Sheaf-based communication network invariants

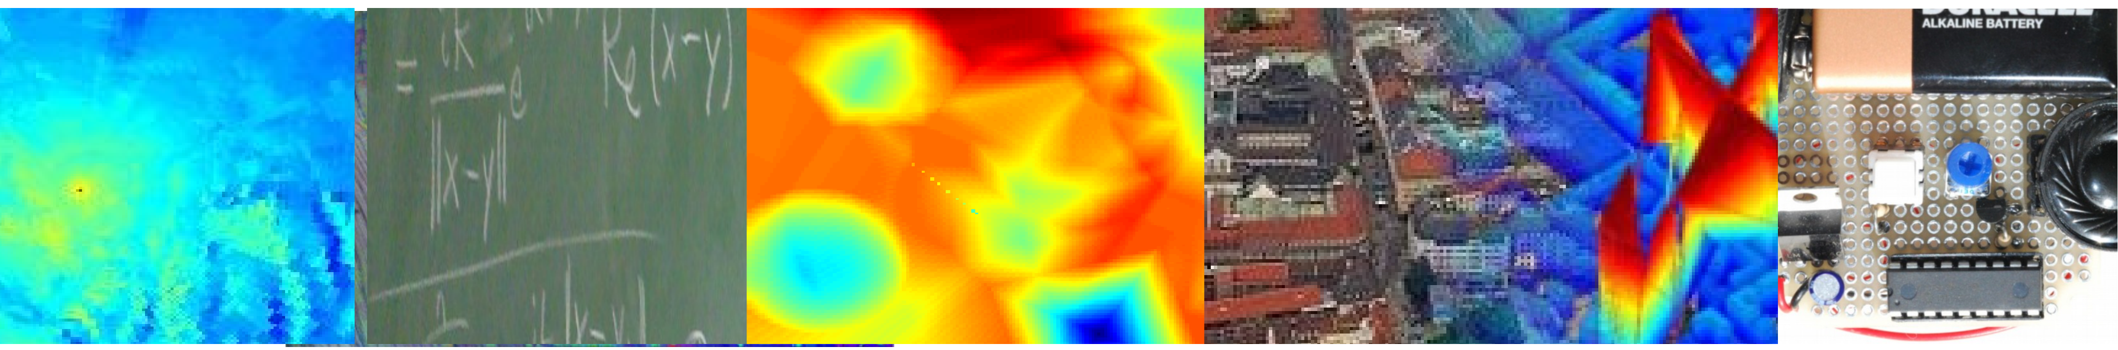

Michael Robinson

Acknowledgements

- Students (networks subgroup)
  - Eyerusalem Abebe
  - Danielle Beard
  - Morgan DeHart (past)
  - Samara Fantie
  - Dhanesh Krishnarao (past)
  - Jimmy Palladino
- Funding:
  - DARPA/STO (Wayne Phoel, current)
  - AFOSR (Bob Bonneau, past)

- Website reference

  http://www.drmichaelrobinson.net/

DARPA Tutorial on Sheaves in Data Analytics

August 25 and 26, 2015 (past)

American University, Washington, DC and online

Websites (includes links to videos, slides, and data):

- http://drmichaelrobinson.net/sheaftutorial/index.html
- http://www.american.edu/cas/darpasheaves/index.cfm

Partial funding:
DARPA SIMPLEX
Reza Ghanadan (DSO)

Problem statement

Assess the vulnerability of an *ad hoc* wireless network to congestion, jamming, or link failure

Challenges:

- The physical layer is extremely variable
- Network connectivity can be complicated
- Connectivity is hard to measure in practice
- Media access models can be subtle

However: Topological effects tend to dominate

An abstracted methodology

- Avoid specifying and committing to a high-fidelity physical or protocol model
- Instead, these models are abstracted into local connectivity information
  - This information is easy to measure
- Local connectivity leads to global inferences about network health
- Once the local connectivity is understood, more detail can be added as it is available

Topology ≠ Topology

- The topology of a wired network is rather concrete – it's a graph (maybe directed)
  - Vertices represent nodes
  - Edges represent **actual wires** between nodes
- The topology of a wireless network should be thought of more abstractly
  - Vertices represent nodes (still)
  - Higher dimensional faces represent collections of nodes that are co-visible to one another, in some sense
  - Dimensionality is a proxy for network density

Link complex

- Model of *ad hoc* wireless network in which all nodes are peers
- Two nodes *i* and *j* can communicate provided their signal strengths are large enough.

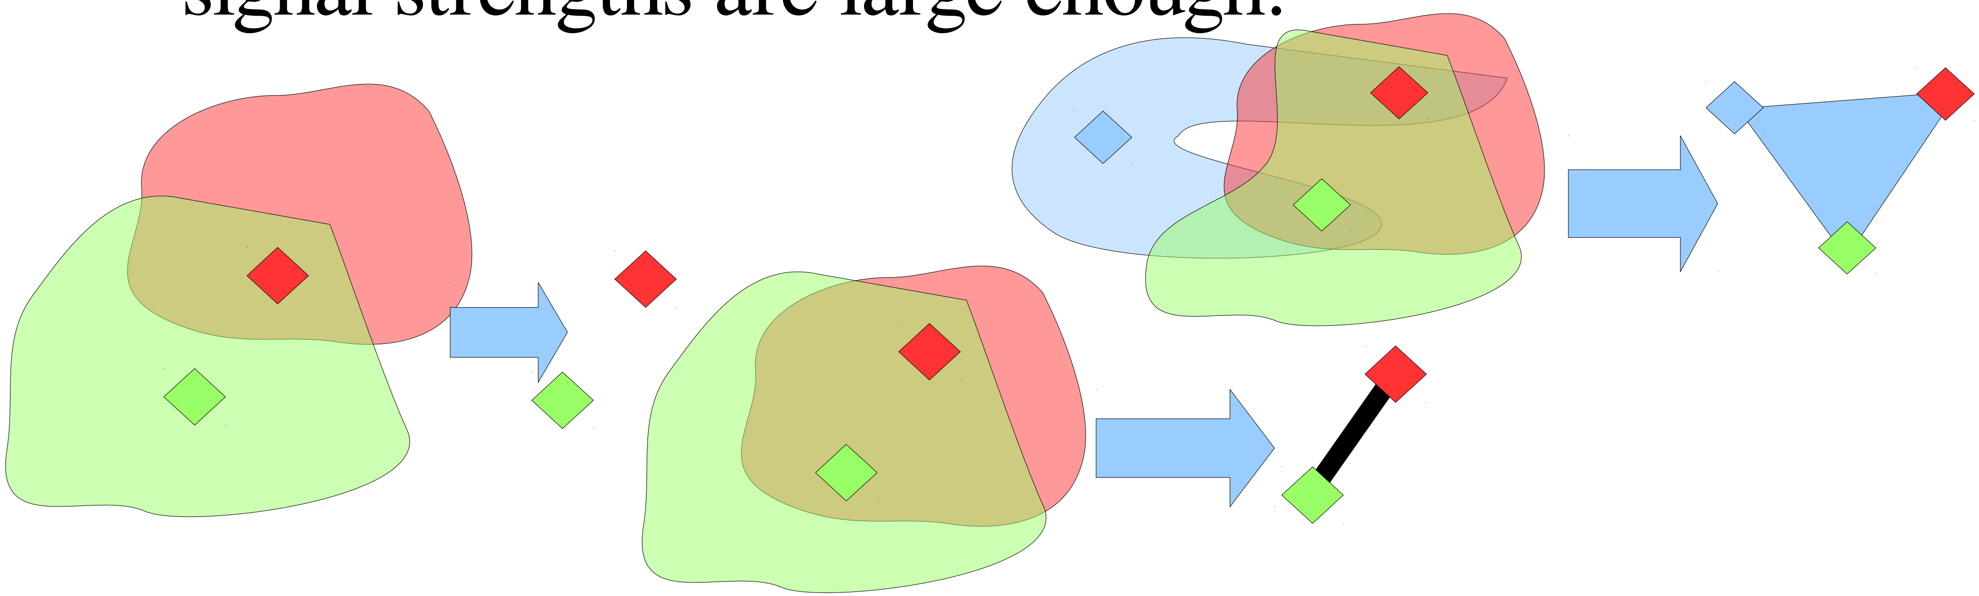

- The link complex is the flag complex all such edges

Network simulation

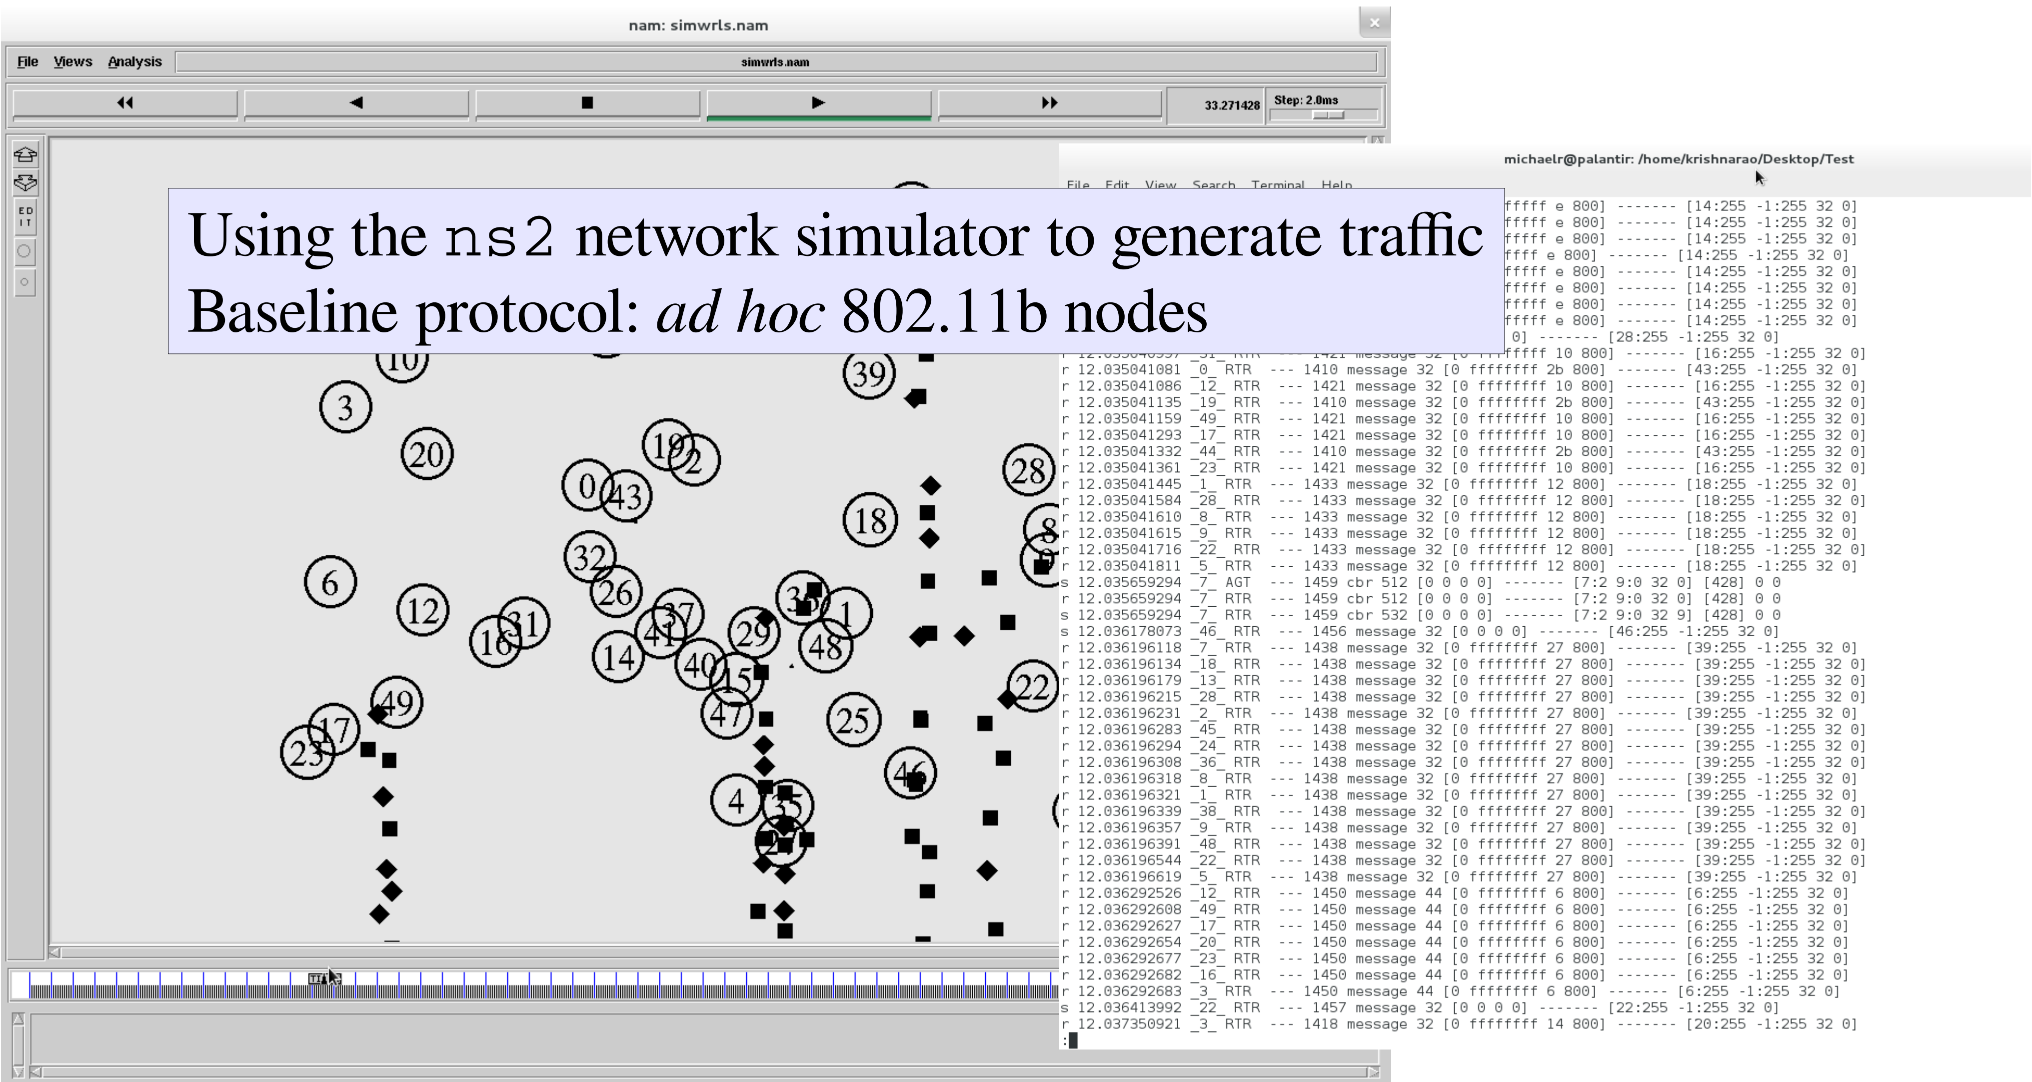

Using the `ns2` network simulator to generate traffic

Baseline protocol: *ad hoc* 802.11b nodes

with Eyerusalem Abebe, Dhanesh Krishnarao, Jimmy Palladino

Michael Robinson

Global topology and packet loss

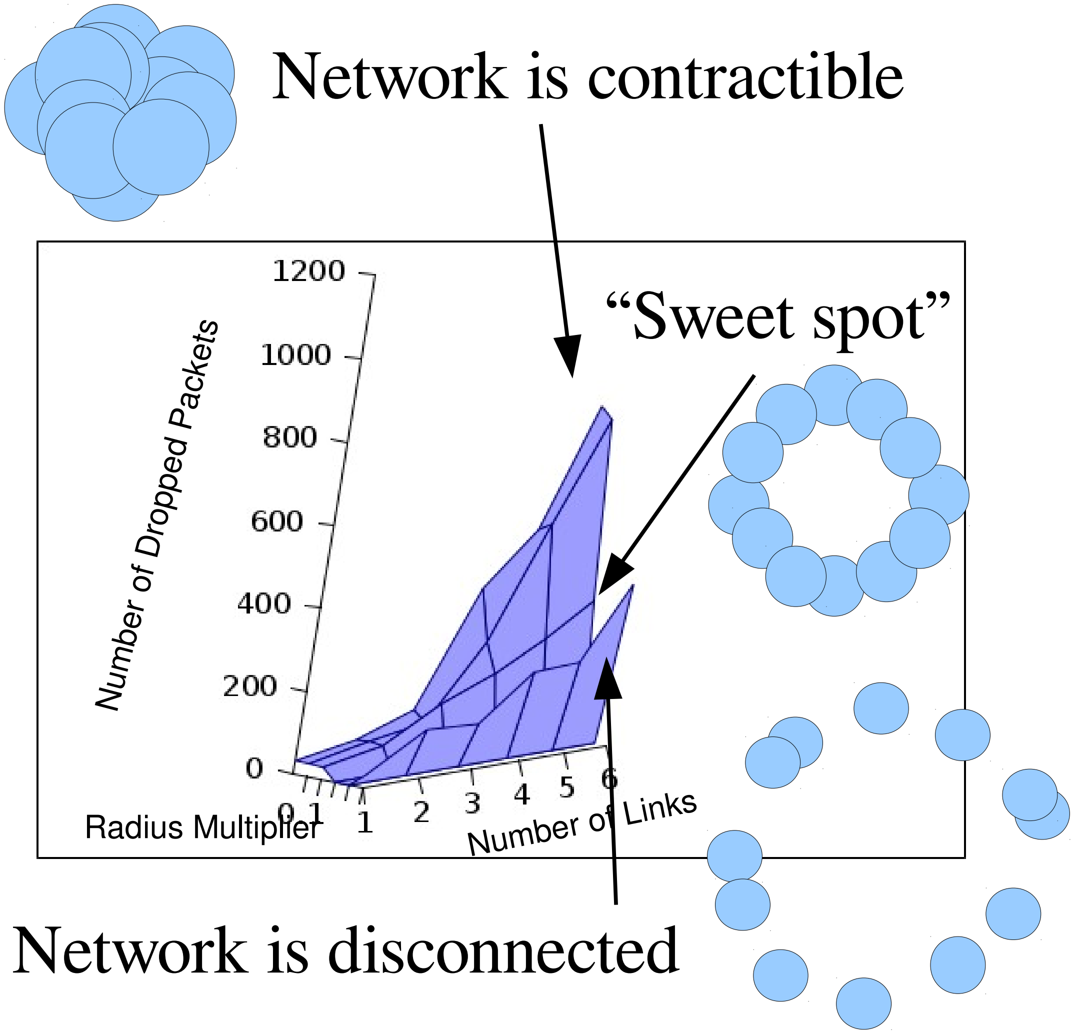

- A disconnected network is obviously bad
- A highly connected, contractible one is also bad:
  - Latency increases due to collisions
  - Many transmissions time out
- Connected, but not contractible network provides a good balance

Note: `ns2` doesn't simulate error rates due to low SNR, so even though we're moving nodes farther apart, the effect is purely topological!

