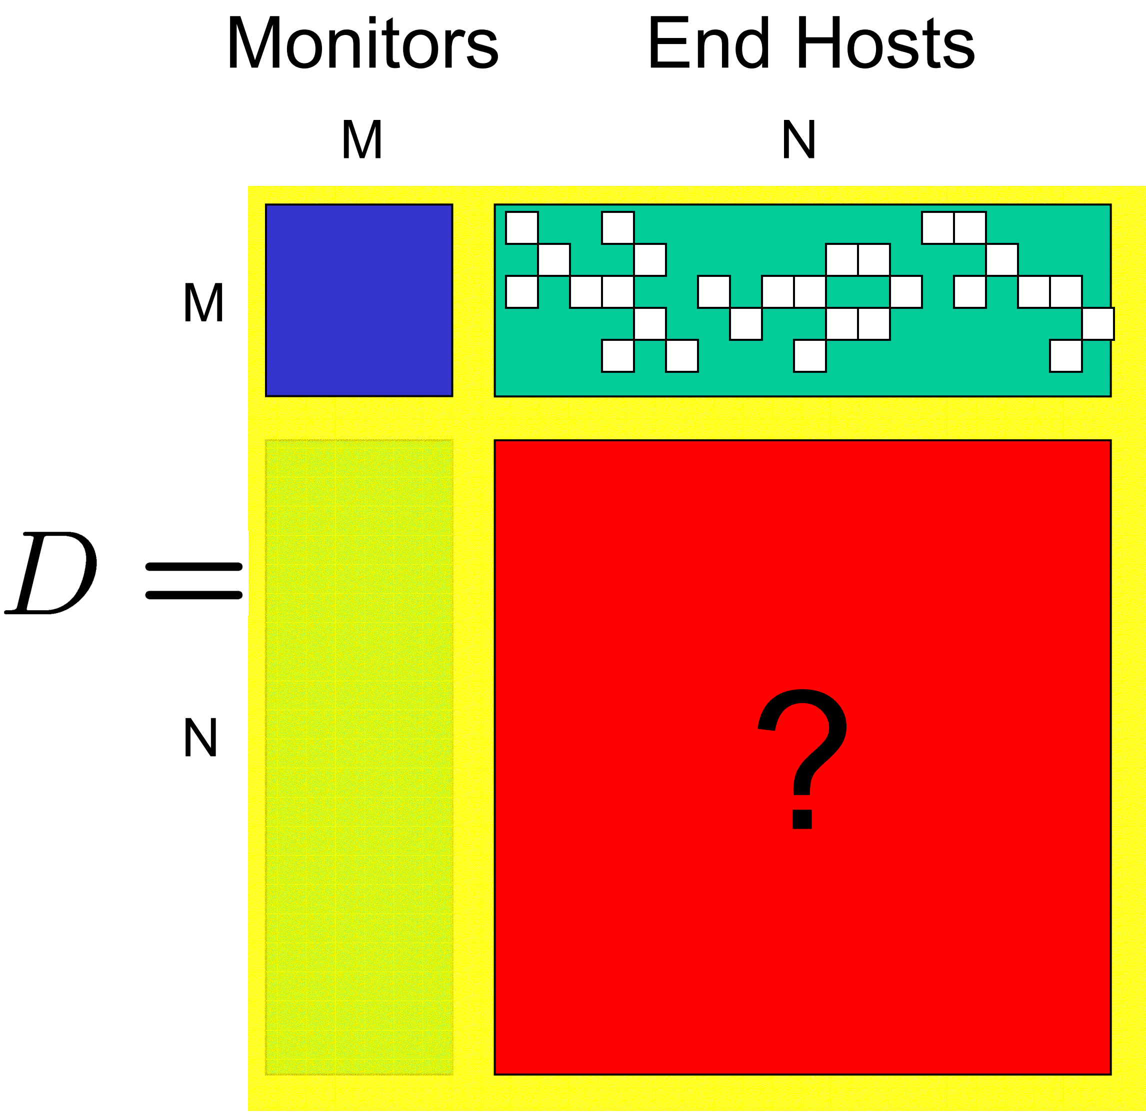

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Incomplete (passive)

Can we still estimate the end host to end host distances?

Dealing with Incomplete Data

- Standard MDS
  - Minimize the sum of squared errors for each embedded point.

Dealing with Incomplete Data

- Standard MDS
  - Minimize the sum of squared errors for each embedded point.
- Incomplete MDS
  - Minimize the <u>weighted</u> sum of squared errors.
  - Weight of each error:
    - Zero, if no hop information observed
    - One, if observed hop count

