

HEADER SPACE ANALYSIS: STATIC CHECKING FOR NETWORKS

Peyman Kazemian (Stanford University)

George Varghese (UCSD, Yahoo Labs)

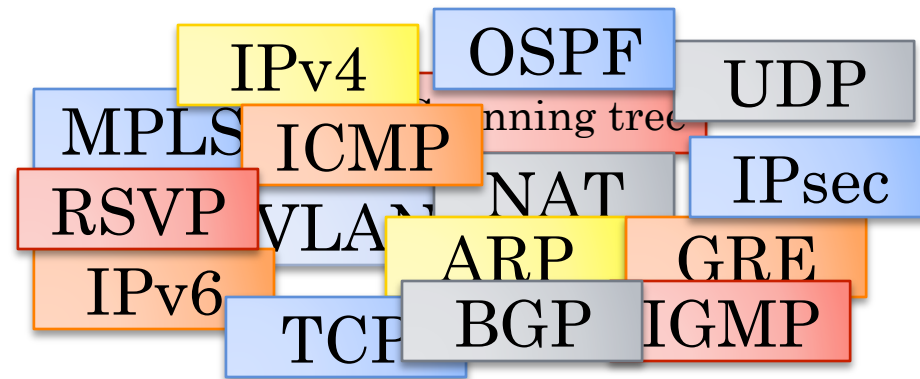
Nick McKeown (Stanford University)

April 25th, 2012

NSDI 2012

TODAY...

- A typical network is a complex mix of protocols:



- Interact in complex ways.
- Cause unforeseen behavior.
- Hard to manage, understand and predict the behavior of networks.

TODAY...

- Even simple questions are hard to answer...
 - Can host A talk to host B?
 - What are all the packet headers from A that can reach B?
 - Are there any loops in the network?
 - Is Slice X isolated totally from Slice Y?
 - What will happen if I remove an entry from a router?

HEADER SPACE ANALYSIS

- A Powerful General Foundation that gives us
 - A unified view of almost all type of boxes.
 - A powerful interface for answering different questions about the network.




HEADER SPACE FRAMEWORK

SIMPLE OBSERVATION: A PACKET IS A POINT IN THE SPACE OF POSSIBLE HEADERS AND A BOX IS A TRANSFORMER ON THAT SPACE.