Topological invariants

Persistent homological

- Global invariant
- Vulnerability to a specific source of interference
- Time independent

Local homological

- Local invariant: it's a sheaf!
- Vulnerability of the network to a link failure
- Time independent

Filtration from network disruption

- Given the set of nodes and their connection radius, the link complex is built
- Each simplex in the link complex is labeled with integers in the order at which it goes down

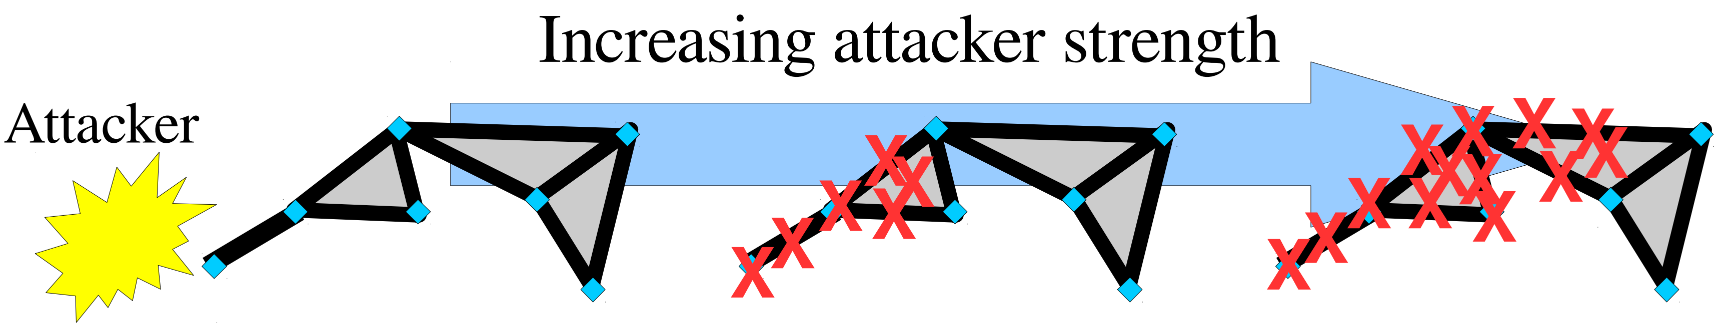

Labeled link complex (death order)

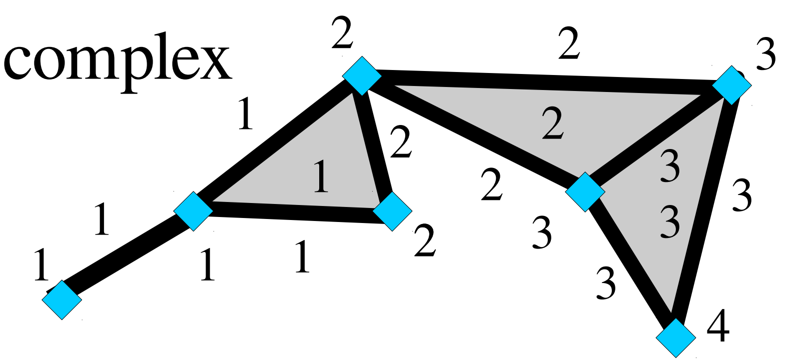

Reversed filtration

- To make `perseus` happy, we relabel the simplicial complex so that it lists birth order
- Not death order as we have initially
- Reorder by reversal of indices

`perseus` input

Labeled link complex (birth order)

Labeled link complex (death order)

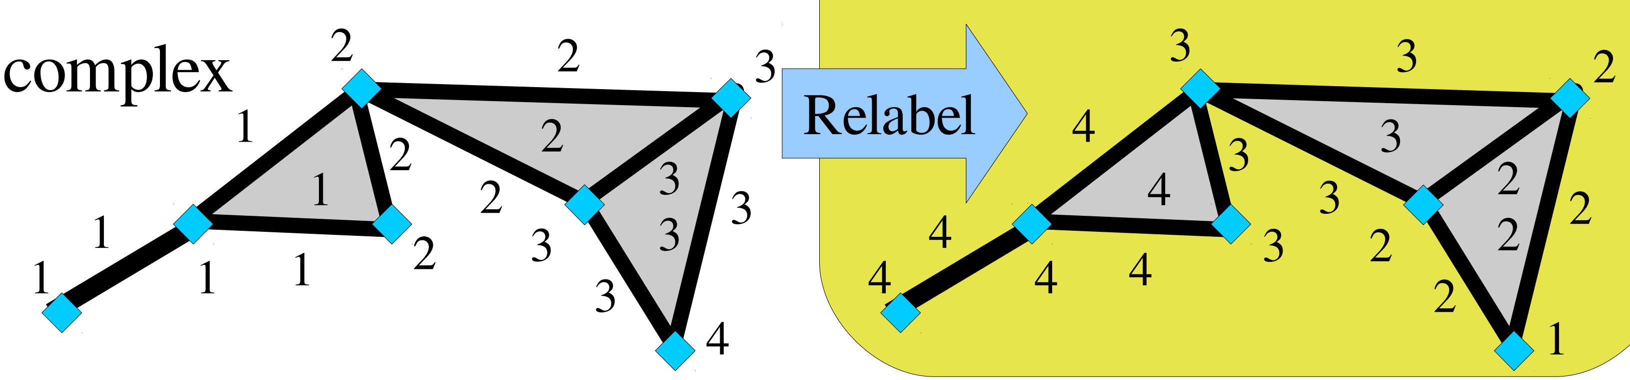

A random network

Notes:
Persistence axes are labeled against the “birth order” labels
Death index = -1 means that generator never dies

Disruption to network connectivity

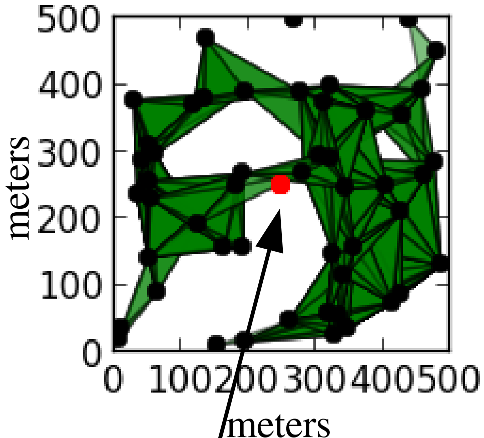

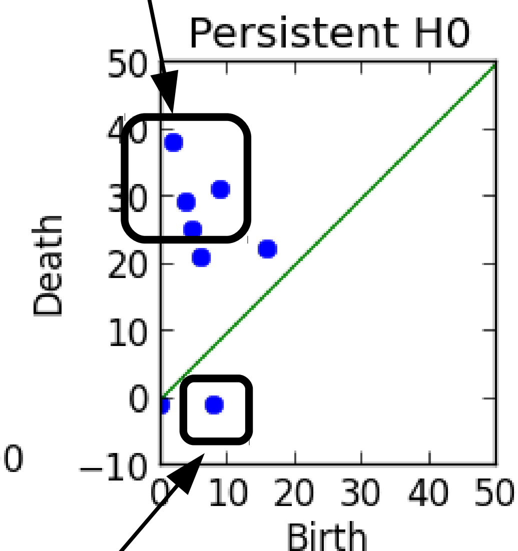

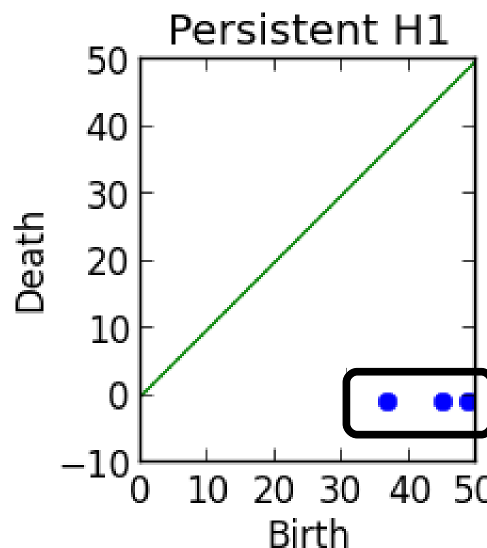

With no interference, the network has nontrivial loops

Persistent H_1 is protective!

Attacker

Even at its weakest, the attacker disrupts the network, since it's right inside a link

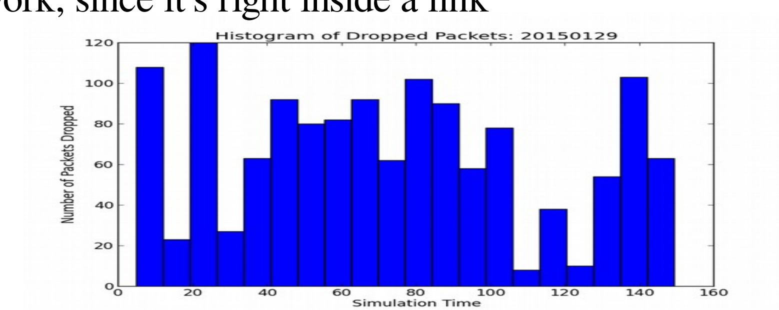

No attack

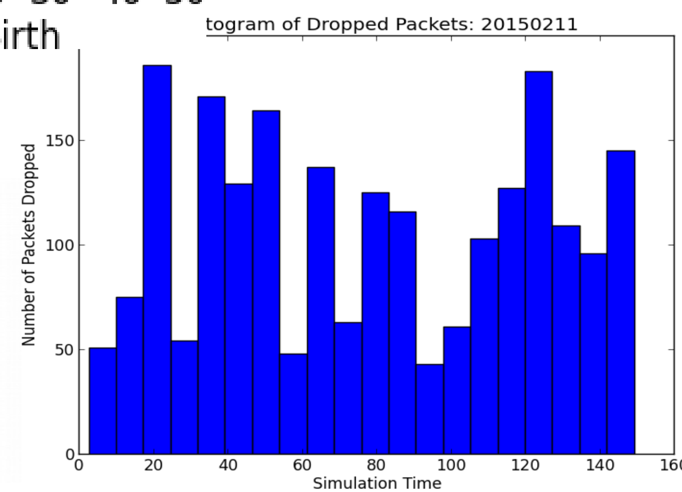

Attack

Michael Robinson

Tree network: aggressive attack

Peripheral attackers

Central attacker

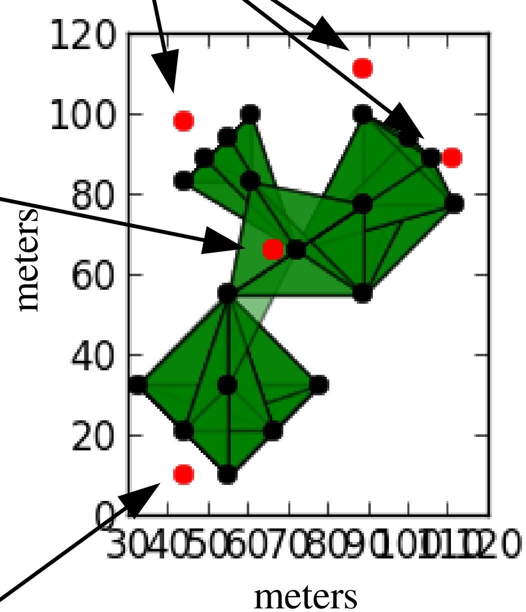

Peripheral attackers

Significant network connectivity disruption

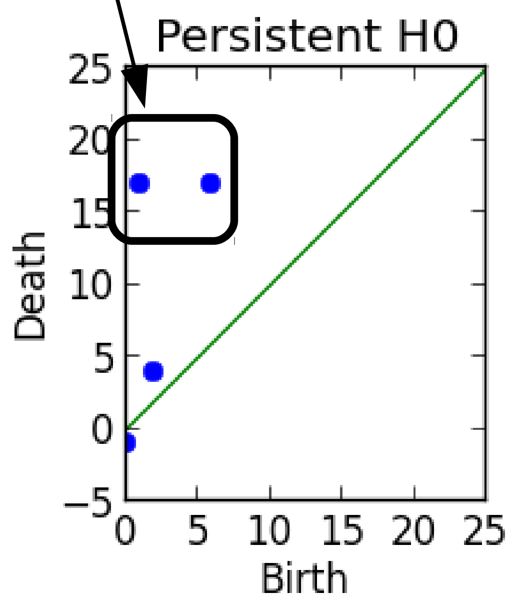

No nontrivial H_1 generators:

This is a serious vulnerability!