AS Information

- Consider two end hosts
  - Same observed passive data
  - In different ASes

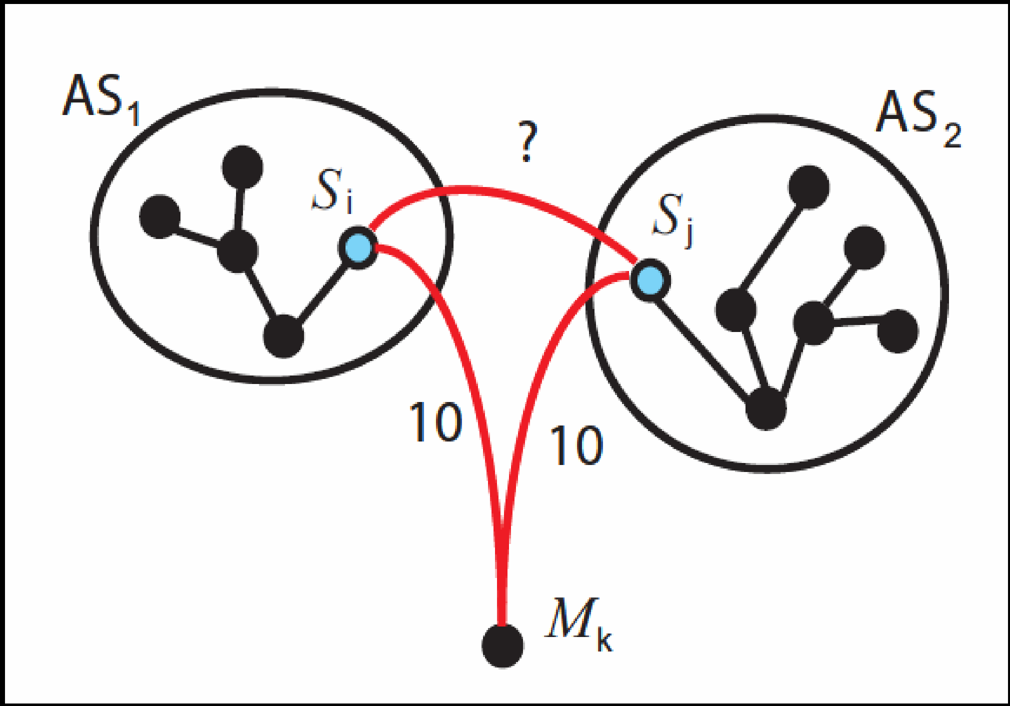

AS Information

- Consider two end hosts
  - Same observed passive data
  - In different ASes

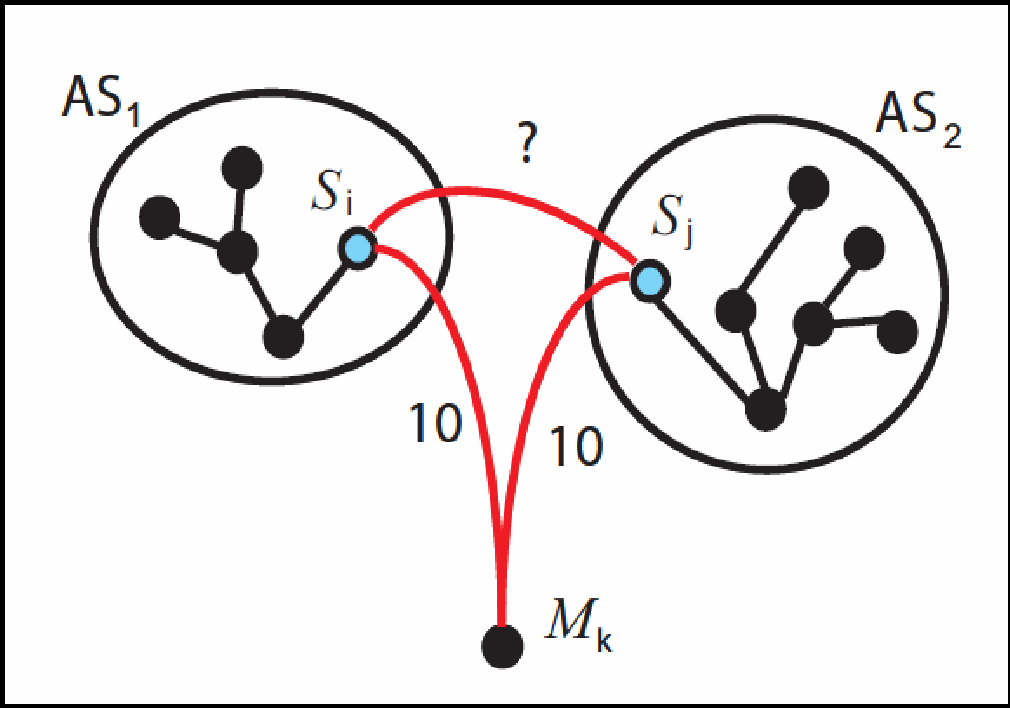

Embedding

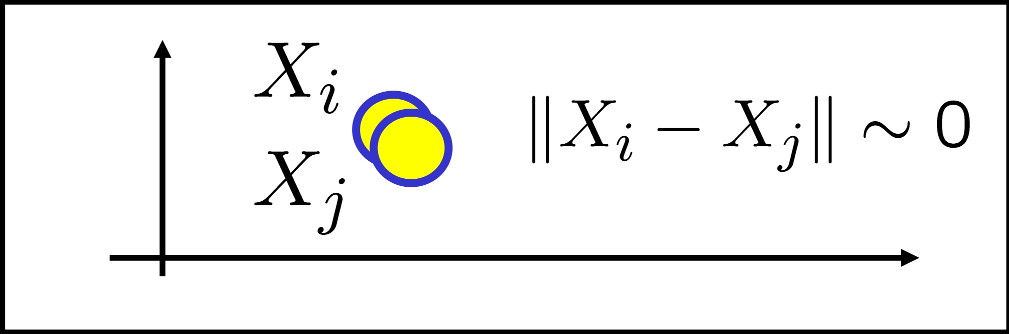

AS Information

- Consider two end hosts
  - Same observed passive data
  - In different ASes

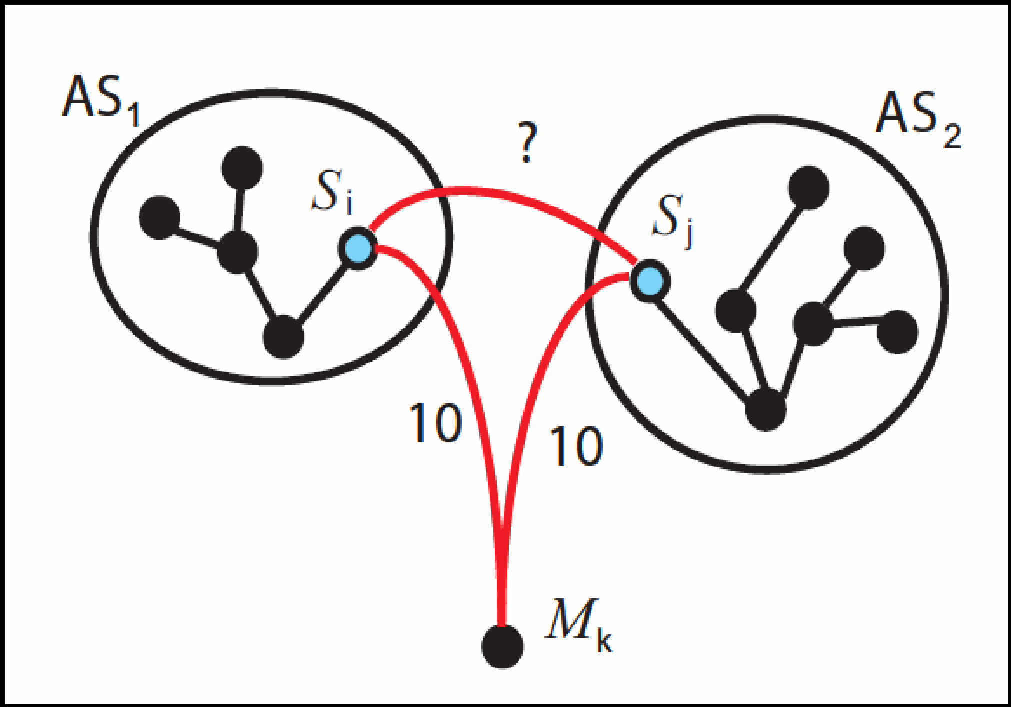

Embedding

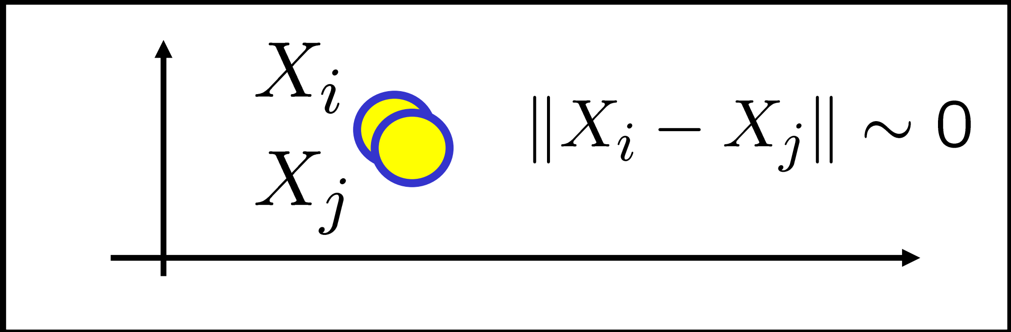

Can we account for AS information?

AS Information

- Given two End Hosts
  - In the same AS
    - Weakly constrain close embedding

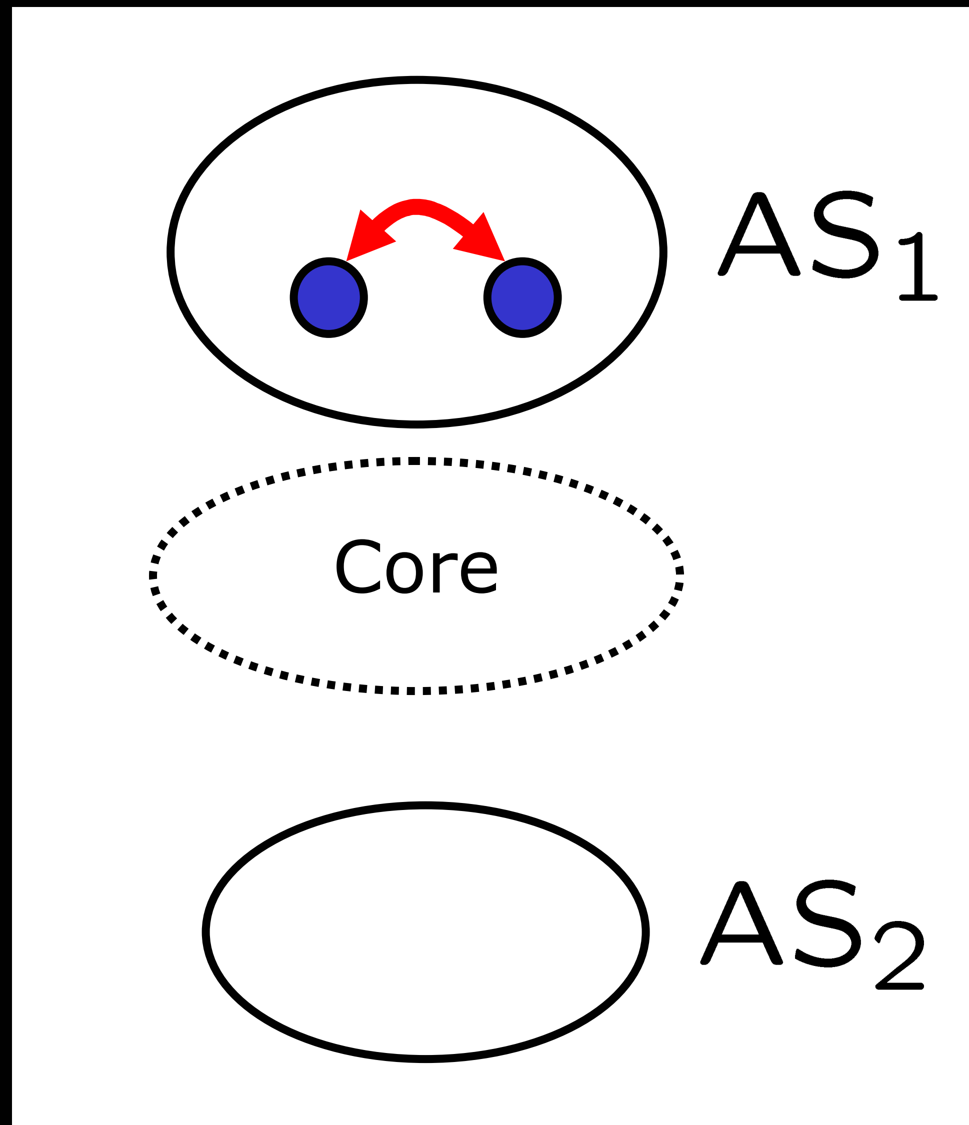

AS Information

- Given two End Hosts
  - In the same AS
    - Weakly constrain close embedding
  - In different ASes
    - Weakly constrain distant embedding

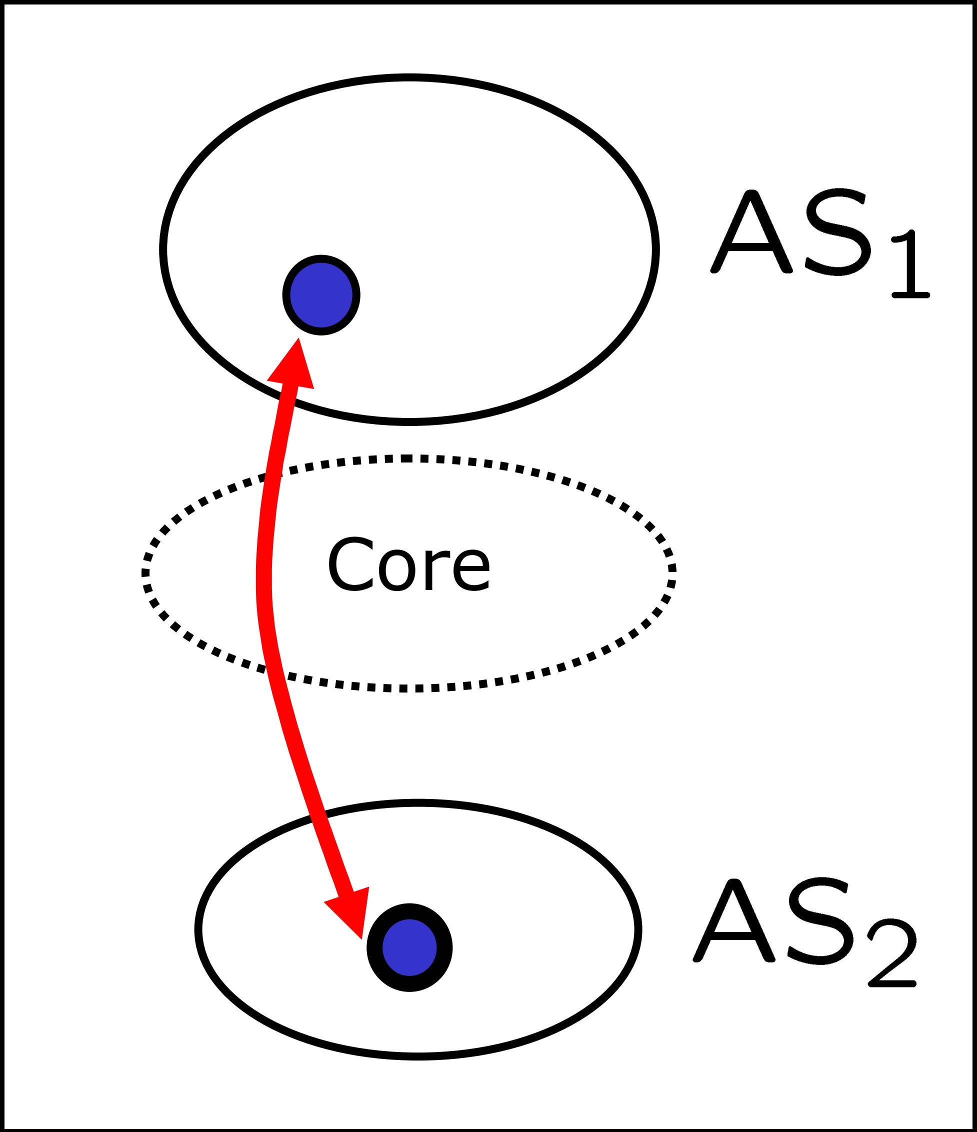

AS Information

- Minimize the embedding stress given:

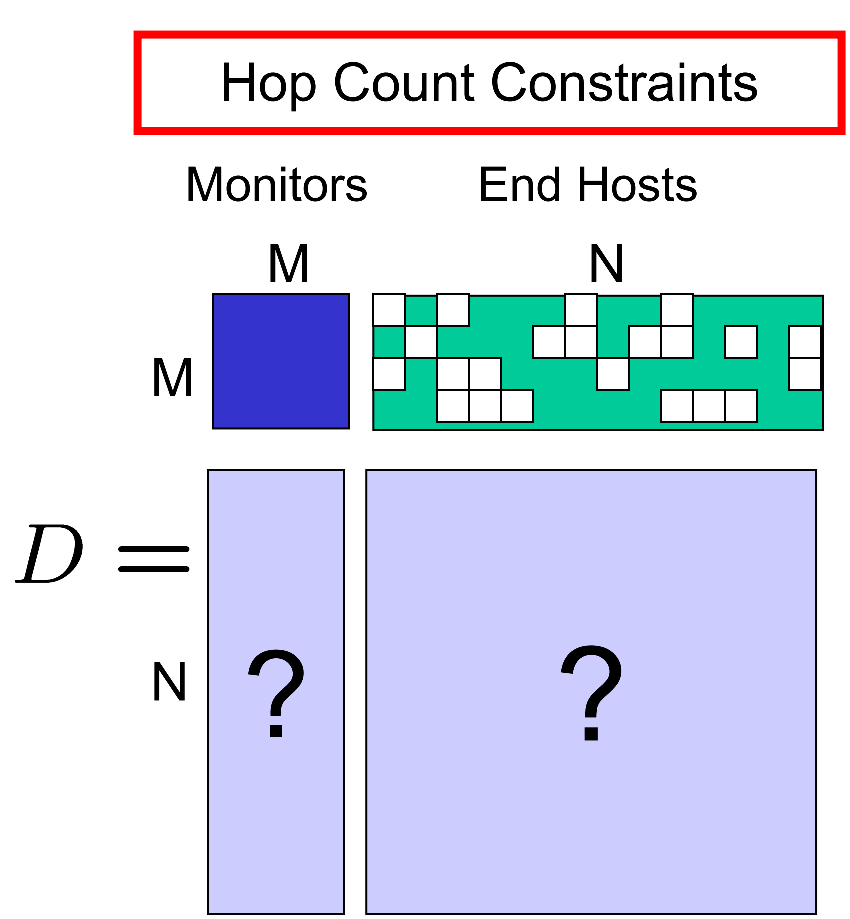

AS Information

- Minimize the embedding stress given:

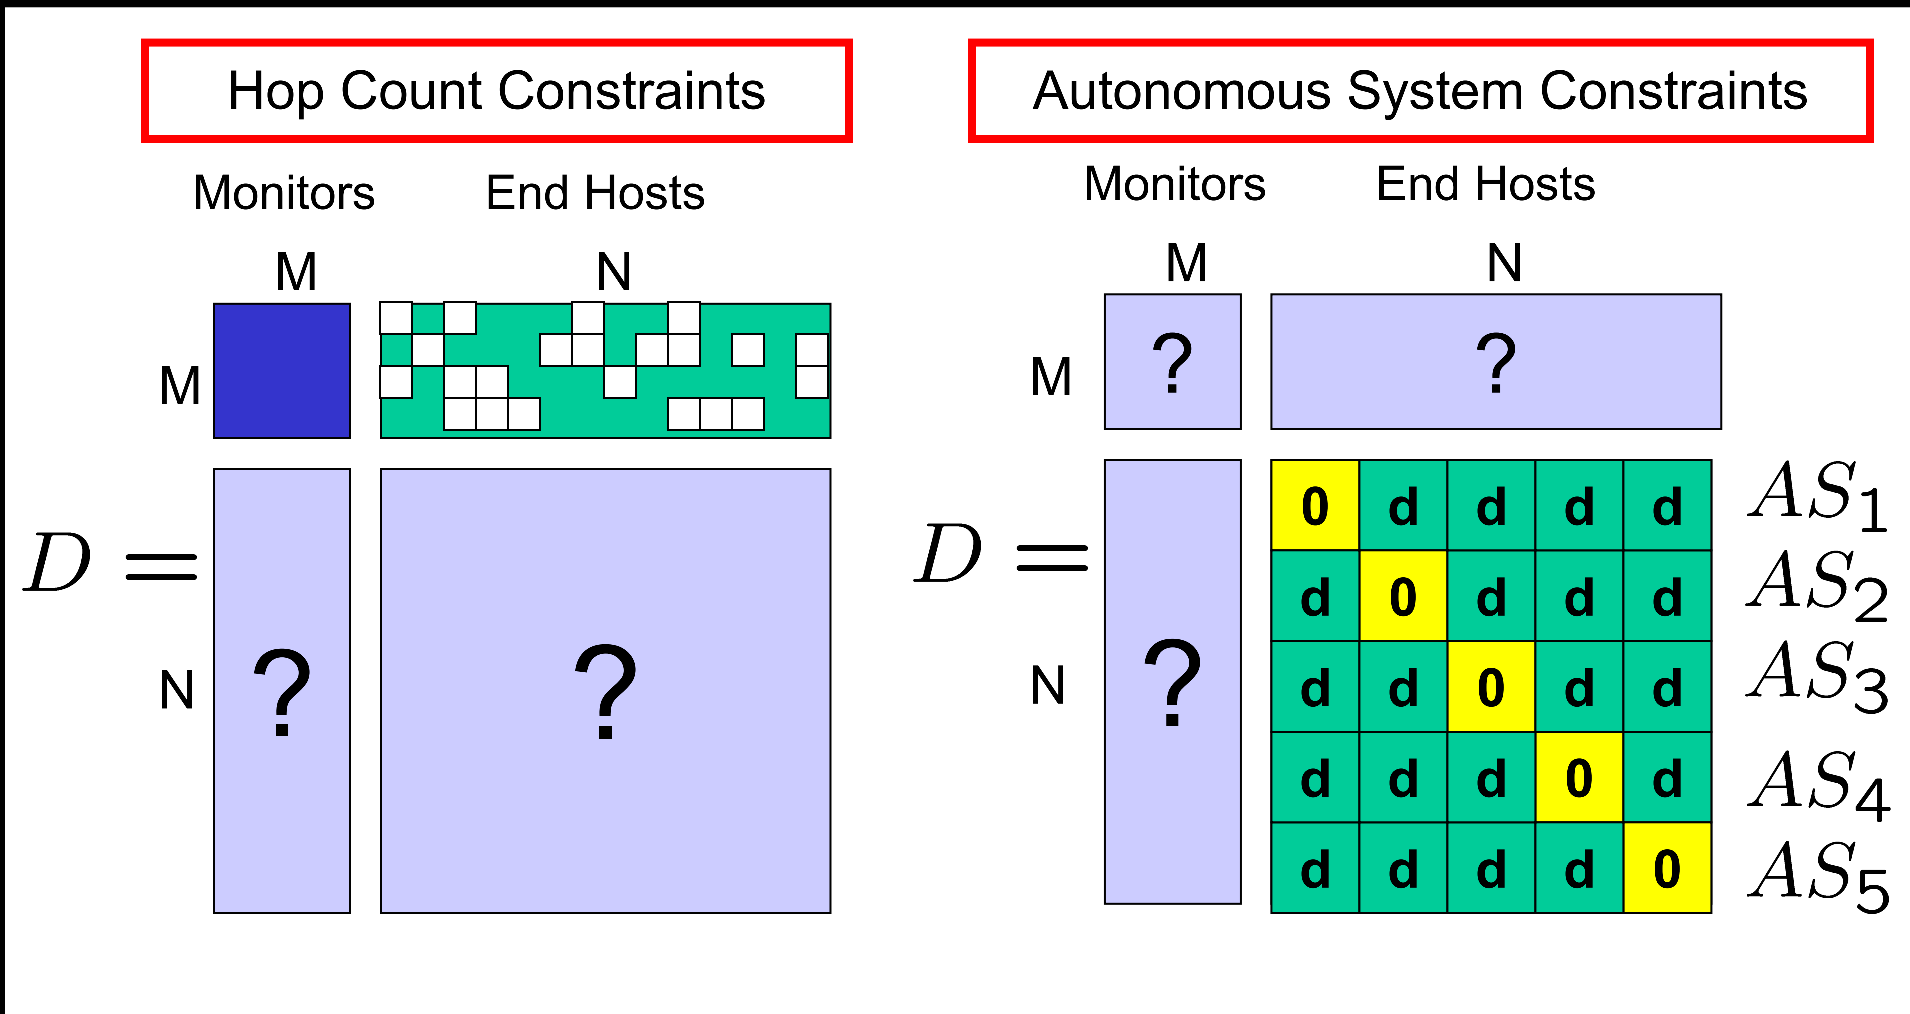

Experimental Setup

- Performance on two topologies

P. Madadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat, “Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies,” in *Proceedings of ACM SIGCOMM ’07*, Kyoto, Japan, August 2007.

Experimental Setup

- Performance on two topologies
  - Synthetic Topology
    - Generated using Orbis Tool
    - Multiple topology sizes
    - AS Information assigned by nearest neighbor

P. Madadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat, "Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies," in *Proceedings of ACM SIGCOMM '07*, Kyoto, Japan, August 2007.

Experimental Setup

- Performance on two topologies
  - Synthetic Topology
    - Generated using Orbis Tool
    - Multiple topology sizes
    - AS Information assigned by nearest neighbor
  - Skitter

P. Madadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat, “Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies,” in *Proceedings of ACM SIGCOMM ’07*, Kyoto, Japan, August 2007.

Experimental Setup

- Performance on two topologies
  - Synthetic Topology
    - Generated using Orbis Tool
    - Multiple topology sizes
    - AS Information assigned by nearest neighbor
  - Skitter
- Weight of AS Information and AS distance assigned by Cross Validation of observed hop values

P. Madadevan, C. Hubble, D. Krioukov, B. Huffaker, and A. Vahdat, “Orbis: Rescaling Degree Correlations to Generate Annotated Internet Topologies,” in *Proceedings of ACM SIGCOMM ’07*, Kyoto, Japan, August 2007.