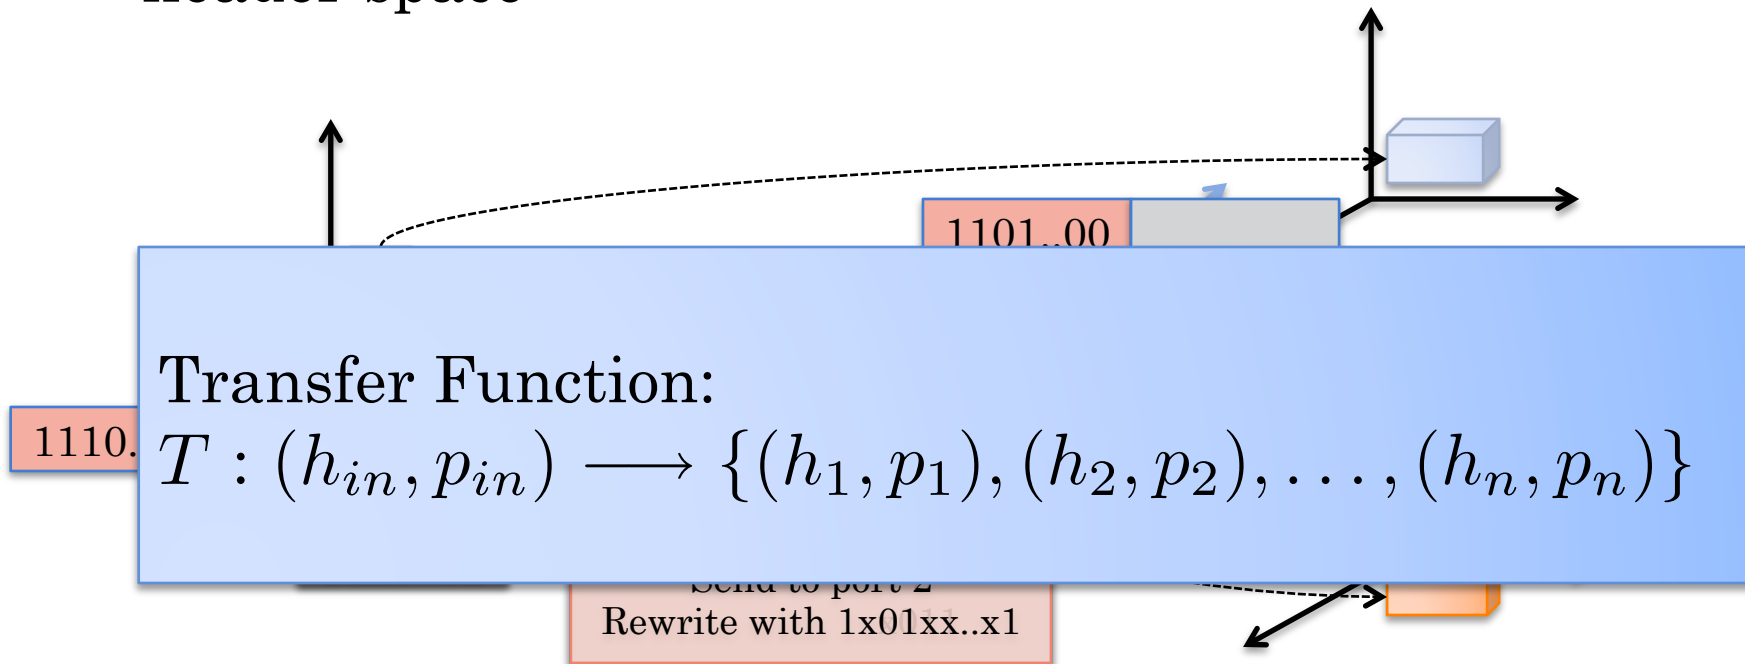
HEADER SPACE FRAMEWORK

- Step 1 - Model packet header as a point in $\{0,1\}^L$ space – The Header Space



HEADER SPACE FRAMEWORK


- Step 2 – Model all networking boxes as transformer of header space



HEADER SPACE FRAMEWORK

○ Example: Transfer Function of an IPv4 Router

- 172.24.74.0 255.255.255.0 Port1
- 172.24.128.0 255.255.255.0 Port2
- 171.67.0.0 255.255.0.0 Port3




$$T(h, p) = \begin{cases} (h,1) & \text{if } \text{dst_ip}(h) = 172.24.74.x \\ (h,2) & \text{if } \text{dst_ip}(h) = 172.24.128.x \\ (h,3) & \text{if } \text{dst_ip}(h) = 171.67.x.x \end{cases}$$

HEADER SPACE FRAMEWORK

○ Example: Transfer Function of an IPv4 Router

- 172.24.74.0 255.255.255.0 Port1
- 172.24.128.0 255.255.255.0 Port2
- 171.67.0.0 255.255.0.0 Port3




$$T(h, p) = \begin{cases} (\text{dec_ttl}(h), 1) & \text{if } \text{dst_ip}(h) = 172.24.74.x \\ (\text{dec_ttl}(h), 2) & \text{if } \text{dst_ip}(h) = 172.24.128.x \\ (\text{dec_ttl}(h), 3) & \text{if } \text{dst_ip}(h) = 171.67.x.x \end{cases}$$

HEADER SPACE FRAMEWORK

○ Example: Transfer Function of an IPv4 Router

- 172.24.74.0 255.255.255.0 Port1
- 172.24.128.0 255.255.255.0 Port2
- 171.67.0.0 255.255.0.0 Port3




$$T(h, p) = \begin{cases} (rw_mac(dec_ttl(h), next_mac), 1) & \text{if } dst_ip(h) = 172.24.74.x \\ (rw_mac(dec_ttl(h), next_mac), 2) & \text{if } dst_ip(h) = 172.24.128.x \\ (rw_mac(dec_ttl(h), next_mac), 3) & \text{if } dst_ip(h) = 171.67.x.x \end{cases}$$