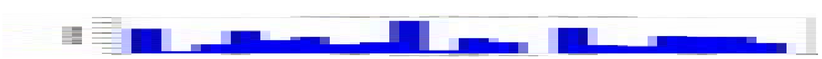

No attack

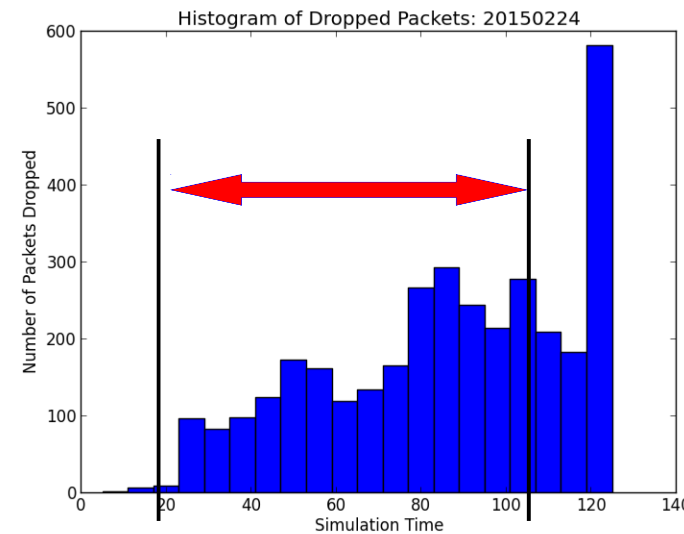

Attack

Tree network: less aggressive attack

Peripheral attackers

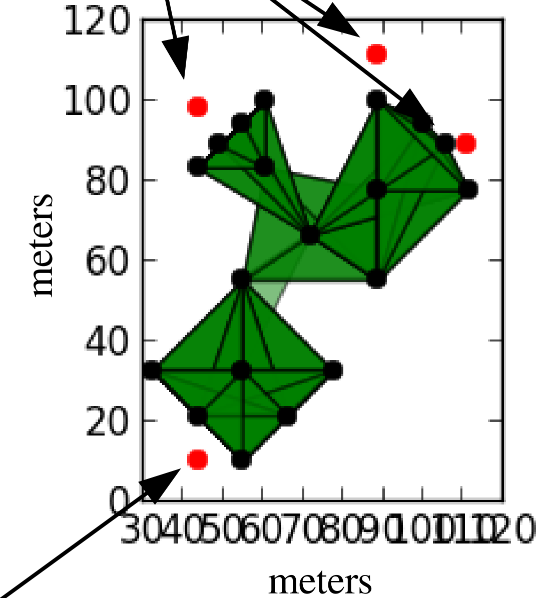

Peripheral attackers

Network connectivity disruption much reduced without central attacker

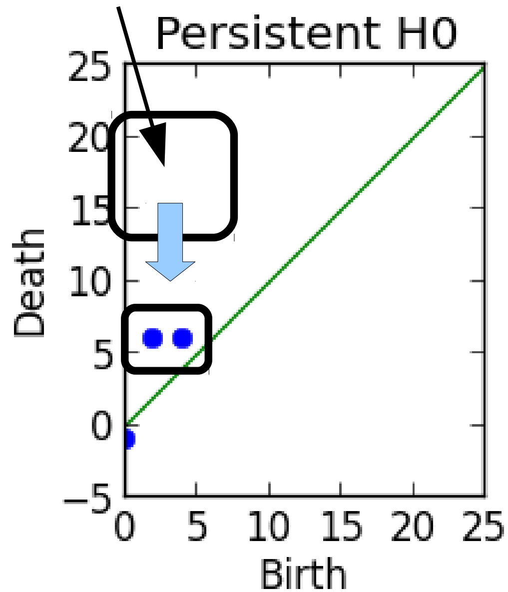

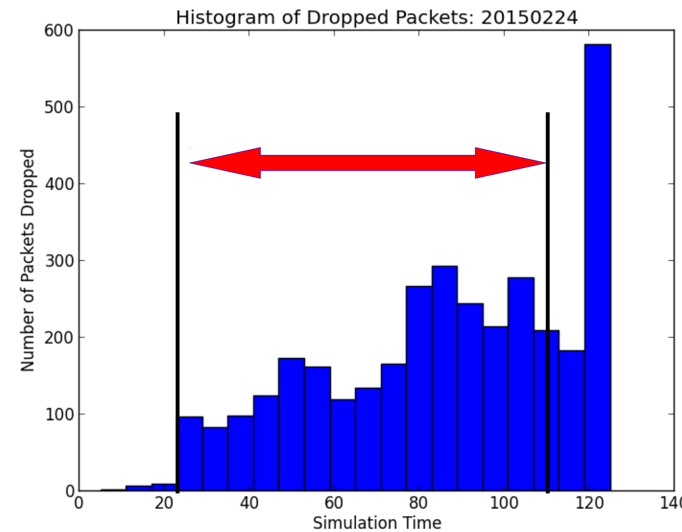

Network performance improved as well

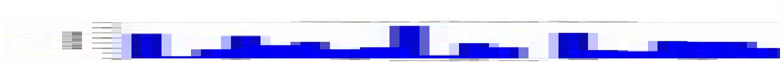

No attack

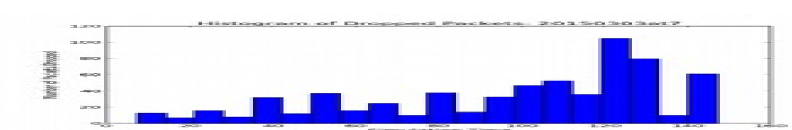

Attack

Sensitivity to the traffic pattern

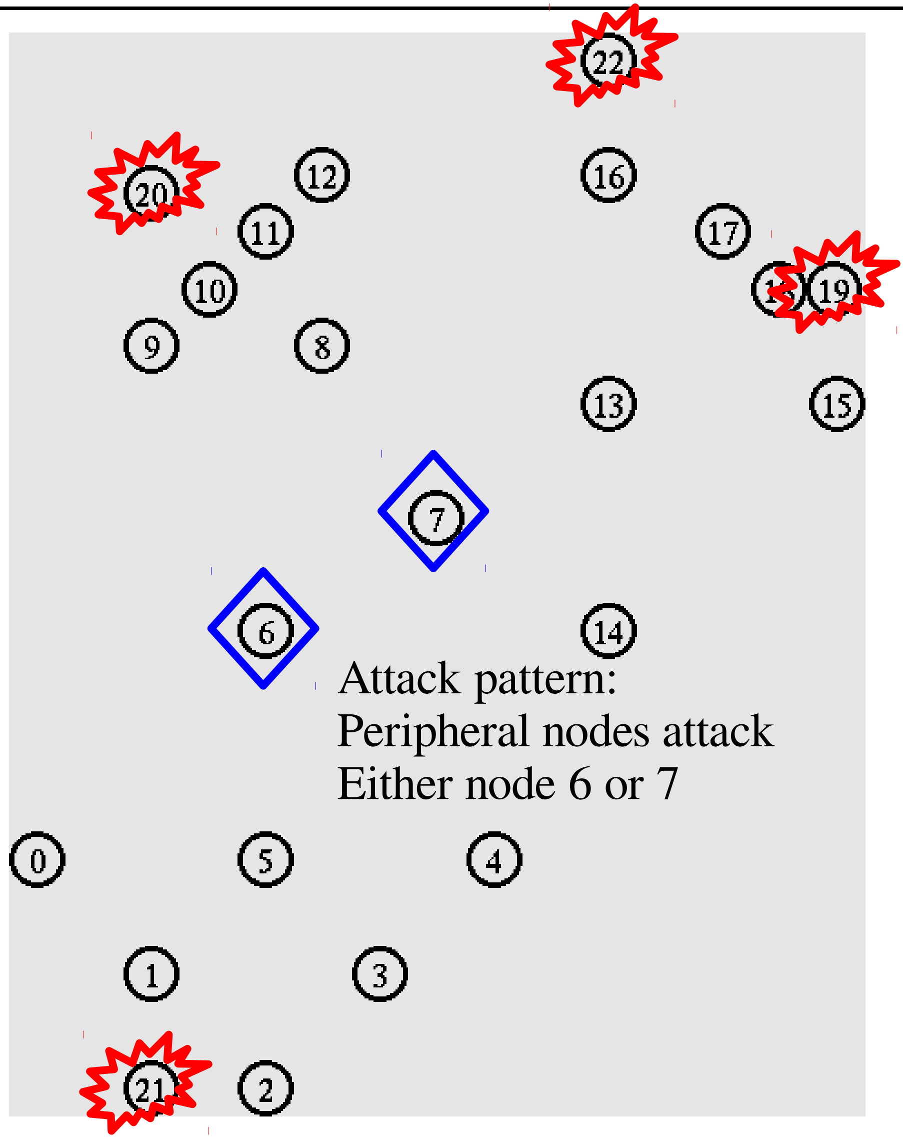

Attack pattern:
Peripheral nodes attack
Either node 6 or 7

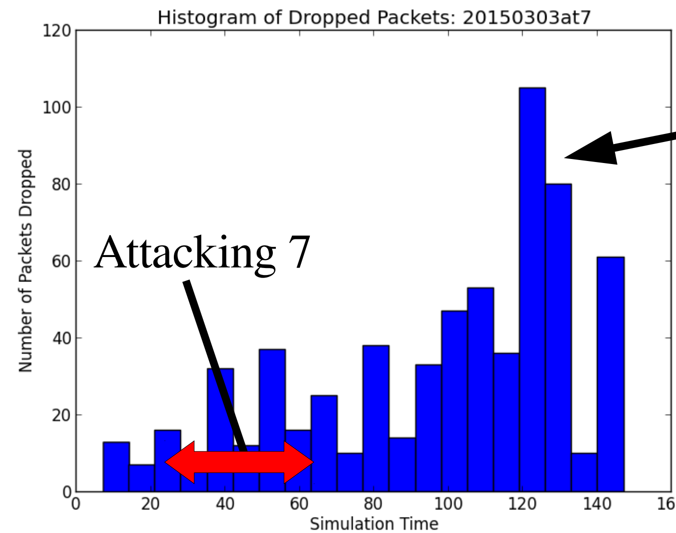

Attacking 7

Lingering packet loss after attack ends

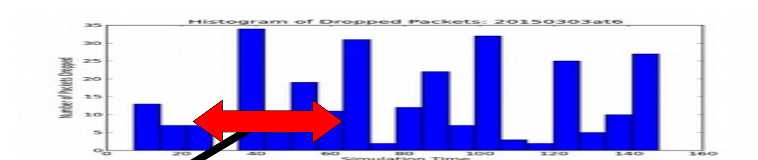

Attacking 6

Observe: Attacking 7 is much worse for network!
Total traffic is the same for both cases

Local homology

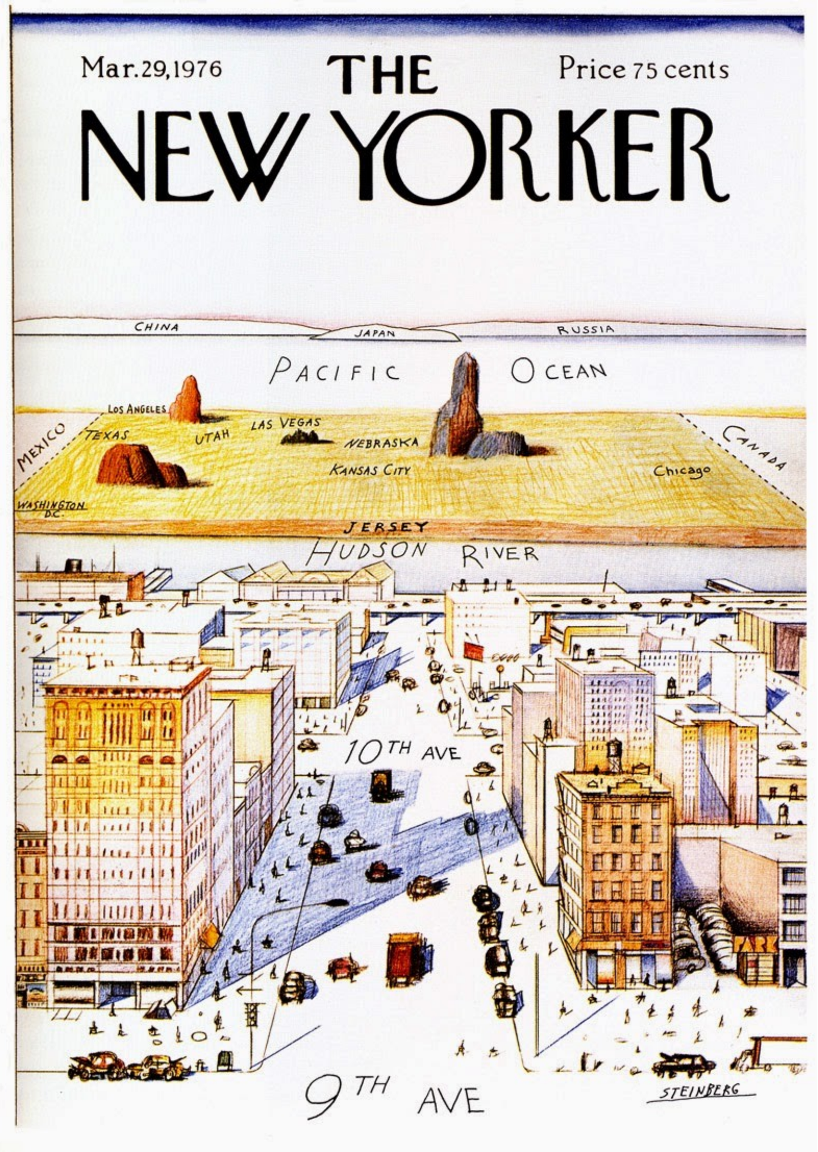

Michael Robinson

Relative homology → algebraic locality

There's a chain complex that computes homology of a simplicial complex X neglecting a particular closed subcomplex A

- Use $C_k(X,A) = C_k(X) / C_k(A)$
- Same boundary maps, just descend to the quotient