Synthetic (Orbis) - Estimation Error Results

End Hosts = 2000   # Monitors = 16   #AS = 10   dim = 5

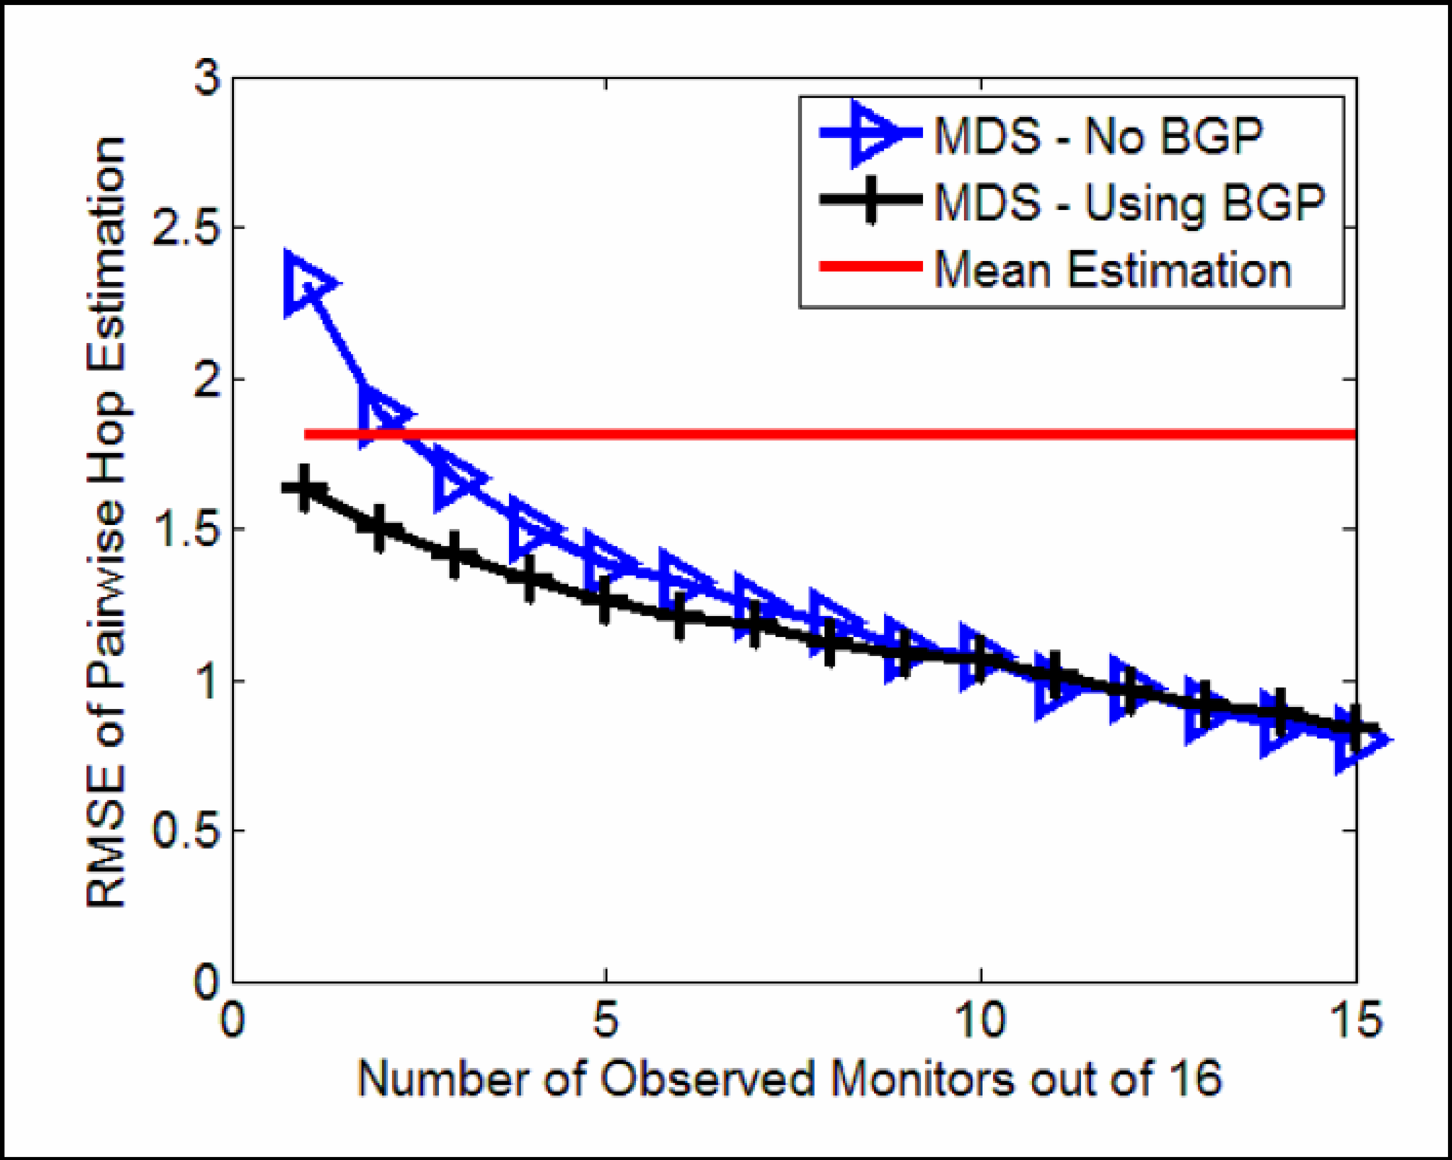

Dimensionality Results

- # End Hosts = 2000   # Monitors = 32

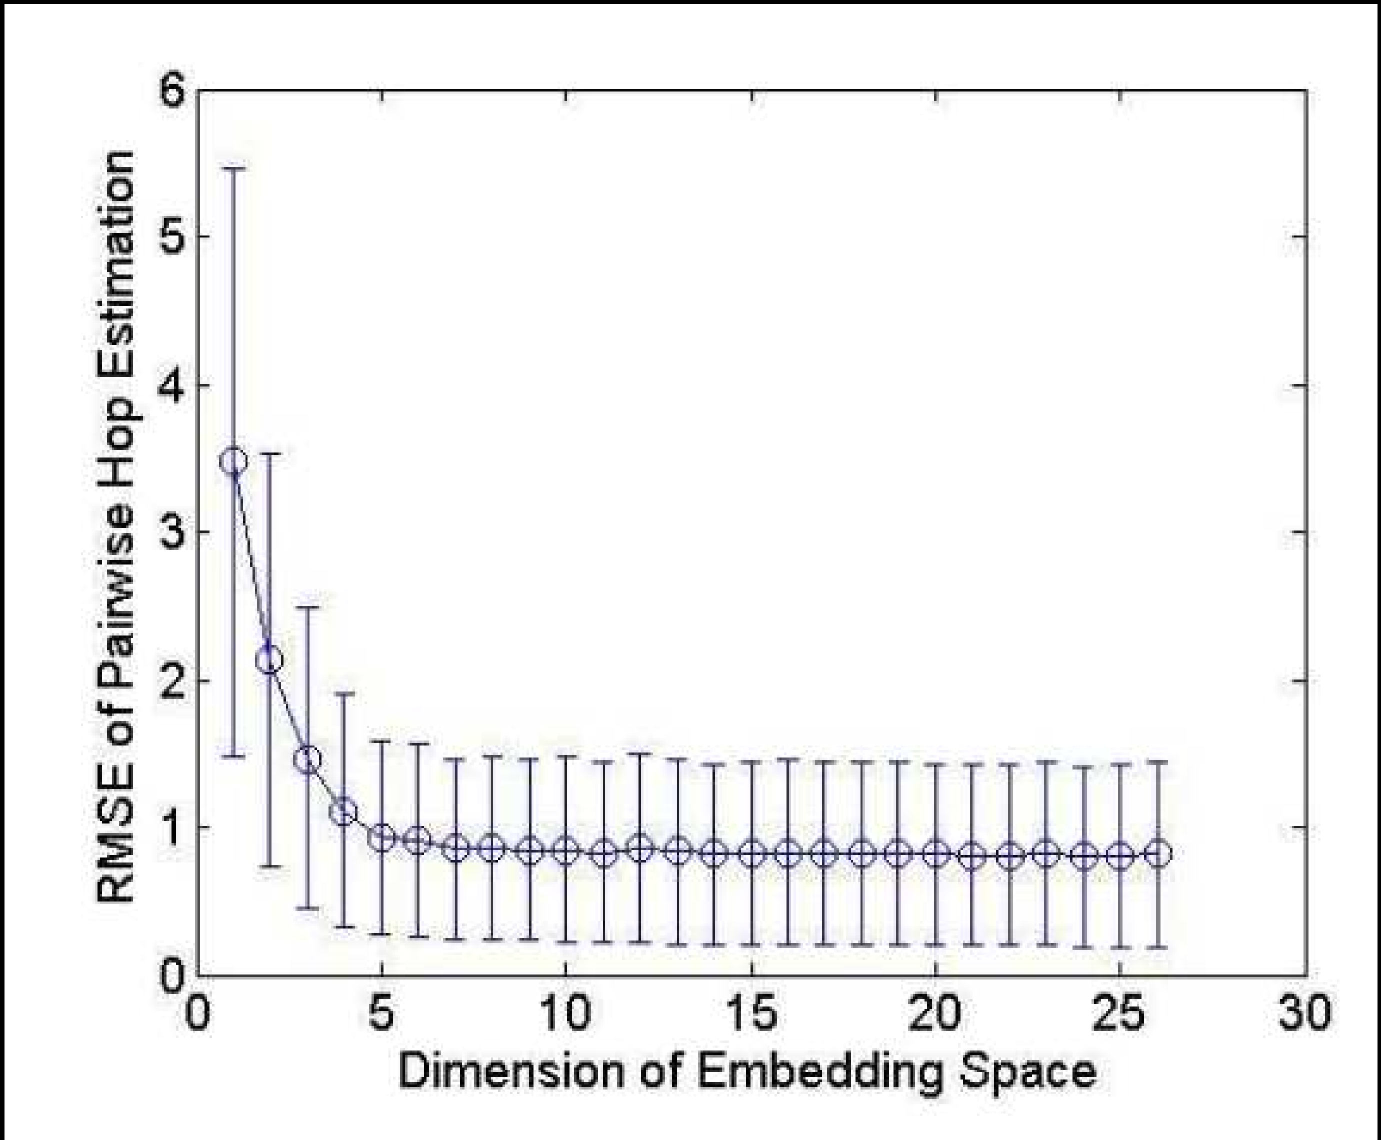

Conclusions

- Presented a methodology to estimate hop distances between arbitrary end hosts

Conclusions

- Presented a methodology to estimate hop distances between arbitrary end hosts
- Using passive measurements, we can significantly reduce the number of active probes
  - Although, Incomplete Data makes embedding more difficult

Conclusions

- Presented a methodology to estimate hop distances between arbitrary end hosts
- Using passive measurements, we can significantly reduce the number of active probes
  - Although, Incomplete Data makes embedding more difficult
- Using AS information, the embedding performance can be improved.

Questions?

Reverse Path Results

- # End Hosts = 2000   # Monitors = 16   #AS = 10

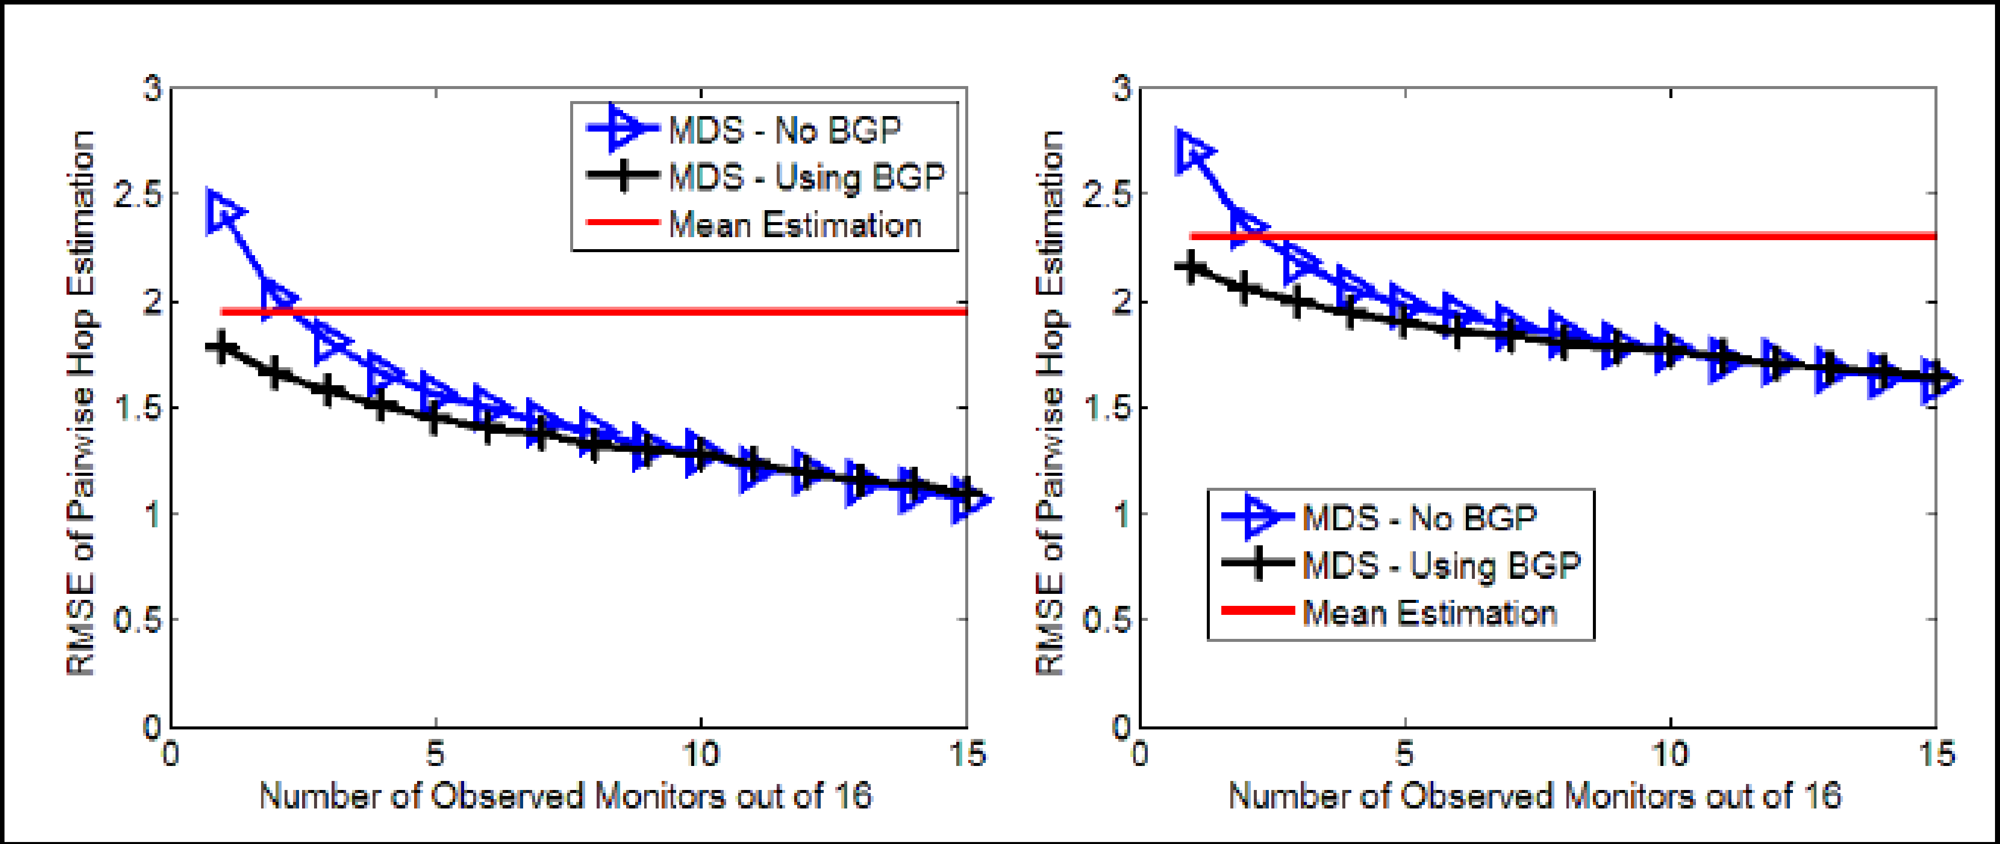

Off by one hop   Off by two hops

LMDS - Incomplete Data

- Standard Multidimensional Scaling:

$$\text{Find} \quad \text{argmin}_{\mathbf{X}} \quad \text{stress}\,(\mathbf{X})$$

LMDS - Incomplete Data

- Standard Multidimensional Scaling:

Find $\operatorname{argmin}_{\mathbf{X}}$ stress $(\mathbf{X})$

$$\text{stress}(\mathbf{X}) = \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \left(\widehat{D}_{i,j} - D_{i,j}\right)^2$$

$$= \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \left(\|X_i - X_j\| - D_{i,j}\right)^2$$

LMDS - Incomplete Data

- Standard Multidimensional Scaling:

$$\text{Find} \quad \text{argmin}_{\mathbf{X}} \quad \text{stress}\,(\mathbf{X})$$

$$\begin{aligned} \text{stress}(\mathbf{X}) &= \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \left(\widehat{D}_{i,j} - D_{i,j}\right)^2 \\ &= \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \left(\|X_i - X_j\| - D_{i,j}\right)^2 \end{aligned}$$

What happens when you don't know all the distance values?