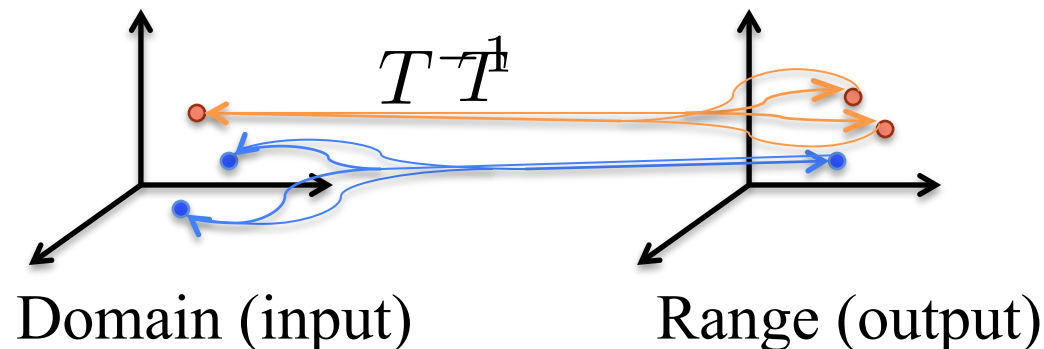
HEADER SPACE FRAMEWORK

- Properties of transfer functions

- Composable: $T_3(T_2(T_1(h, p)))$




- Invertible:



HEADER SPACE FRAMEWORK

- Step 3 - Develop an algebra to work on these spaces.
- Every object in Header Space, can be described by union of Wildcard Expressions.
- We want to perform the following set operations on wildcard expressions:
 - Intersection
 - Complementation
 - Difference



HEADER SPACE FRAMEWORK

○ Finding Intersection:

- Bit by bit intersect using intersection table:
 - Example: $10xx \cap 1xx0 = 10x0$
 - If result has any 'z', then intersection is **empty**:
 - Example: $10xx \cap 0xx1 = z0x1 = \phi$

$b_i \backslash b_j$	0	1	x
0	0	z	0
1	z	1	1
x	0	1	x

- See the paper for how to find complement and difference.