Theorem: $H_k(X, A) \cong H_k(X / A)$ for $k > 0$

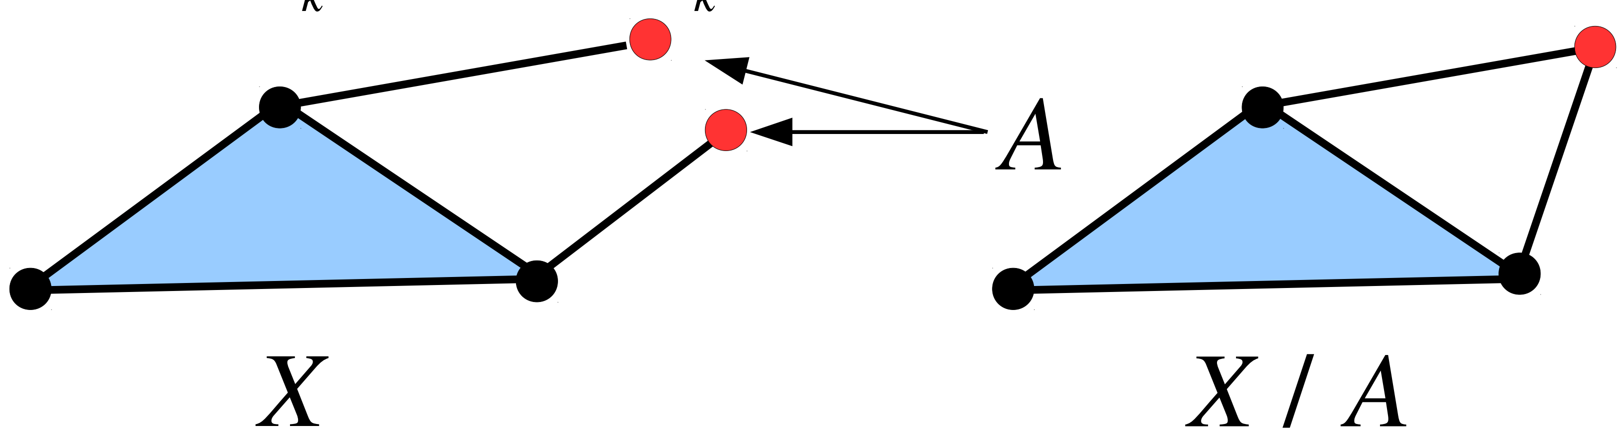

Local homological invariant

- We "delete" an open neighborhood of a simplex of interest

$$H_k(X, X \setminus \text{star } a)$$

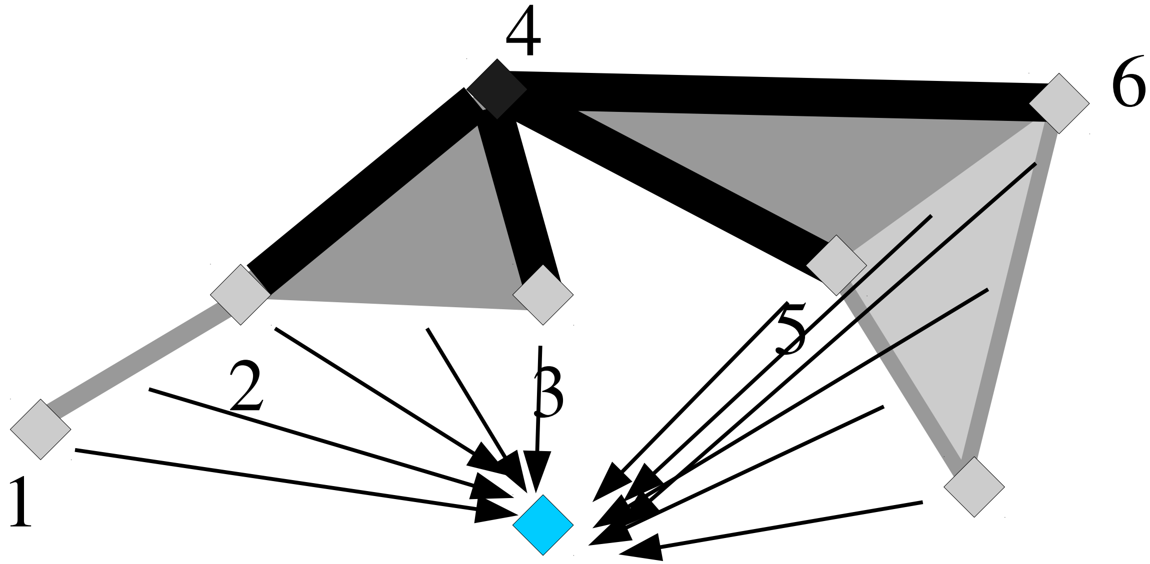

Local homology and excision

Let A = cl star a and $B = X \setminus$ star a

- Both A and B are closed subcomplexes
- $A \cup B = X$

Thus $H_k(X, X \setminus \text{star } a) = H_k(X,B) \cong H_k(A, A \cap B)$

$\cong H_k(\text{cl star } a, \partial \text{ star } a)$

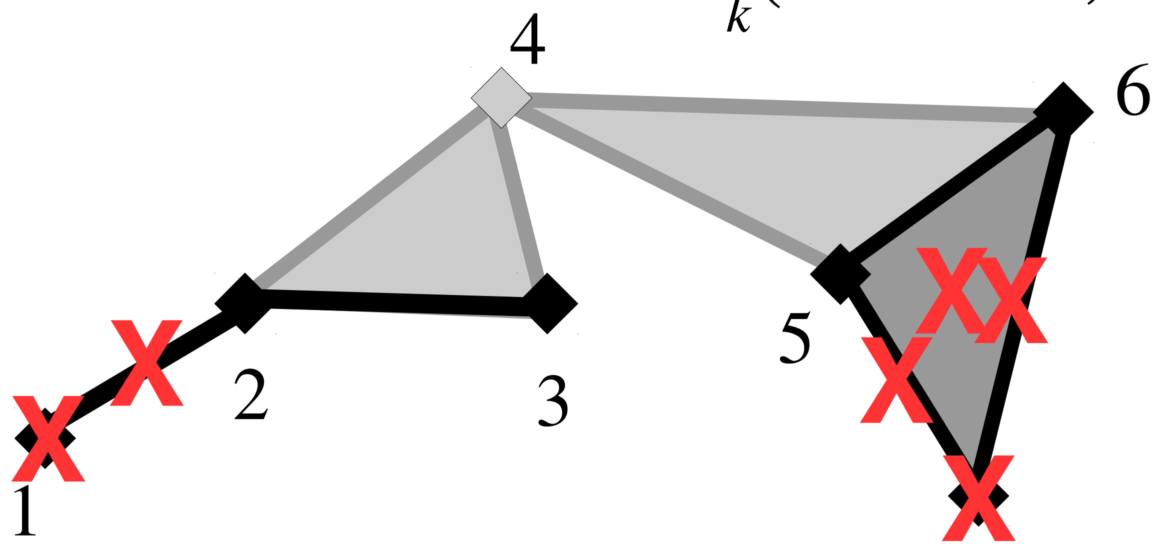

Local homology and excision

Let A = cl star a and $B = X \setminus$ star a

- Both A and B are closed subcomplexes
- $A \cup B = X$

Thus $H_k(X, X \setminus \text{star } a) = H_k(X,B) \cong H_k(A, A \cap B)$

$\cong H_k(\text{cl star } a, \partial \text{ star } a)$

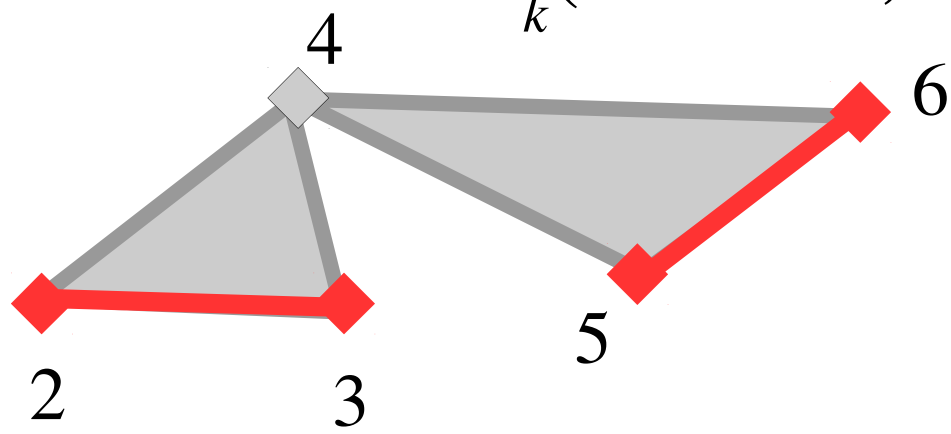

Local homology is a sheaf

Suppose that a is a k-face of a (k+1)-simplex b

Then,

$$a \subset b$$

$$\text{star } a \subset \text{star } b$$

$$X \setminus \text{star } a \supset X \setminus \text{star } b$$

Which induces a linear map (depending on a and b)

$$H_k(X, X \setminus \text{star } a) \to H_k(X, X \setminus \text{star } b).$$

And the gluing axioms hold, too.

Local homology is a sheaf

Here is a simplicial complex

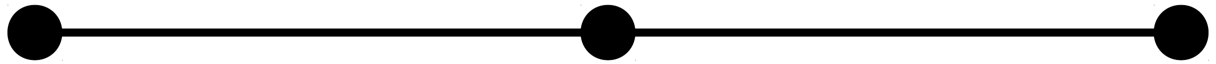

Local homology is a sheaf

Here are the local pair complexes

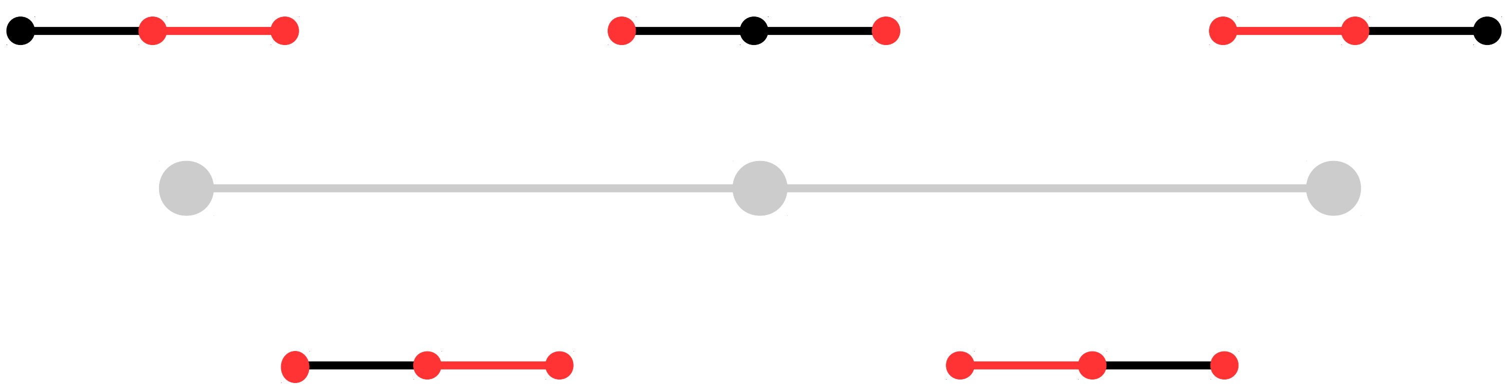

Key: each diagram shows a topological pair (X,A)

Local homology is a sheaf

Here are the local pair maps

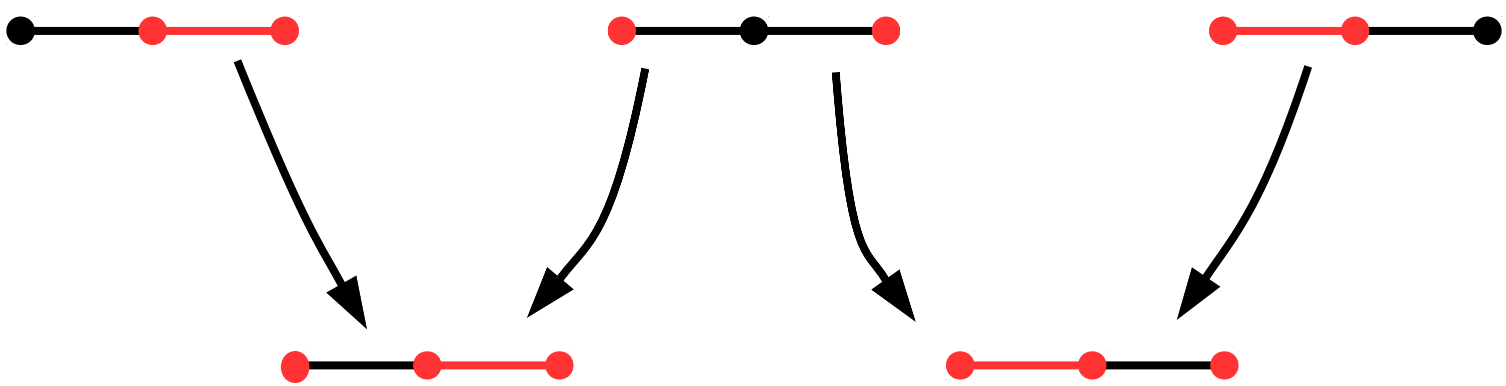

Key: each diagram shows a topological pair (X,A)

Local homology is a sheaf

Local H_1: (all others vanish)

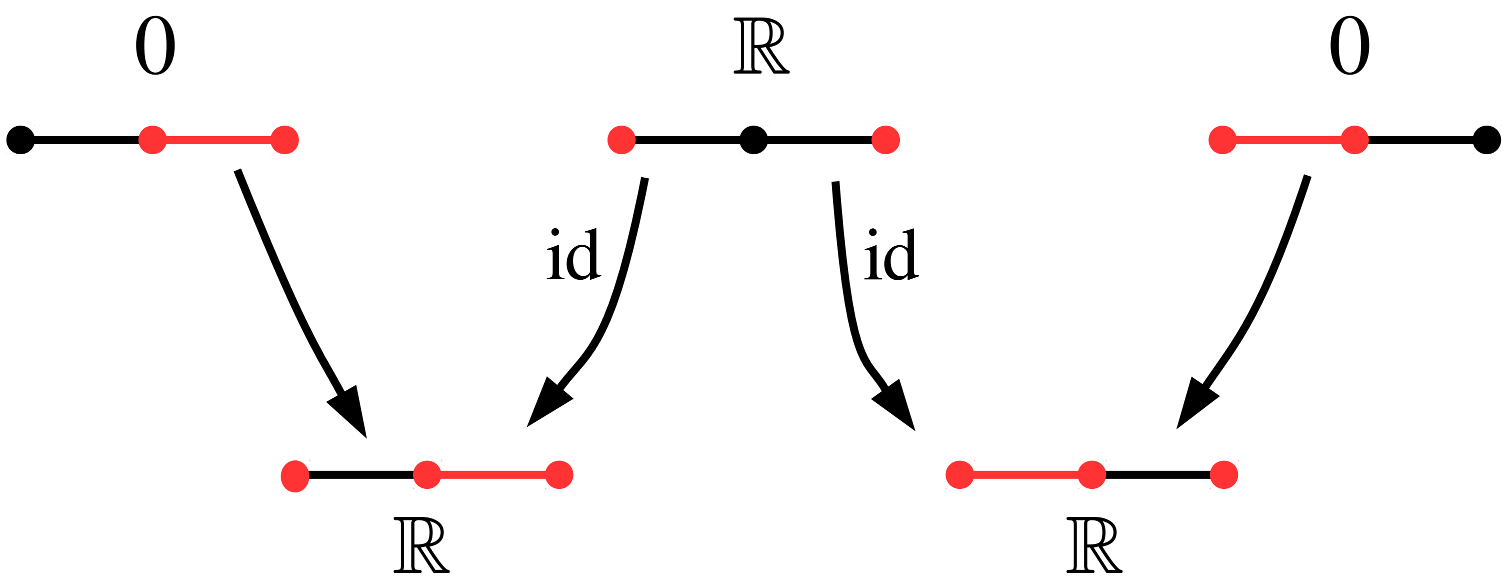

Key: each space is $H_1(X,A)$ with real coefficients

The arrows are the induced maps on homology

Local homology is a sheaf

Here is another simplicial complex...