LMDS - Incomplete Data

- Modified Multidimensional Scaling:

$$\text{Find} \quad \text{argmin}_{\mathbf{X}} \quad \text{stress}\,(\mathbf{X})$$

LMDS - Incomplete Data

- Modified Multidimensional Scaling:

$$\text{Find} \quad \text{argmin}_{\mathbf{X}} \quad \text{stress}\,(\mathbf{X})$$

$$\text{stress}(\mathbf{X}) = \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} W_{i,j} \left(\widehat{D}_{i,j} - D_{i,j}\right)^2$$

$$= \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} W_{i,j} \left(\|X_i - X_j\| - D_{i,j}\right)^2$$

LMDS - Incomplete Data

- Modified Multidimensional Scaling:

$$\text{Find } \operatorname{argmin}_{\mathbf{X}} \ \text{stress}(\mathbf{X})$$

$$\begin{aligned}\text{stress}(\mathbf{X}) &= \sum_{i=1}^{N+M}\sum_{j=1}^{N+M} \boxed{W_{i,j}} \left(\widehat{D}_{i,j} - D_{i,j}\right)^2 \\ &= \sum_{i=1}^{N+M}\sum_{j=1}^{N+M} \boxed{W_{i,j}} \left(\|X_i - X_j\| - D_{i,j}\right)^2\end{aligned}$$

$$\boxed{W_{i,j} = \begin{cases} 1: & \text{if } h_{i,j} \text{ is known} \\ 0: & \text{if } h_{i,j} \text{ is missing} \end{cases}}$$

Multidimensional Scaling

	Atlanta	Chicago	Denver	Houston	LA	Miami	NYC	San Fran	Seattle	Washington DC
Atlanta	0	587	1212	701	1936	604	748	2139	2182	543
Chicago	587	0	920	940	1745	1188	713	1858	1737	597
Denver	1212	920	0	879	831	1726	1631	949	1021	1494
Houston	701	940	879	0	1374	968	1420	1645	1891	1220
LA	1936	1745	831	1374	0	2339	2451	347	959	2300
Miami	604	1188	1726	968	2339	0	1092	2594	2734	923
NYC	748	713	1631	1420	2451	1092	0	2571	2408	205
San Fran	2139	1858	949	1645	347	2594	2571	0	678	2442
Seattle	2182	1737	1021	1891	959	2734	2408	678	0	2329
Washington DC	543	597	1491	1220	2300	923	205	2442	2329	0

Scree Plot

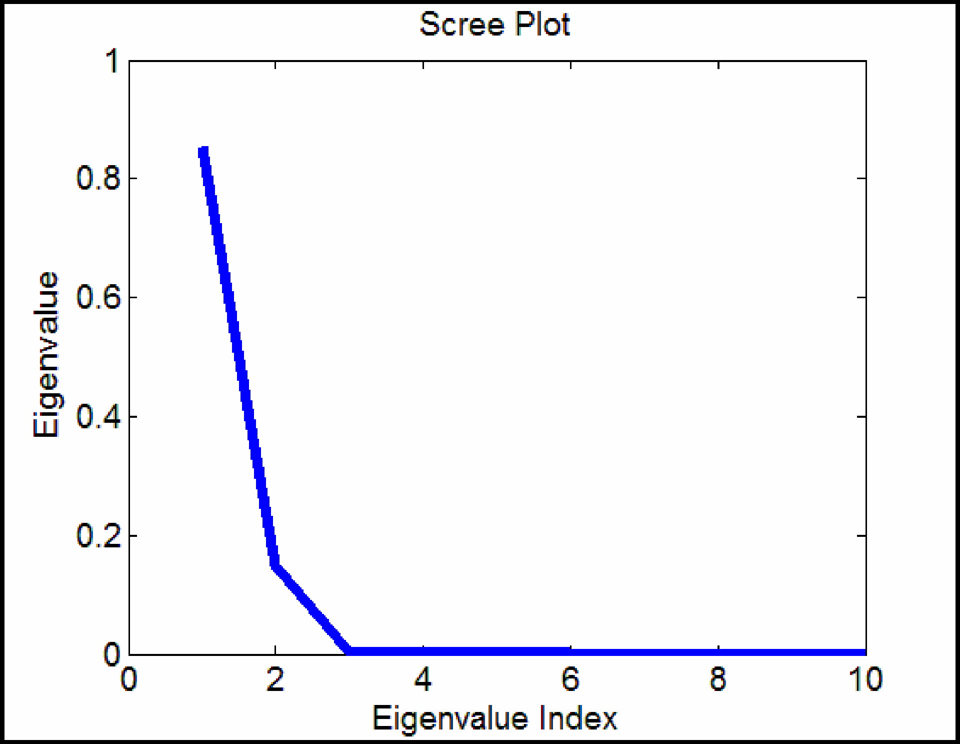

- All the cities lie on the Earth's surface, intuitively 2 dimensions.

Scree Plot

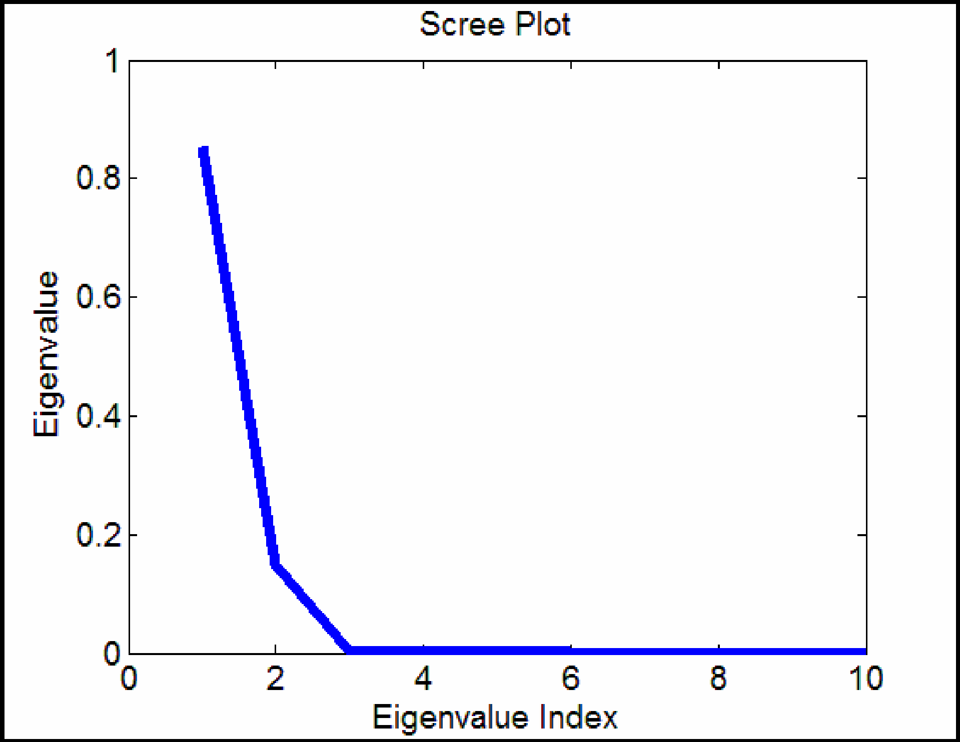

- All the cities lie on the Earth's surface, intuitively 2 dimensions.
- We will choose the dimension based on the number of eigenvalues to capture 90% of the variance