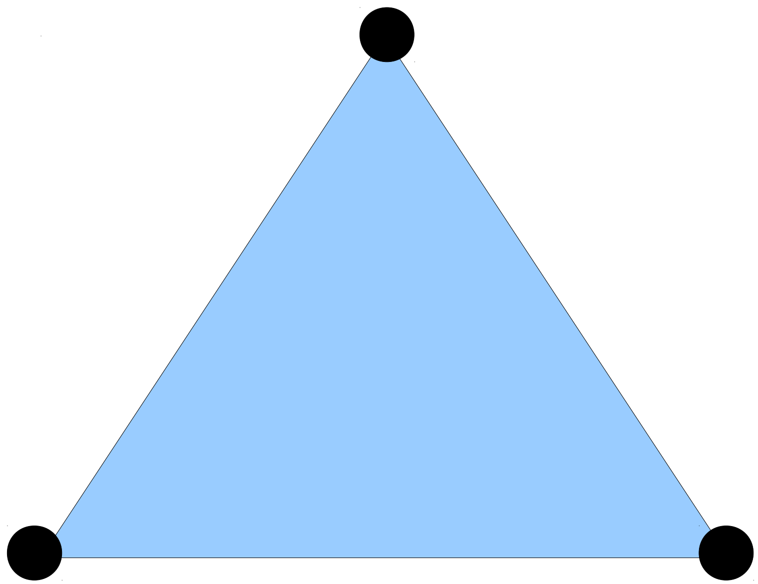

Local homology is a sheaf

Local pair complexes

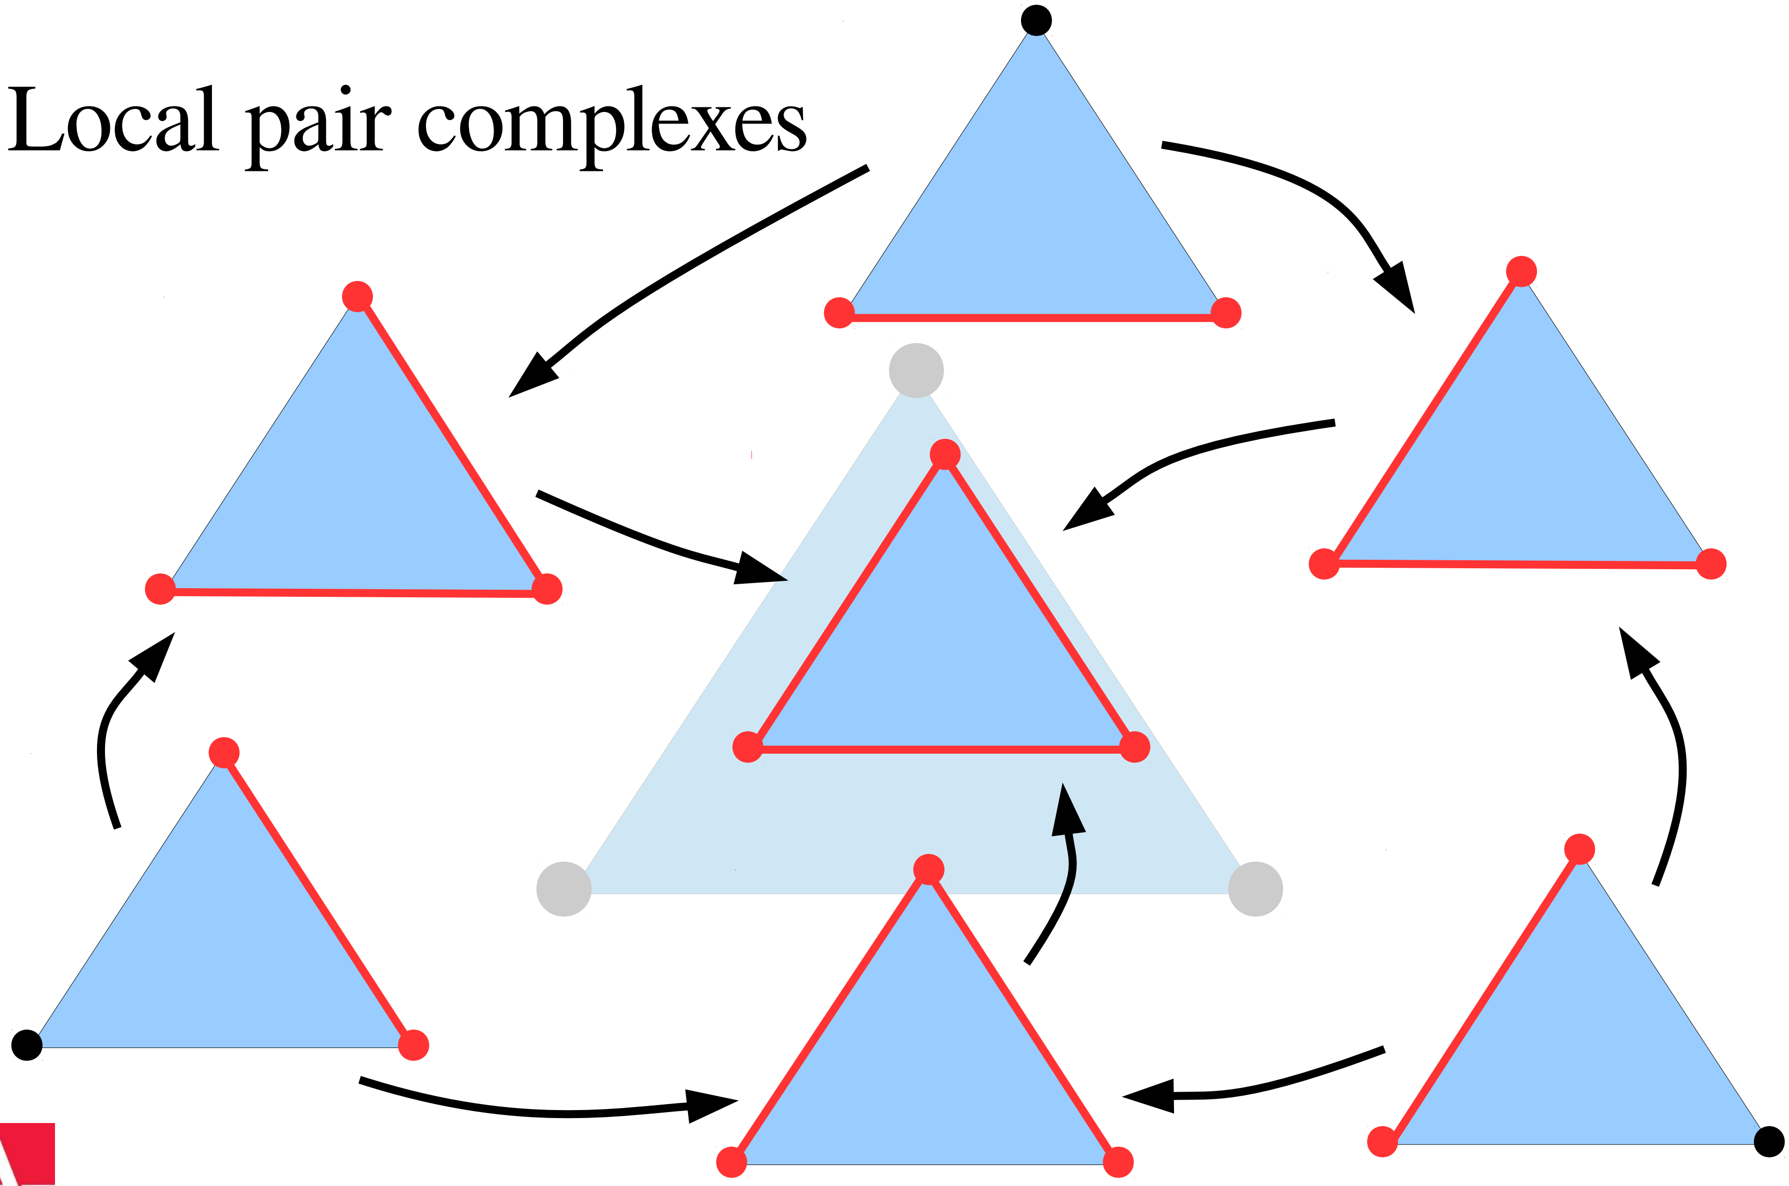

Local homology is a sheaf

Local H_1

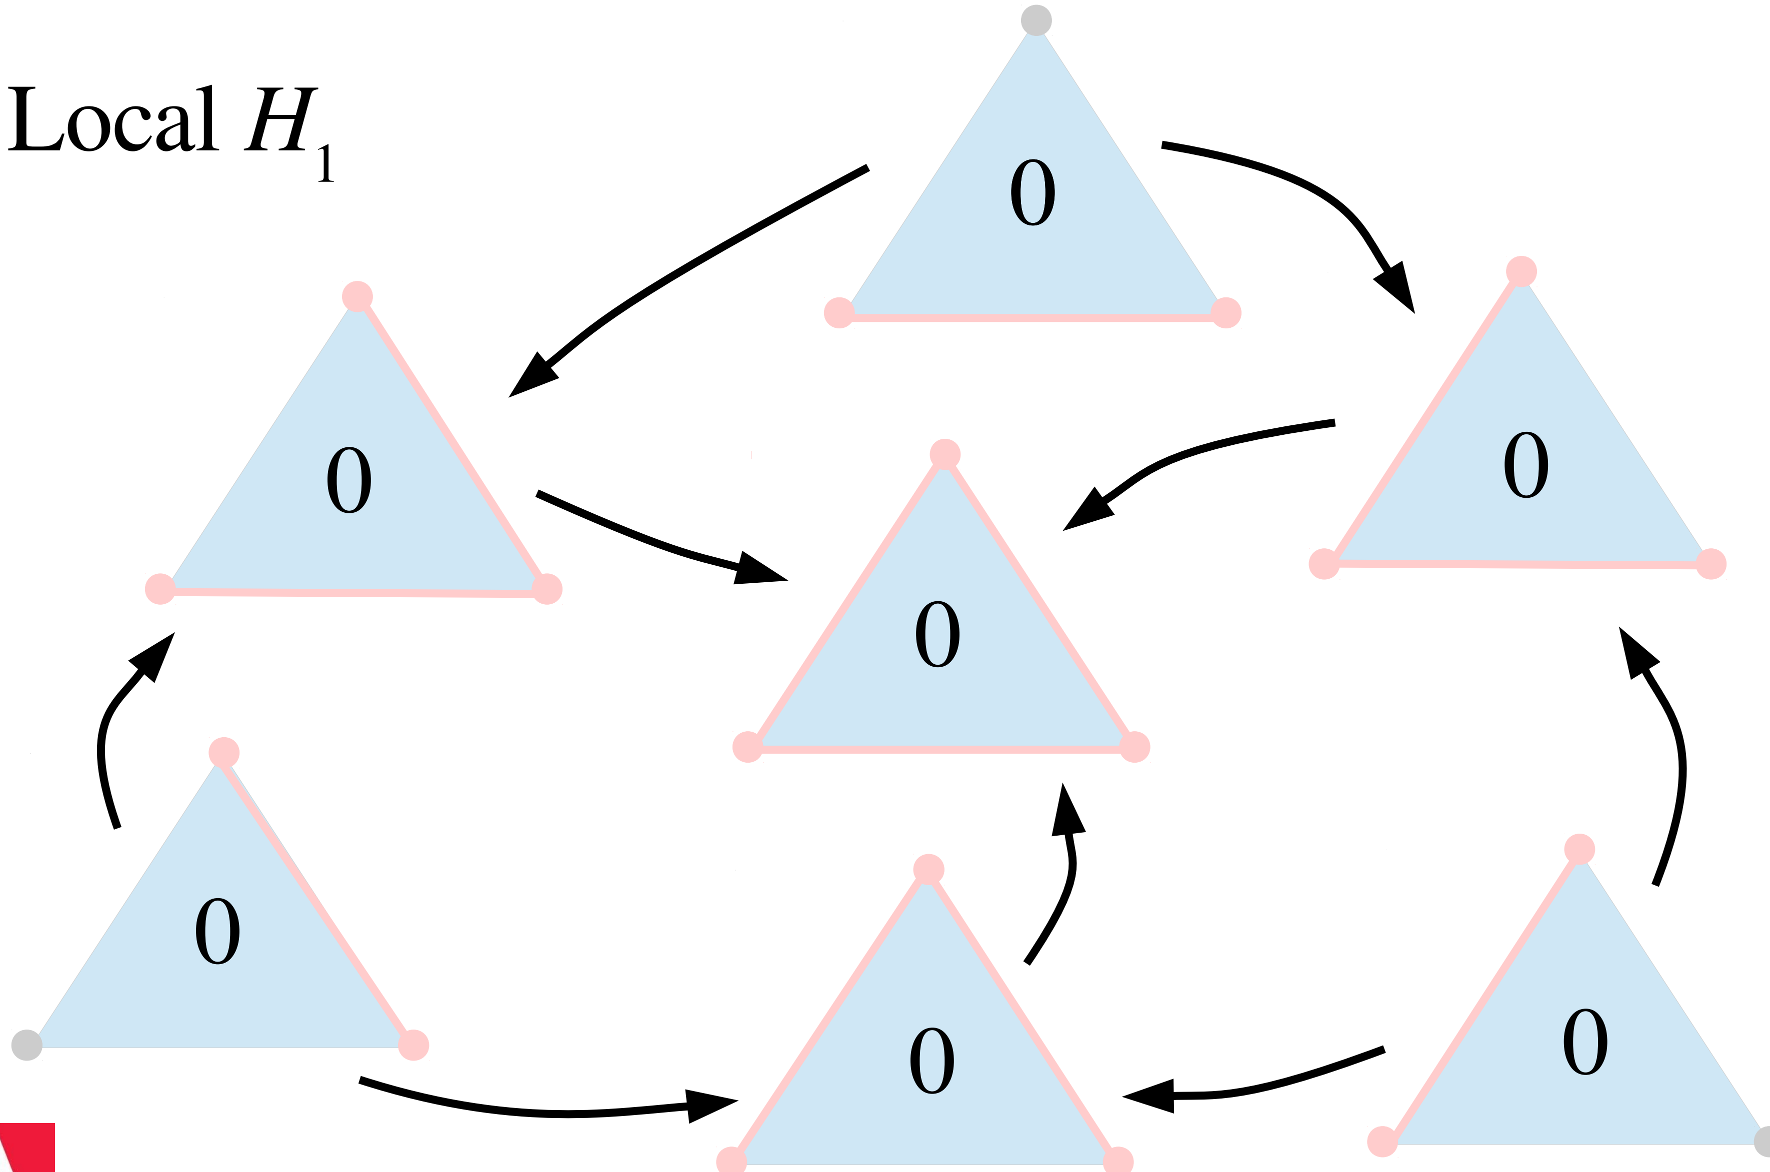

Local homology is a sheaf

Local H_2

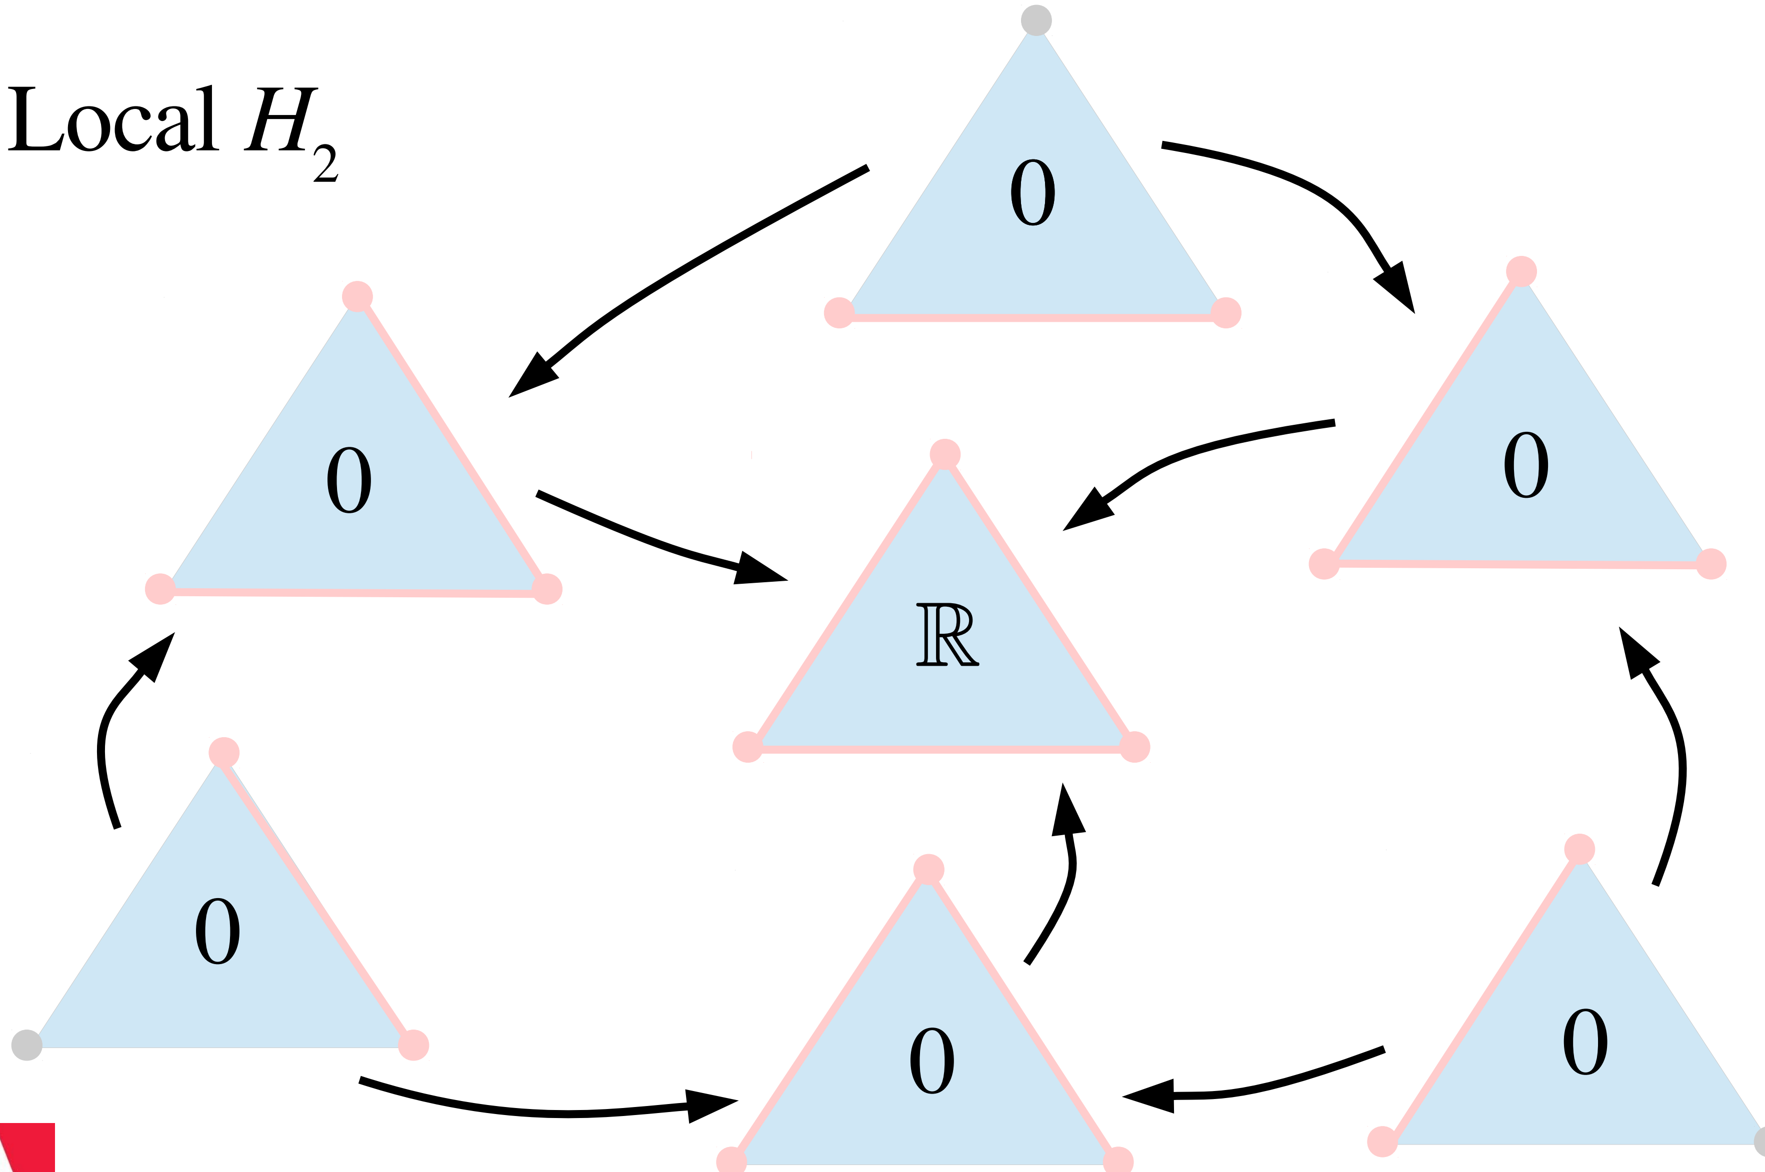

Local homology and graph degree

If X is a graph (a 1-dimensional simplicial complex), then for any vertex v,

$$\dim H_1(X, X \setminus \text{star } v) = \deg v - 1$$

But while degree is not very informative if X has higher dimensional simplices, $H_1(X, X \setminus \text{star } v)$ is still a topological invariant

Local homology and graph degree

- Local homology is **not** graph degree for simplicial complexes, though

Magenta = 0
Blue = 1
Cyan = 2
Green = 3
Yellow = 4
Red = 5+

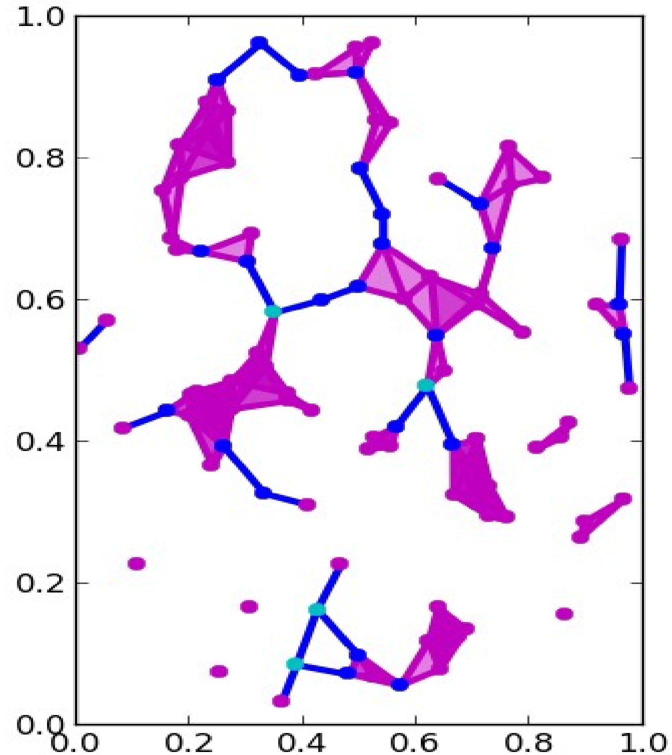

Local H_1

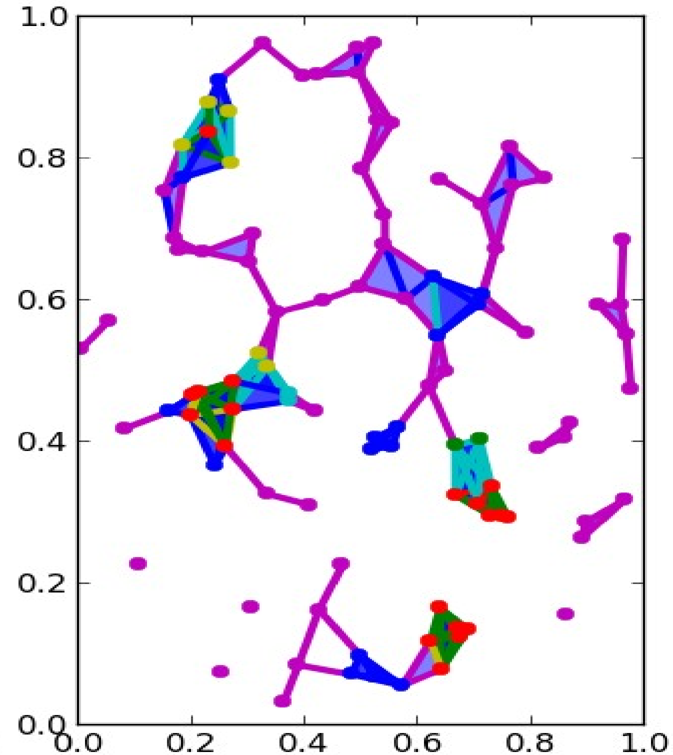

Local H_2

Local homology and vulnerability

More vulnerable link around node 7

Also a high dimensional simplex; many nodes use this link

Local H_1 dimension:
0 = blue
1 = green
2 = red

Attack pattern:
Peripheral nodes attack
Either node 6 or 7

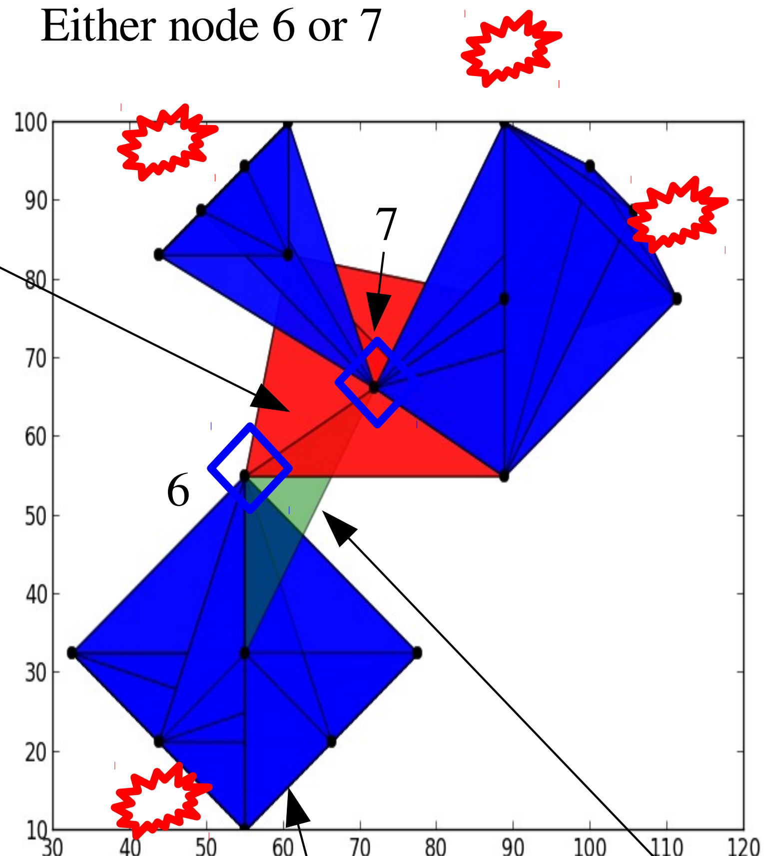

End clusters not very vulnerable

Somewhat less vulnerable link around node 6

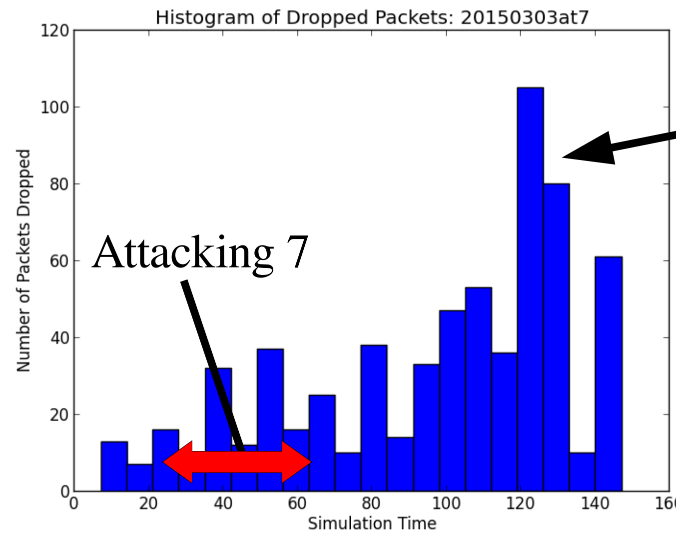

Attacking 7

Lingering packet loss after attack ends

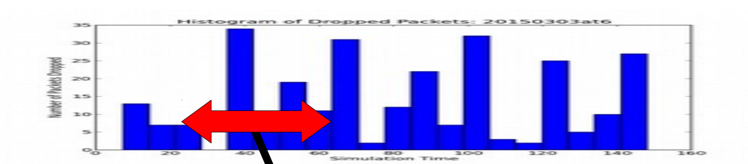

Attacking 6

Michael Robinson

Forwarded packet distributions

Forwarded packet distributions

- The number of packets forwarded by a node appears to depend on its position in the network

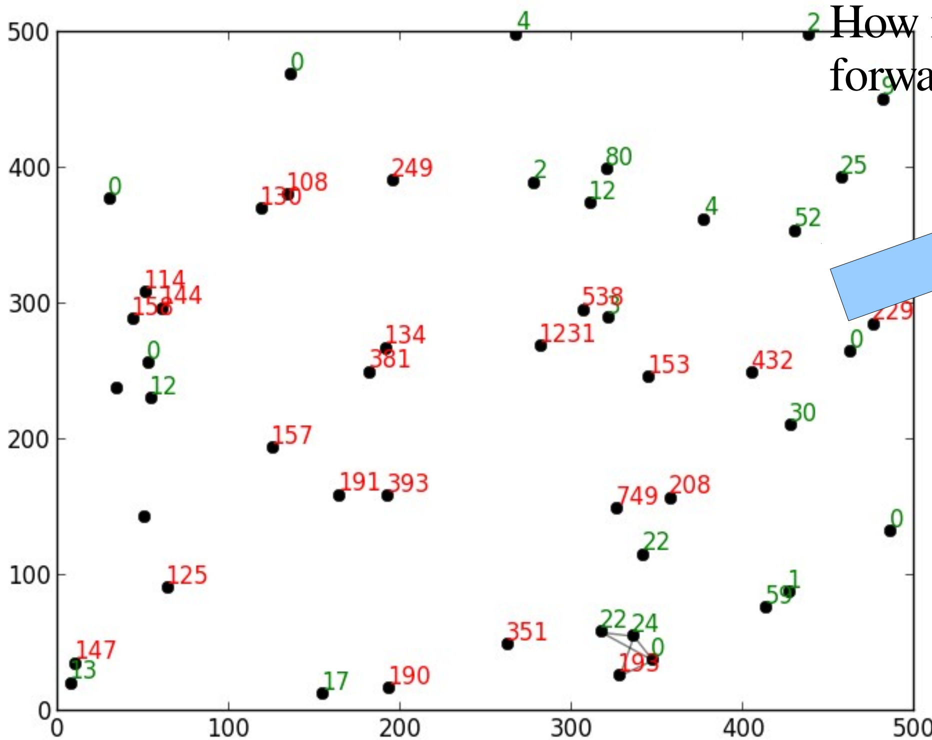

How many packets forwarded by each node

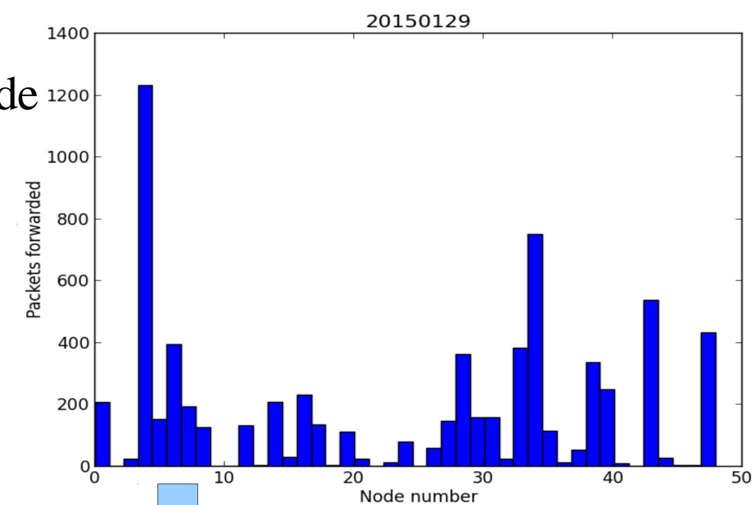

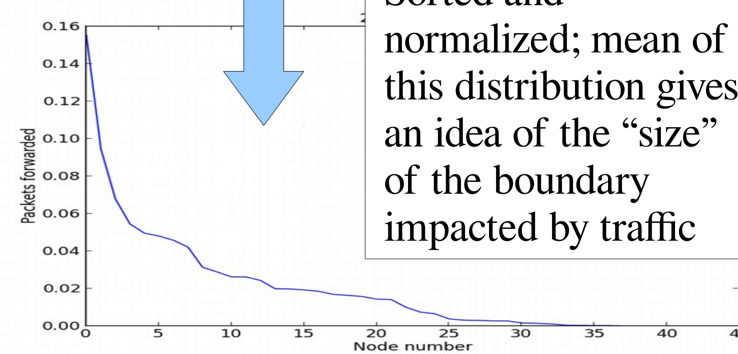

Sorted and normalized; mean of this distribution gives an idea of the "size" of the boundary impacted by traffic

Numbers indicate number of packets forwarded: Red if over 100, green if less than 100

Forwarded packets vs. network "pinch points"

- Local homology compared to forwarded packet distribution

Magenta = 0
Blue = 1
Cyan = 2

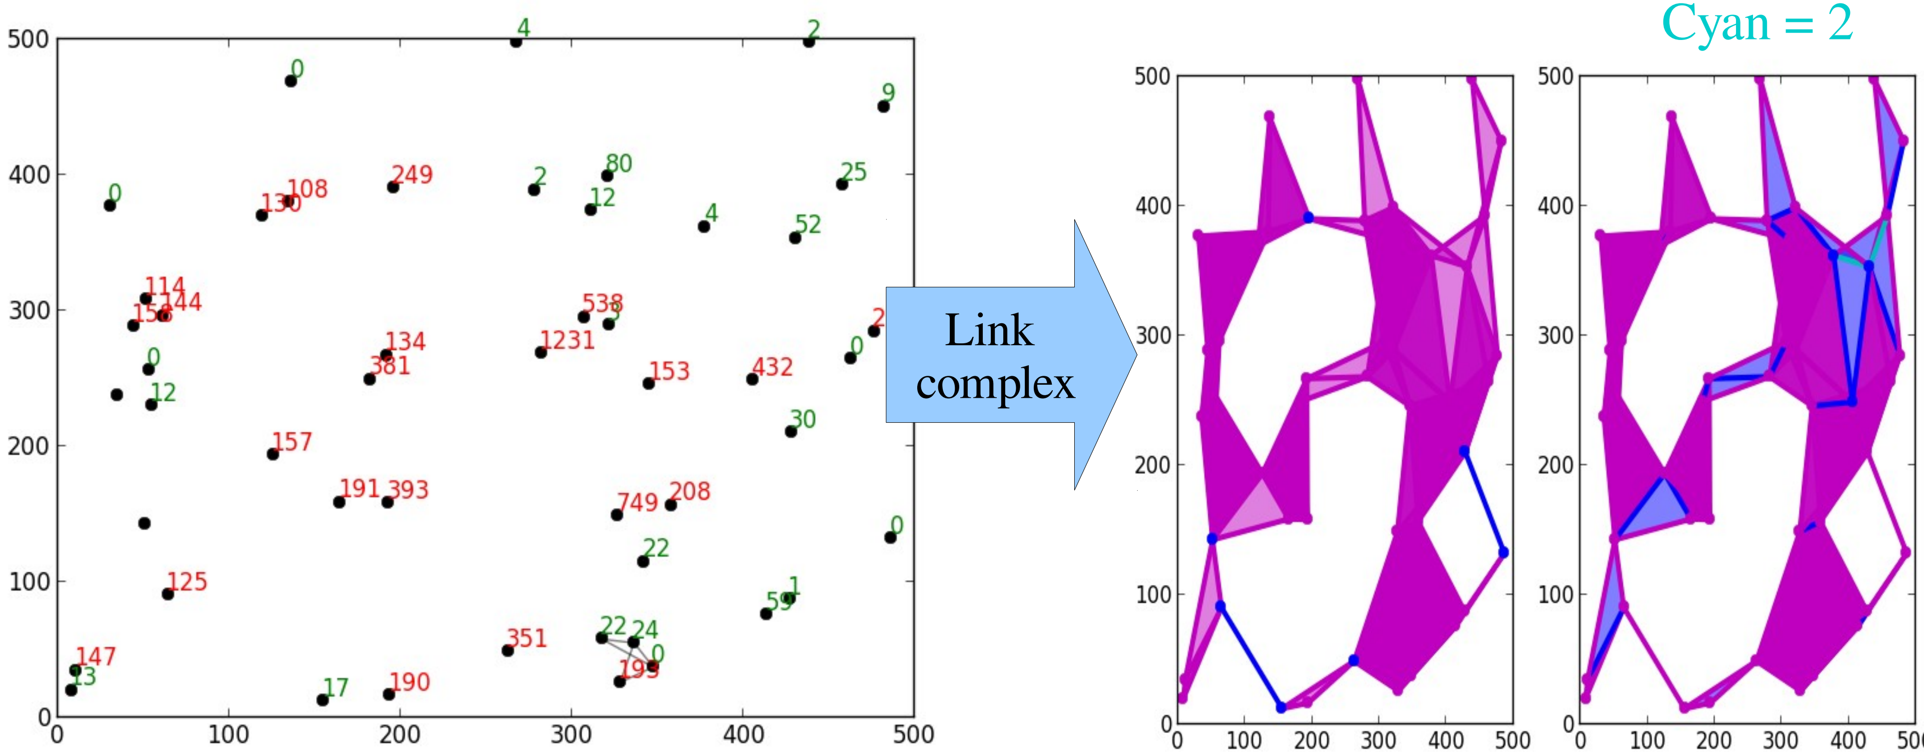

Local H_1 Local H_2

Numbers indicate number of packets forwarded:
Red if over 100, green if less than 100

Forwarded packets vs. network "pinch points"