Multidimensional Scaling

- MDS Results

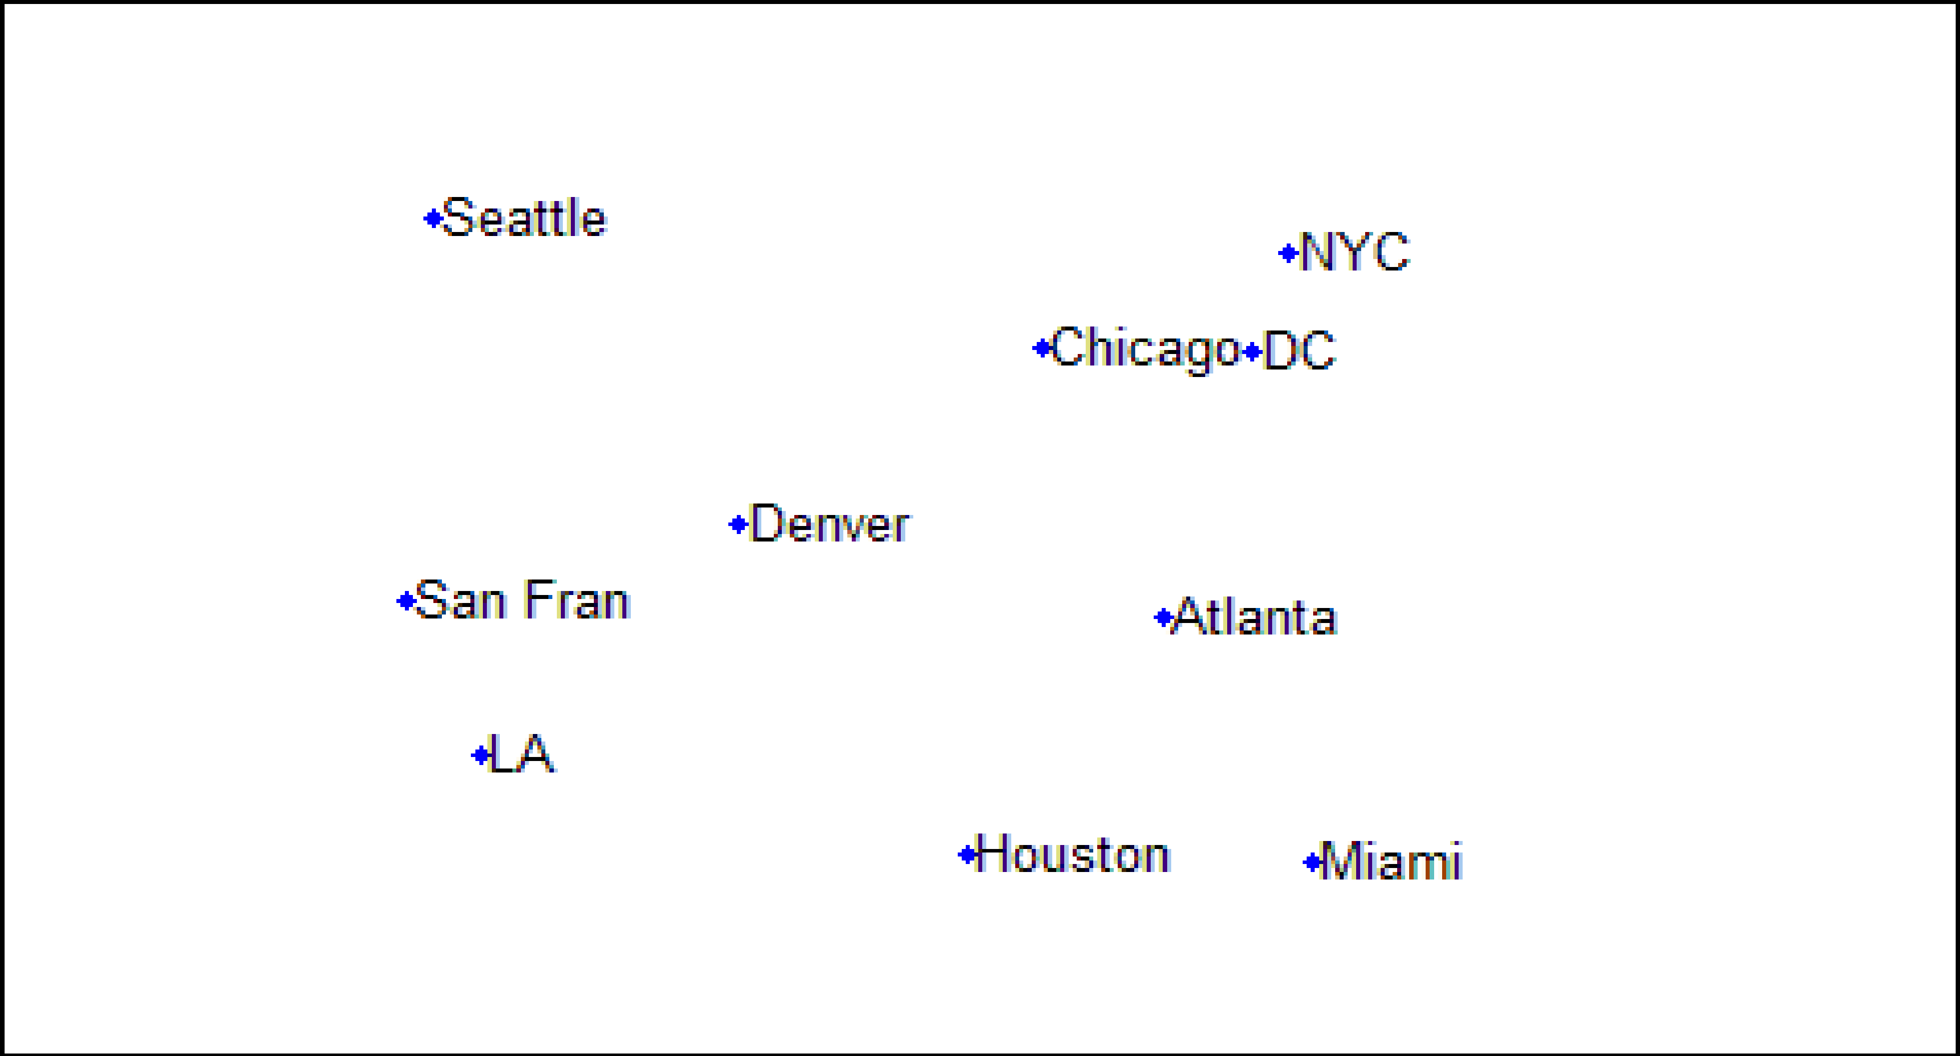

Multidimensional Scaling

- MDS Results

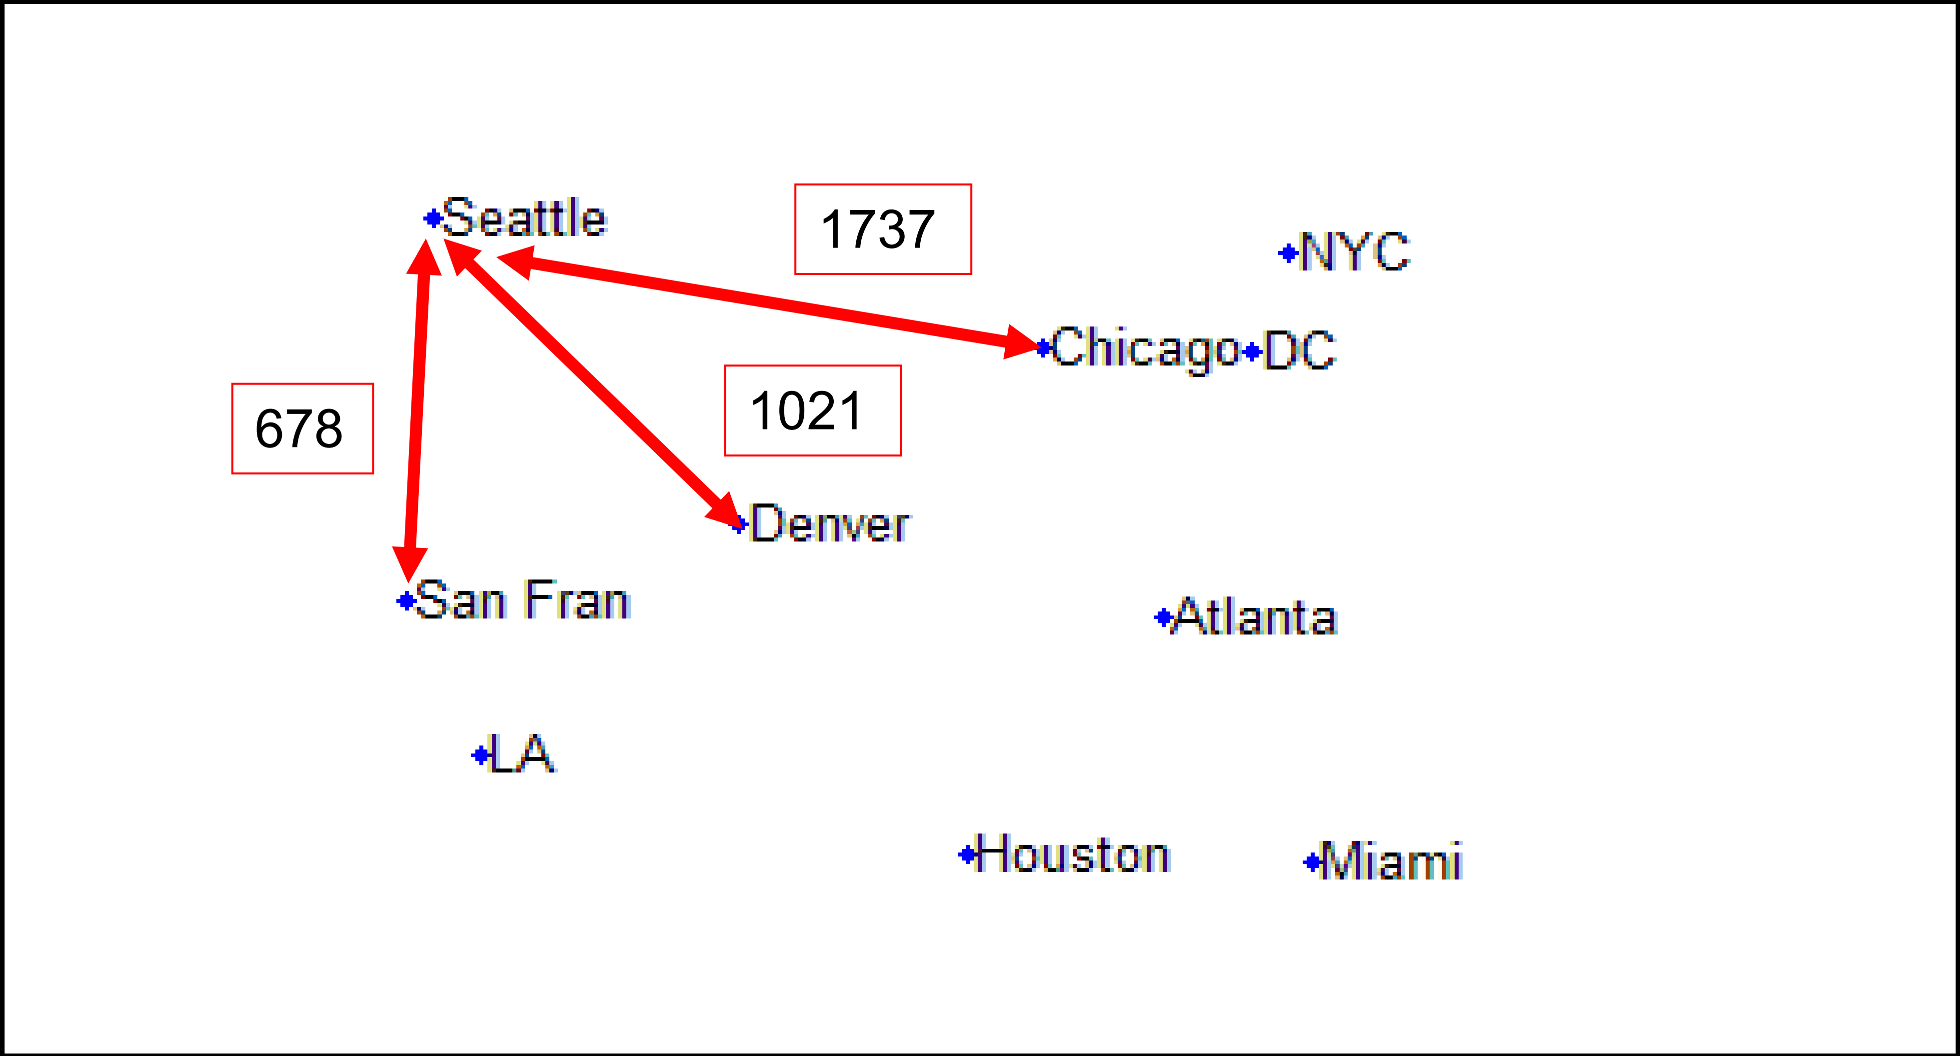

End to End Estimation

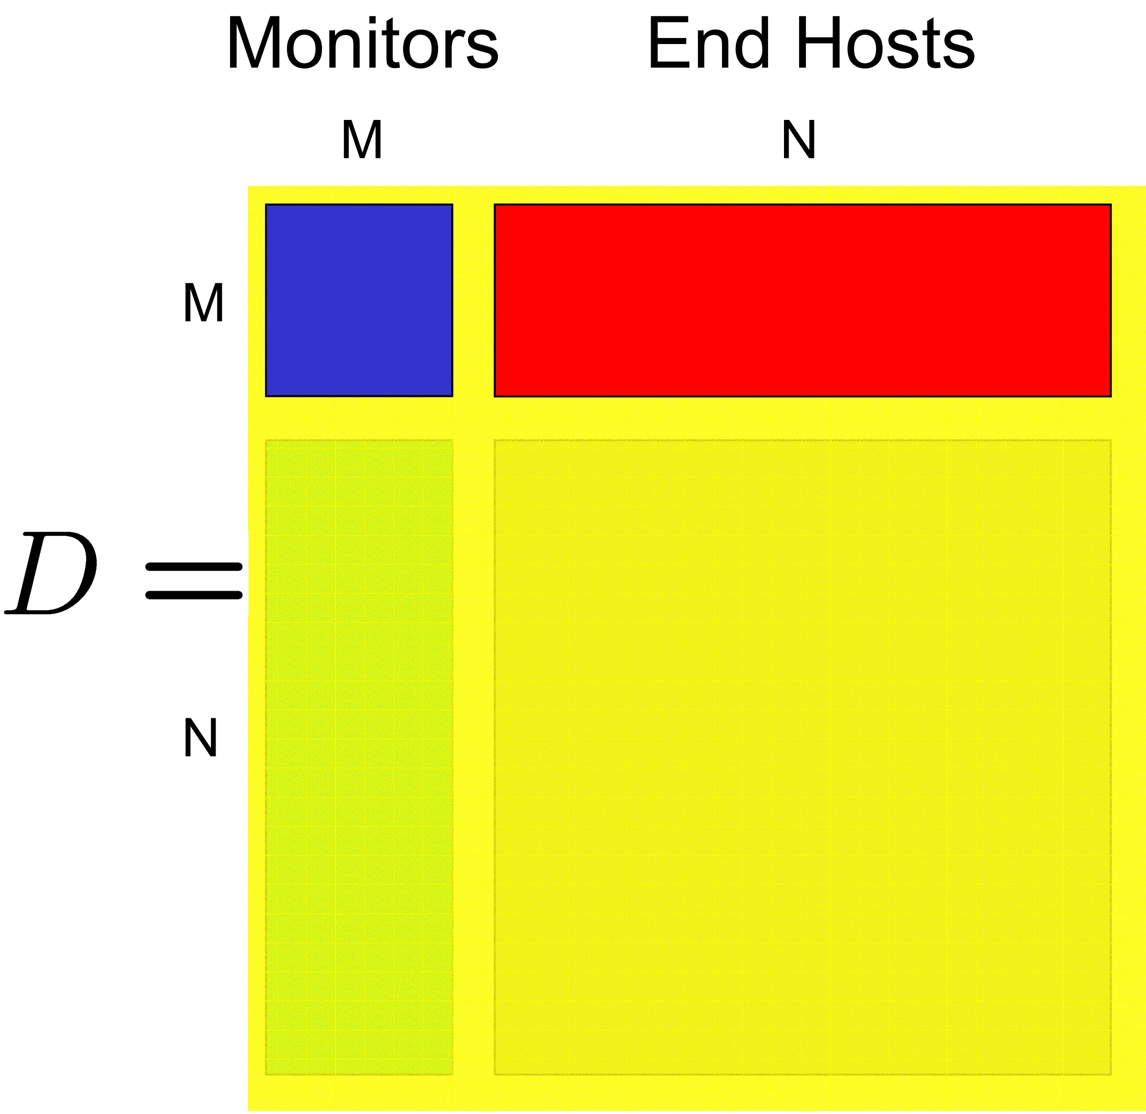

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Known (<u>passive</u>)

End to End Estimation

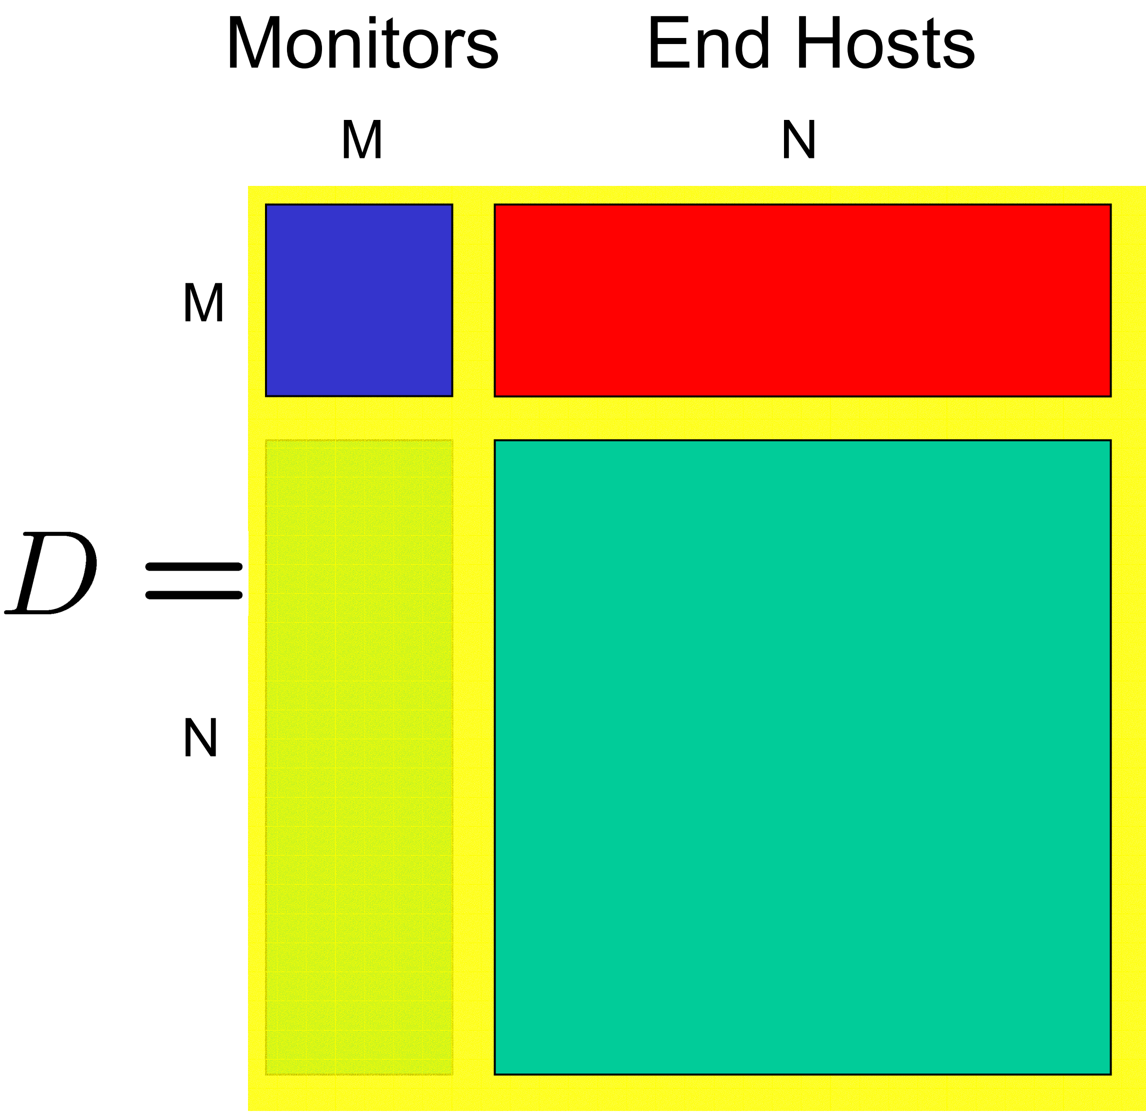

Monitor to Monitor Hop Distances – Known (active)

Monitor to End Host Distances – Known (passive)