High local H_1 dimension is a topological pinch point

All nodes that forward many packets are at pinch points

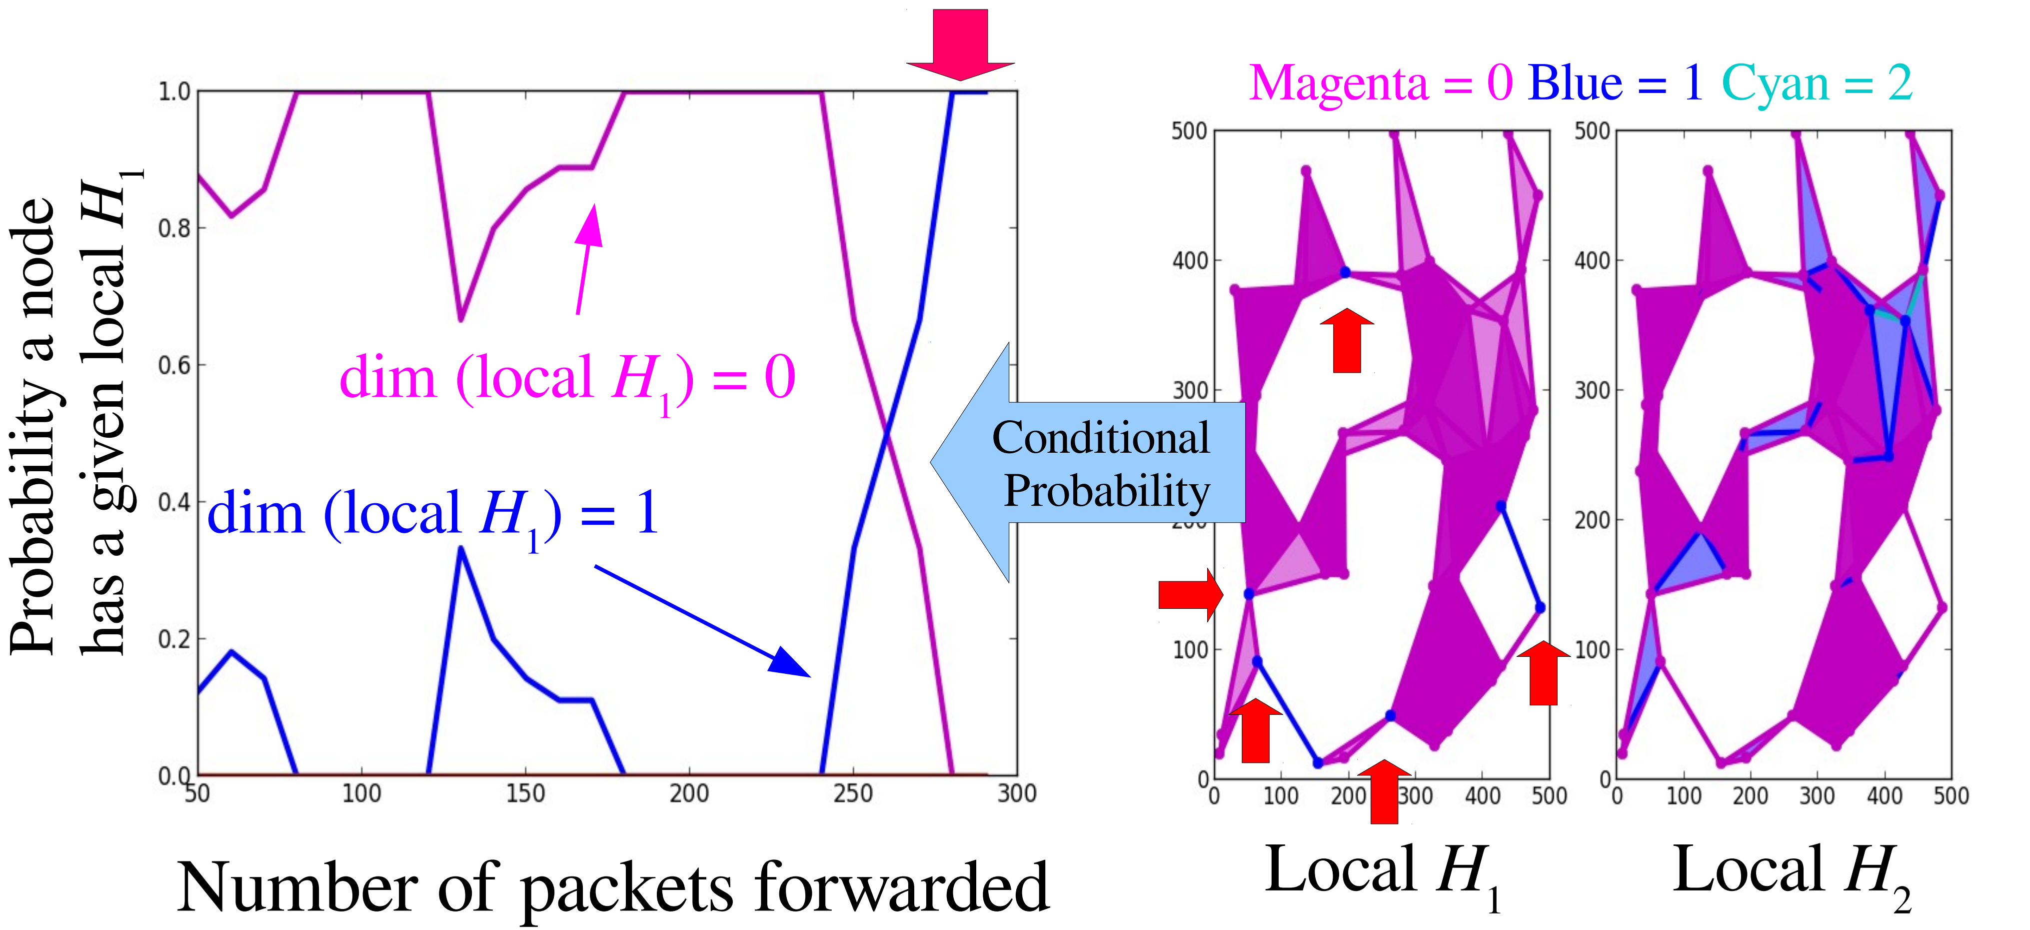

Forwarded packet distributions

- The number of packets forwarded by a node appears to depend on its position in the network

More packets forwarded in topological interior of network

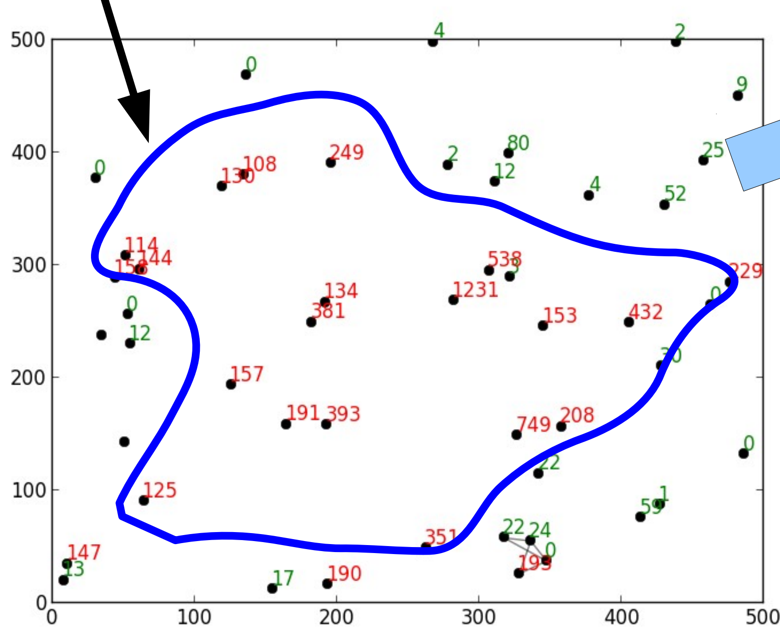

How many packets forwarded by each node

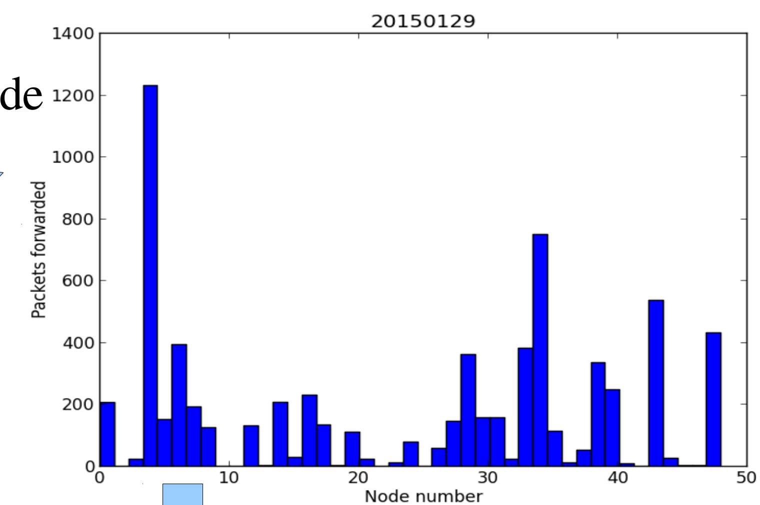

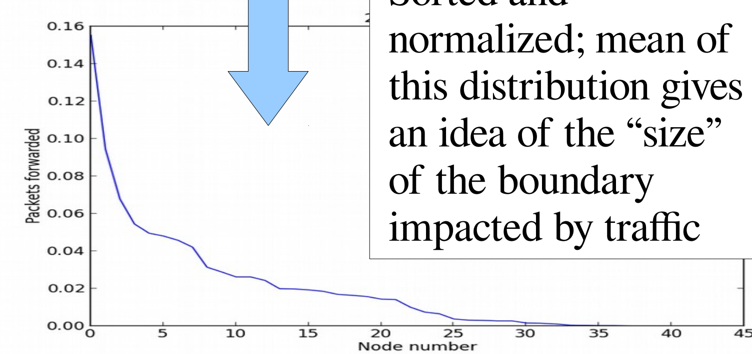

Sorted and normalized; mean of this distribution gives an idea of the "size" of the boundary impacted by traffic

Numbers indicate number of packets forwarded:
Red if over 100, green if less than 100

Homological dimension

If X is a connected manifold of dimension n (or a model of one), then

$$\dim H_k(X, X \setminus \mathrm{star}\ a) = \begin{cases} 1 & \text{if } k = n \\ 0 & \text{otherwise} \end{cases}$$

- You can use this to detect the intrinsic dimension of an embedded manifold...
- … and the **local** dimension of a stratified manifold!

Paul Bendich, Bei Wang , and Sayan Mukherjee, "Local Homology Transfer and Stratification Learning", *Proc. 24th Sympos. on Discrete Algorithms*, pages 1355-1370, 2012

Homological dimension

- Dimension and boundary detection via local homology

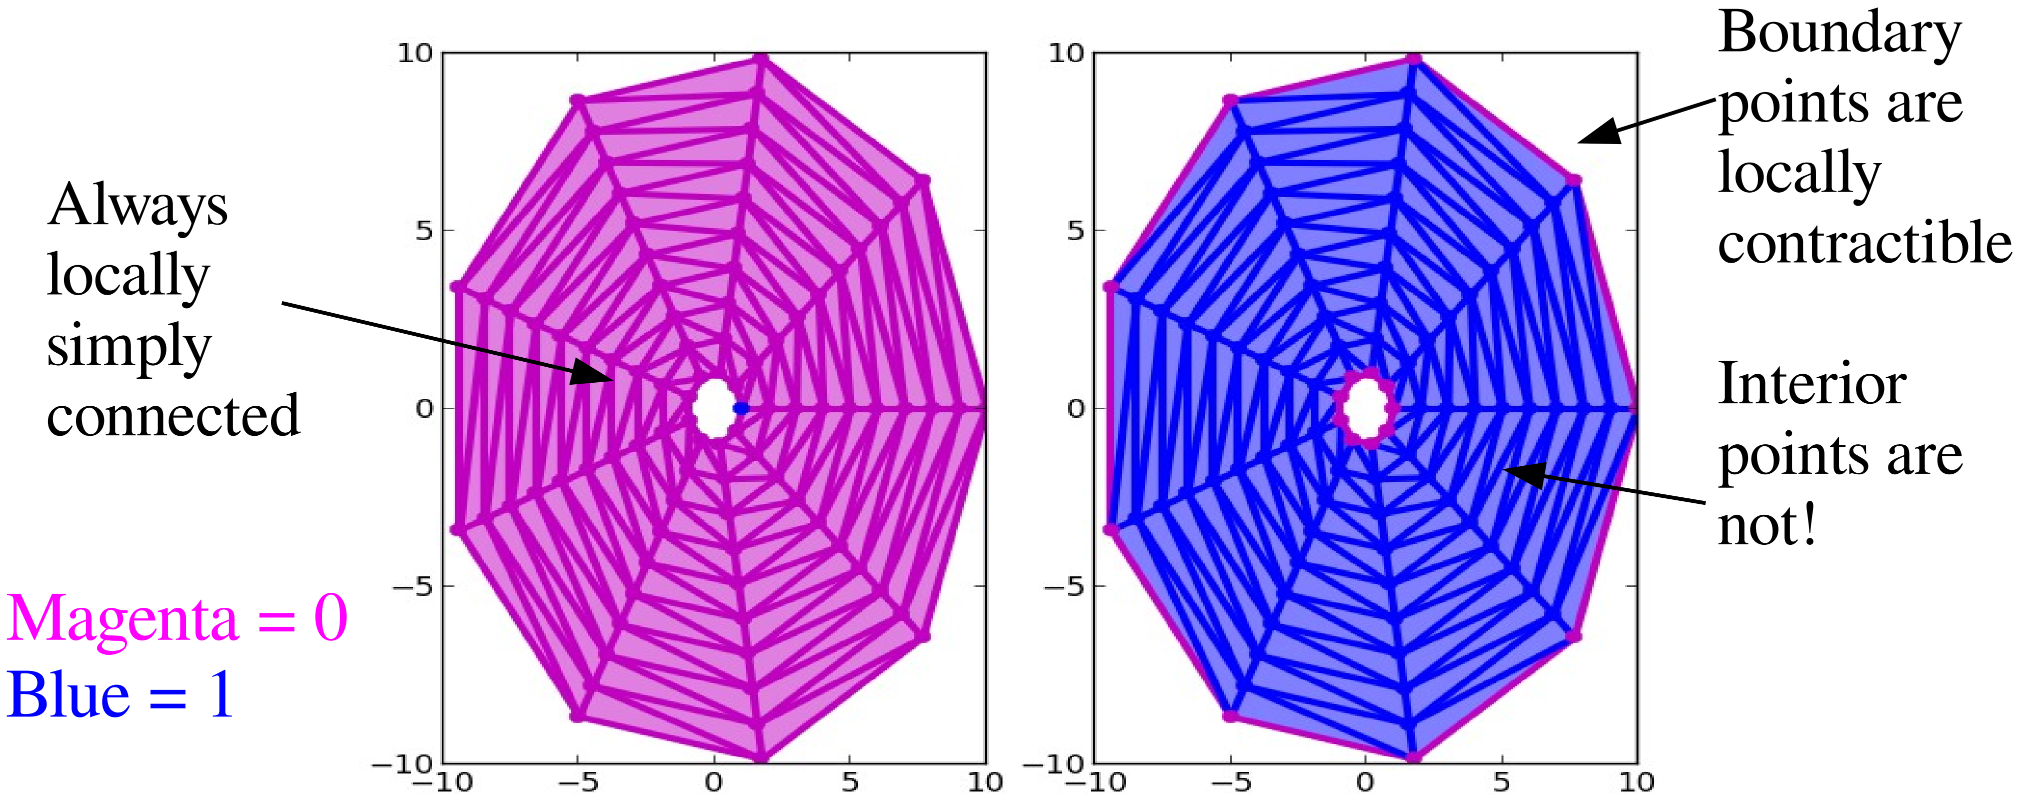

Local H_1 Local H_2

Forwarded packets vs. network "interior"

- Anticipate higher forwarded packet count in the interior: larger local H_2 dimension

Magenta = 0
Blue = 1
Cyan = 2

More packets forwarded in topological interior of network

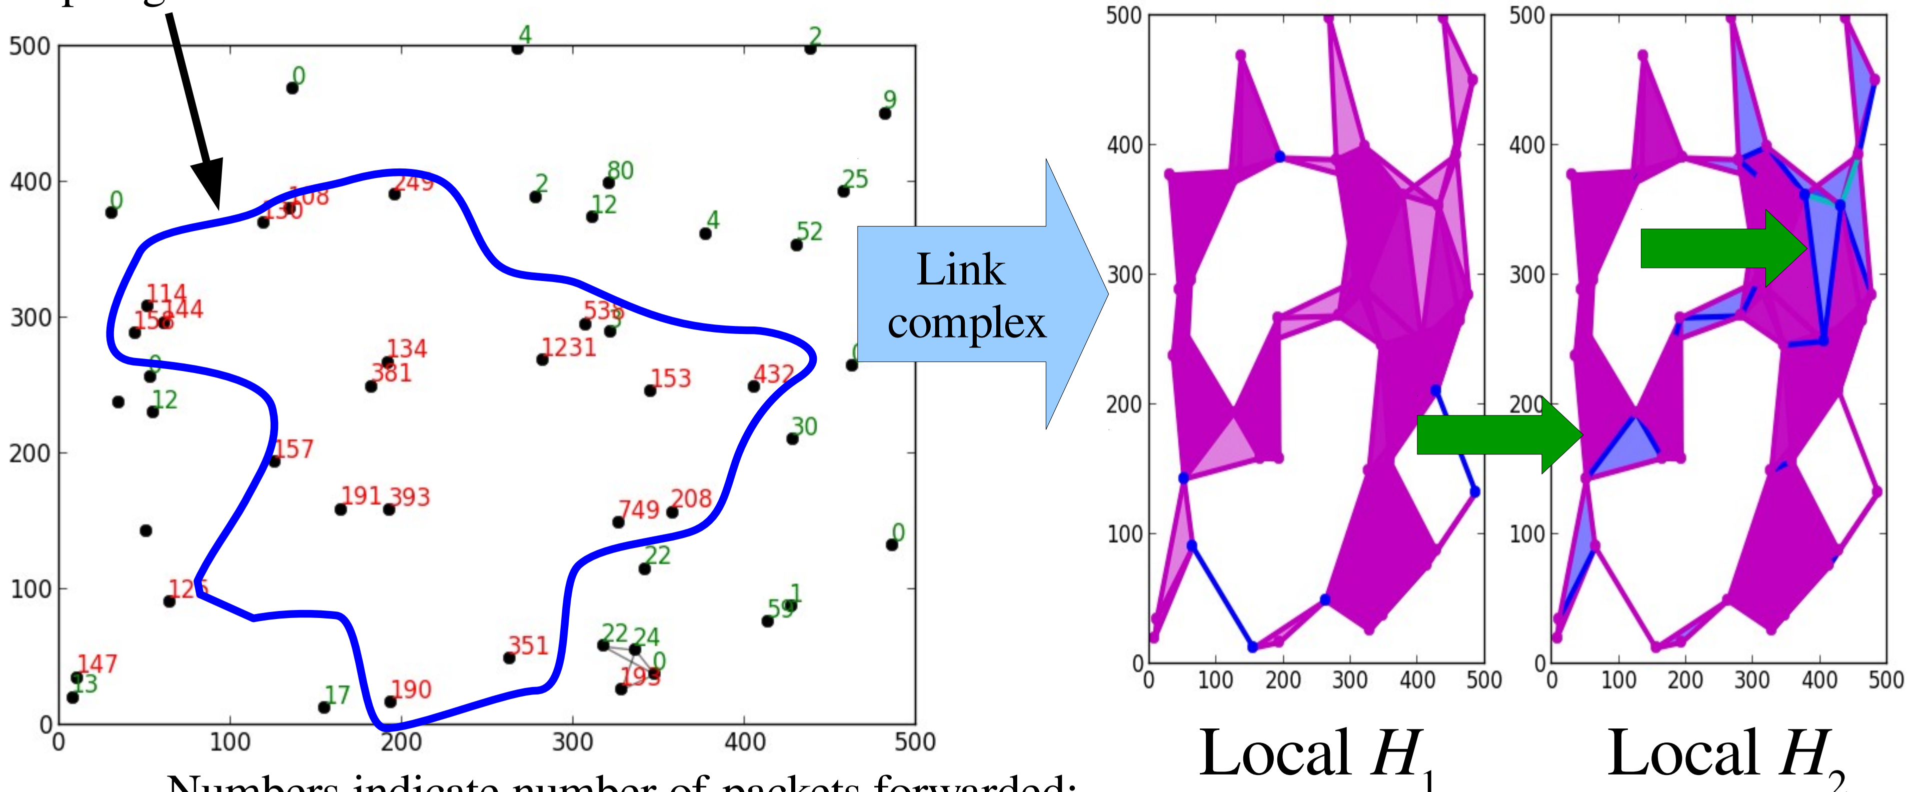

Numbers indicate number of packets forwarded:
Red if over 100, green if less than 100