End host to end host distances – Estimated by MDS

Passive Measurements

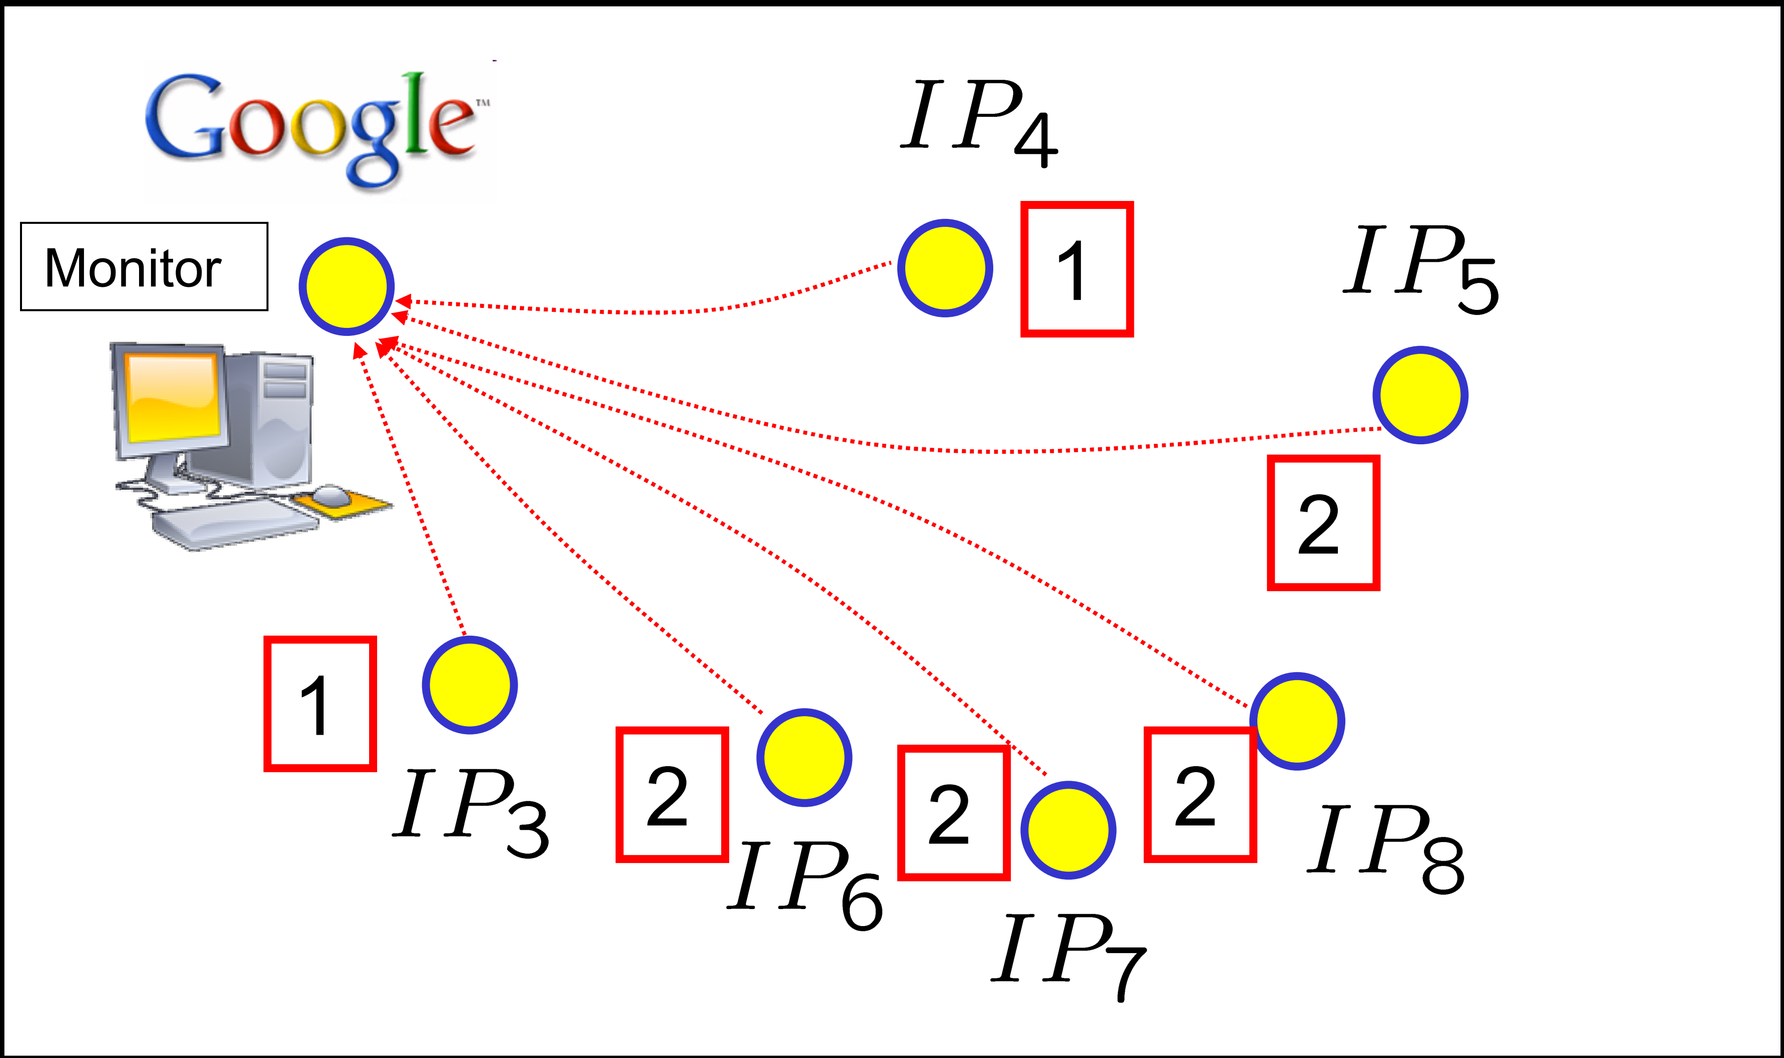

Incomplete Passive Measurements

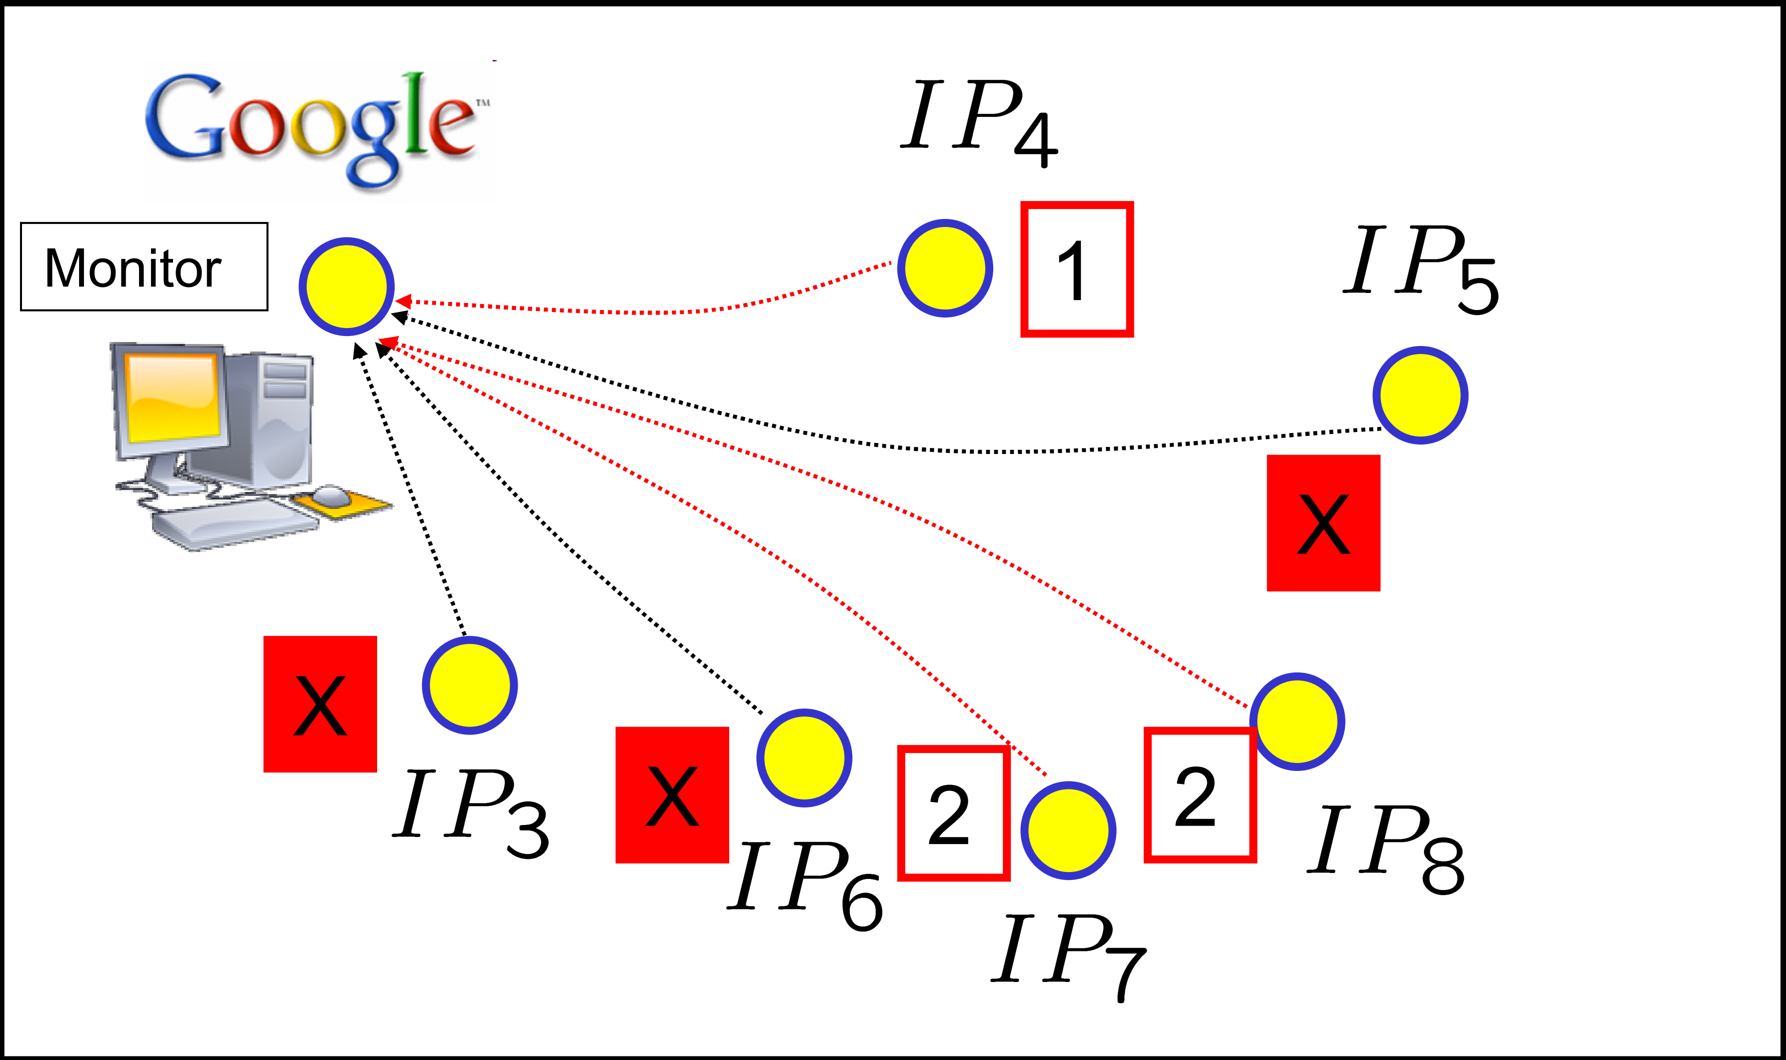

Passive Measurements

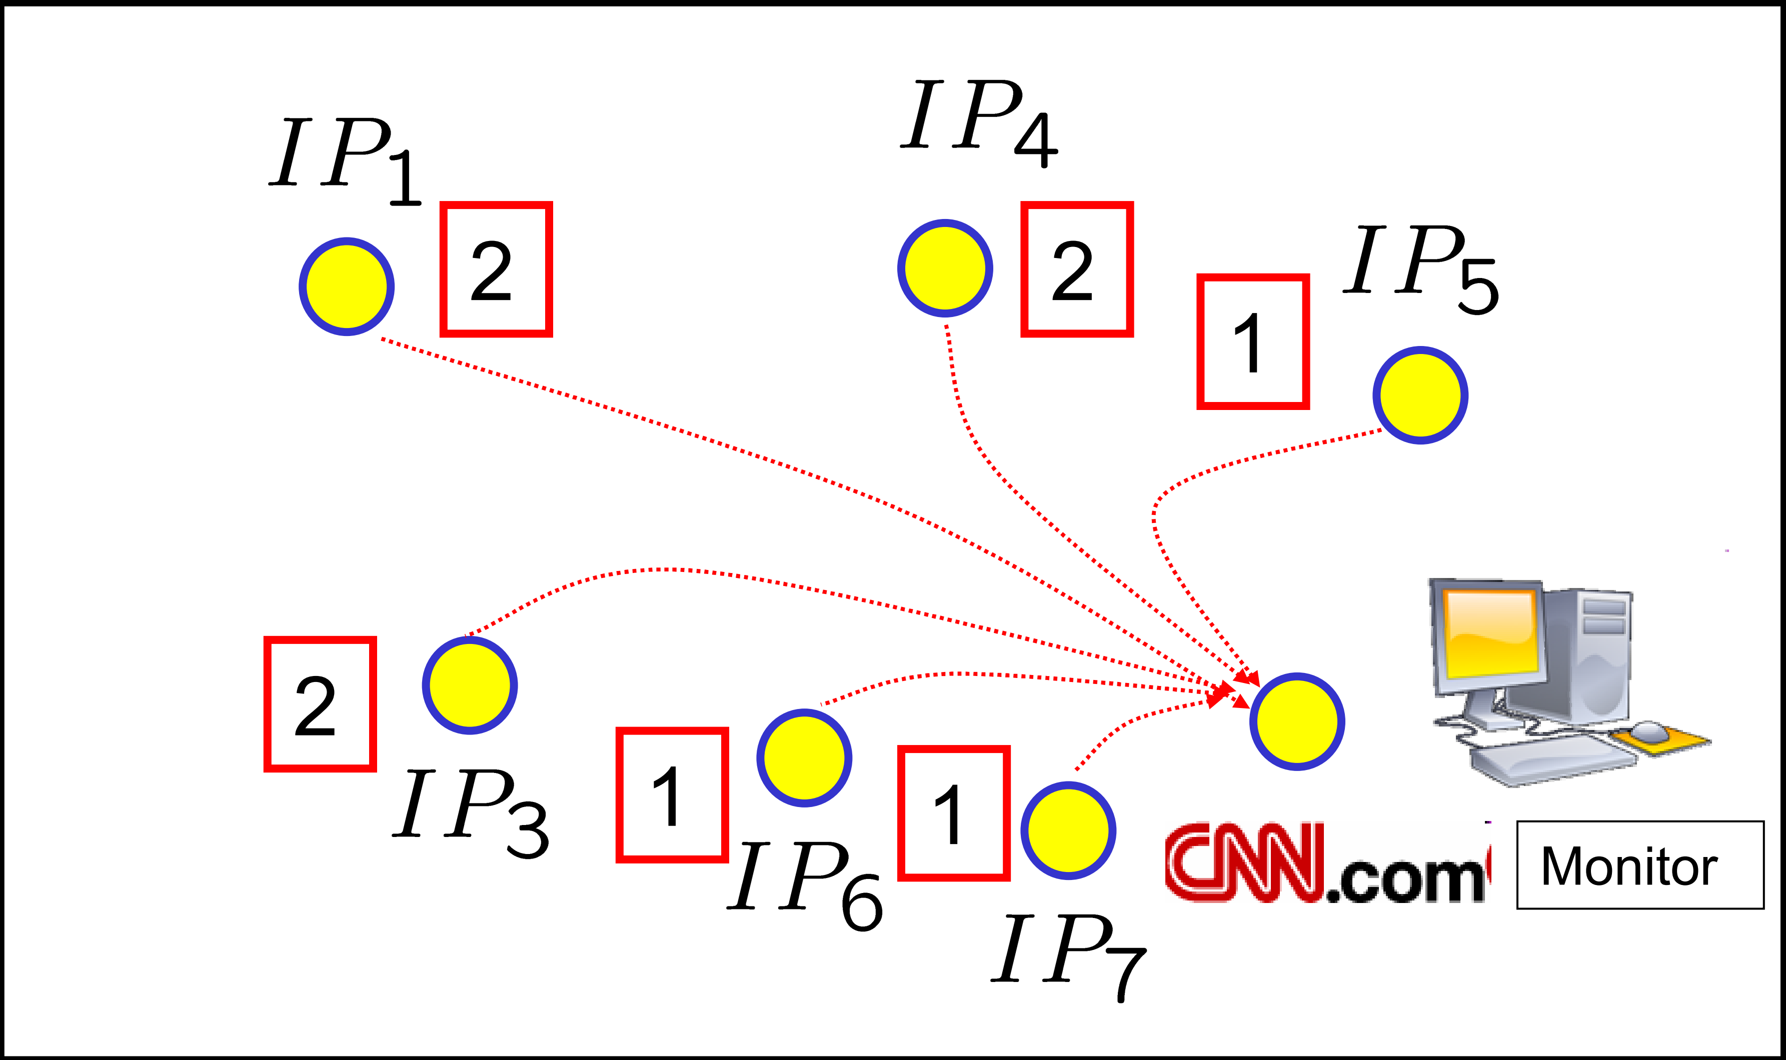

Incomplete Passive Measurements