Forwarded packets vs. network "interior"

- Anticipate higher forwarded packet count in the interior: larger local H_2 dimension

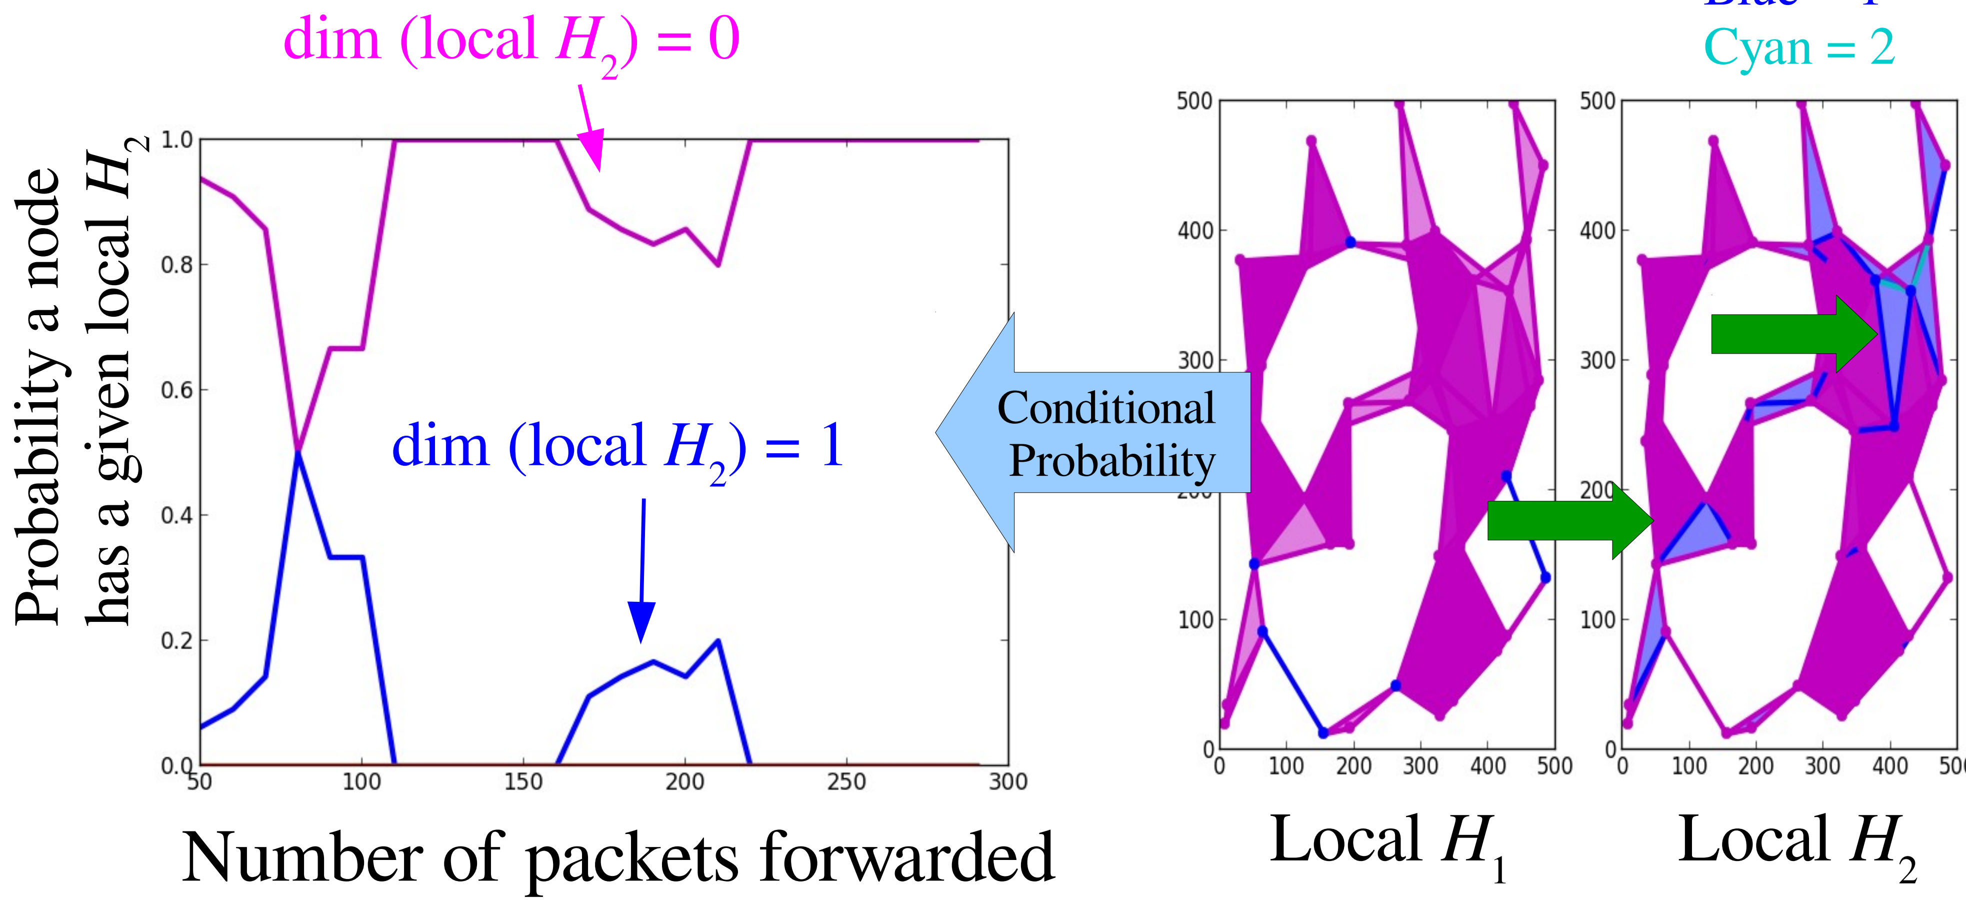

Forwarded packets vs. network "interior"

- Anticipate higher forwarded packet count in the interior: larger local H_2 dimension

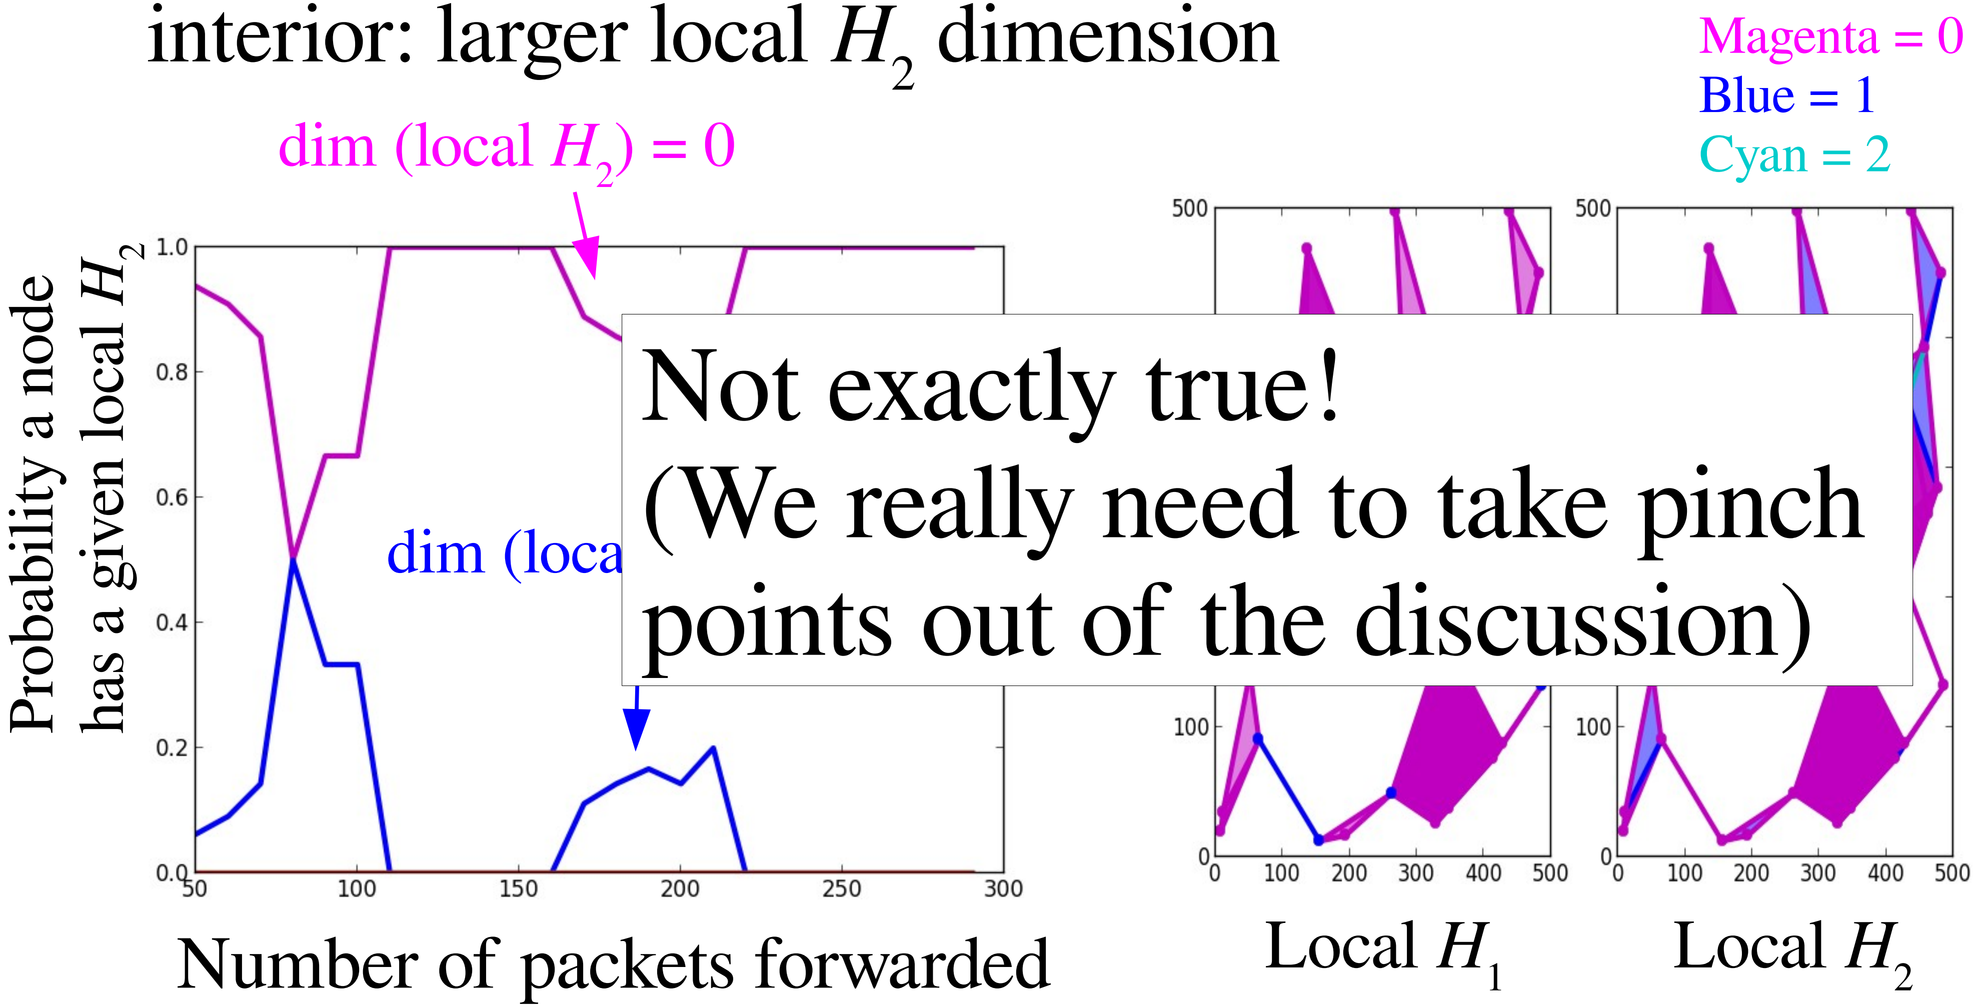

Not exactly true!
(We really need to take pinch points out of the discussion)

Next steps

- Tease apart boundary effects in forwarded packet distributions
- How is topology **of traffic patterns** reflected in traffic statistics?
- Test higher fidelity sheaf models of media access
- Study the topological dependence of dropped packet transients

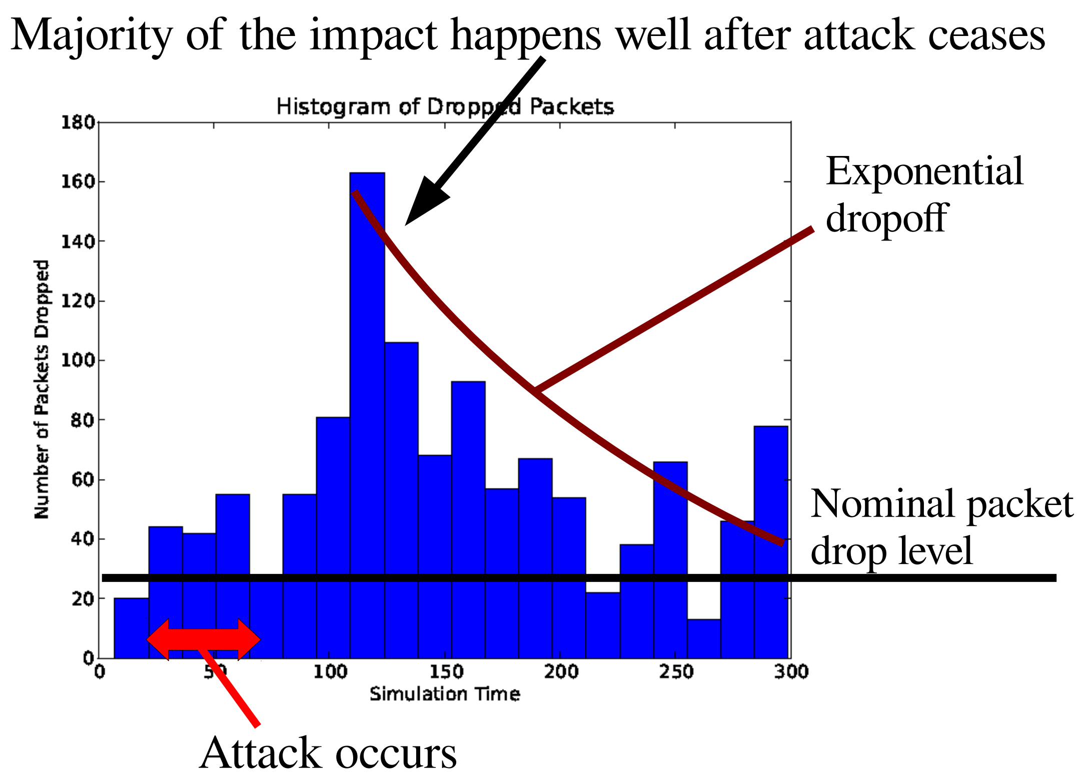

For more information

Michael Robinson

michaelr@american.edu

Preprints available from my website:

http://www.drmichaelrobinson.net/