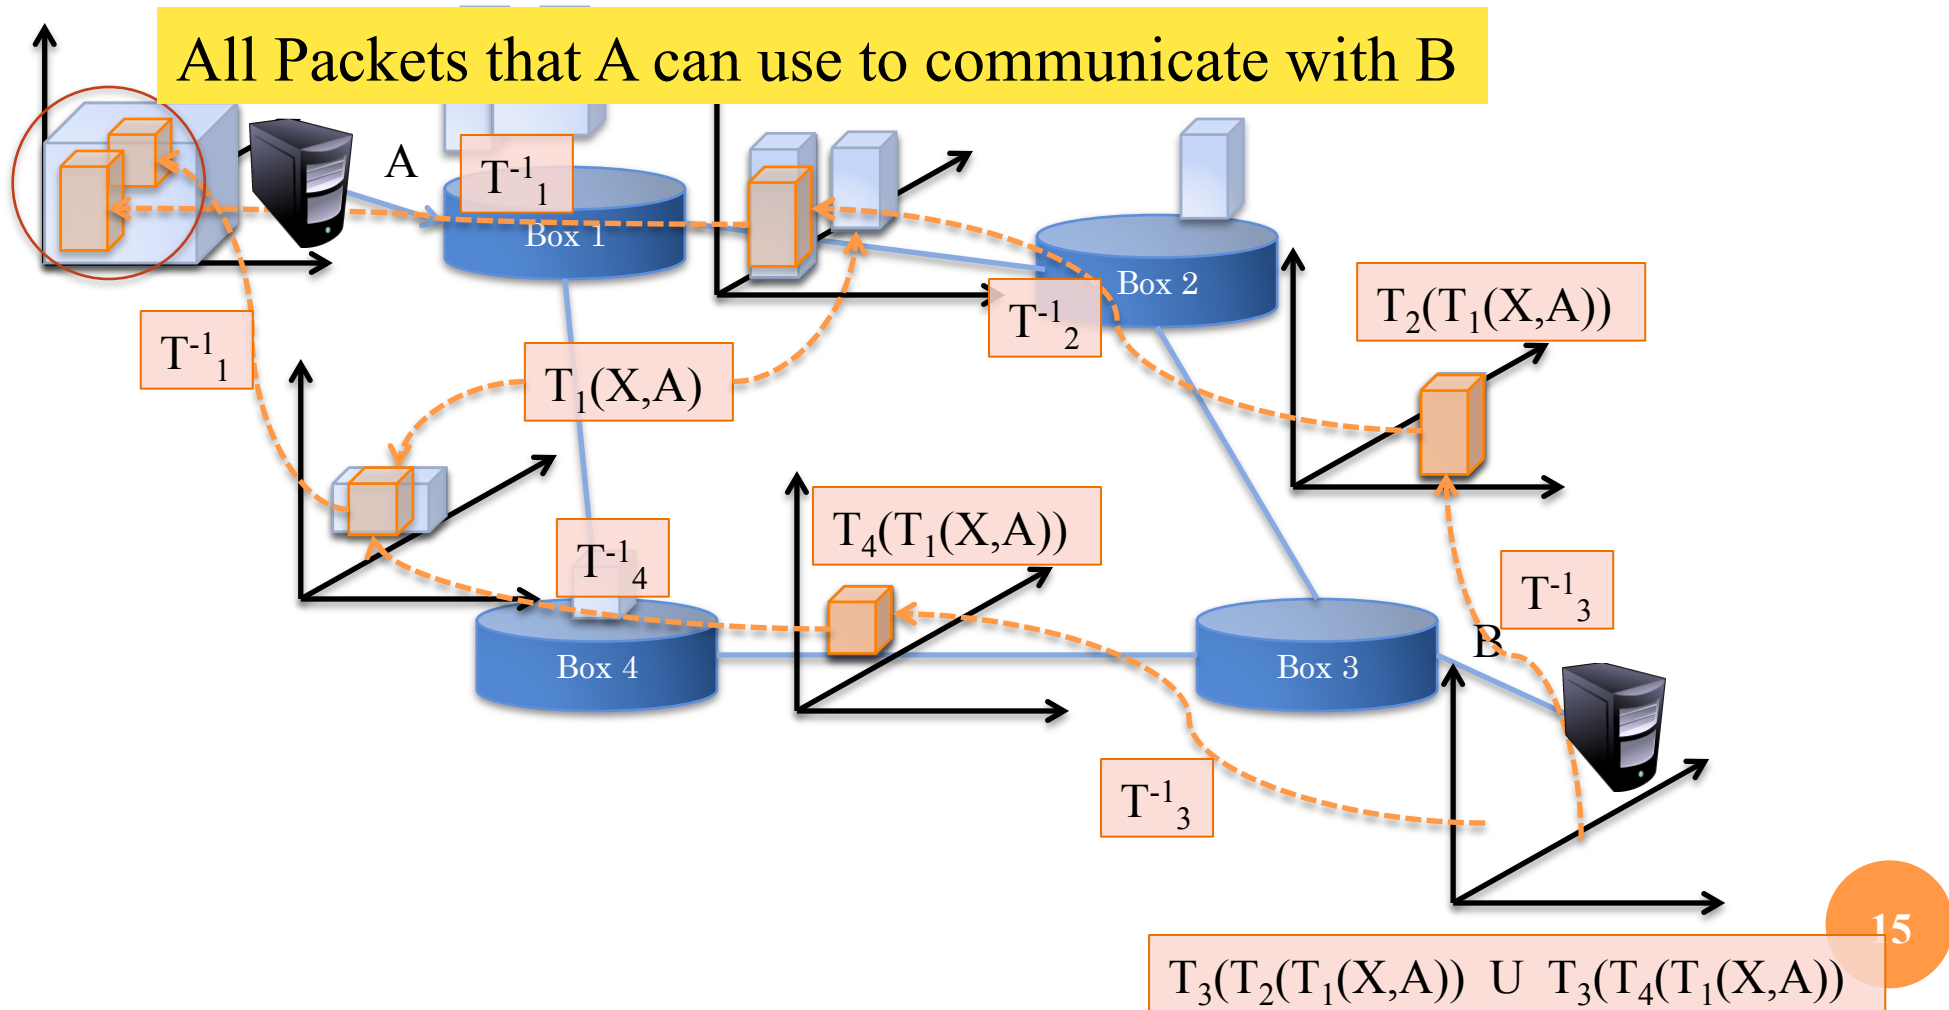


USE CASES OF HEADER SPACE FRAMEWORK

THESE ARE ONLY SOME EXAMPLE USE CASES THAT
WE DEVELOPED SO FAR...

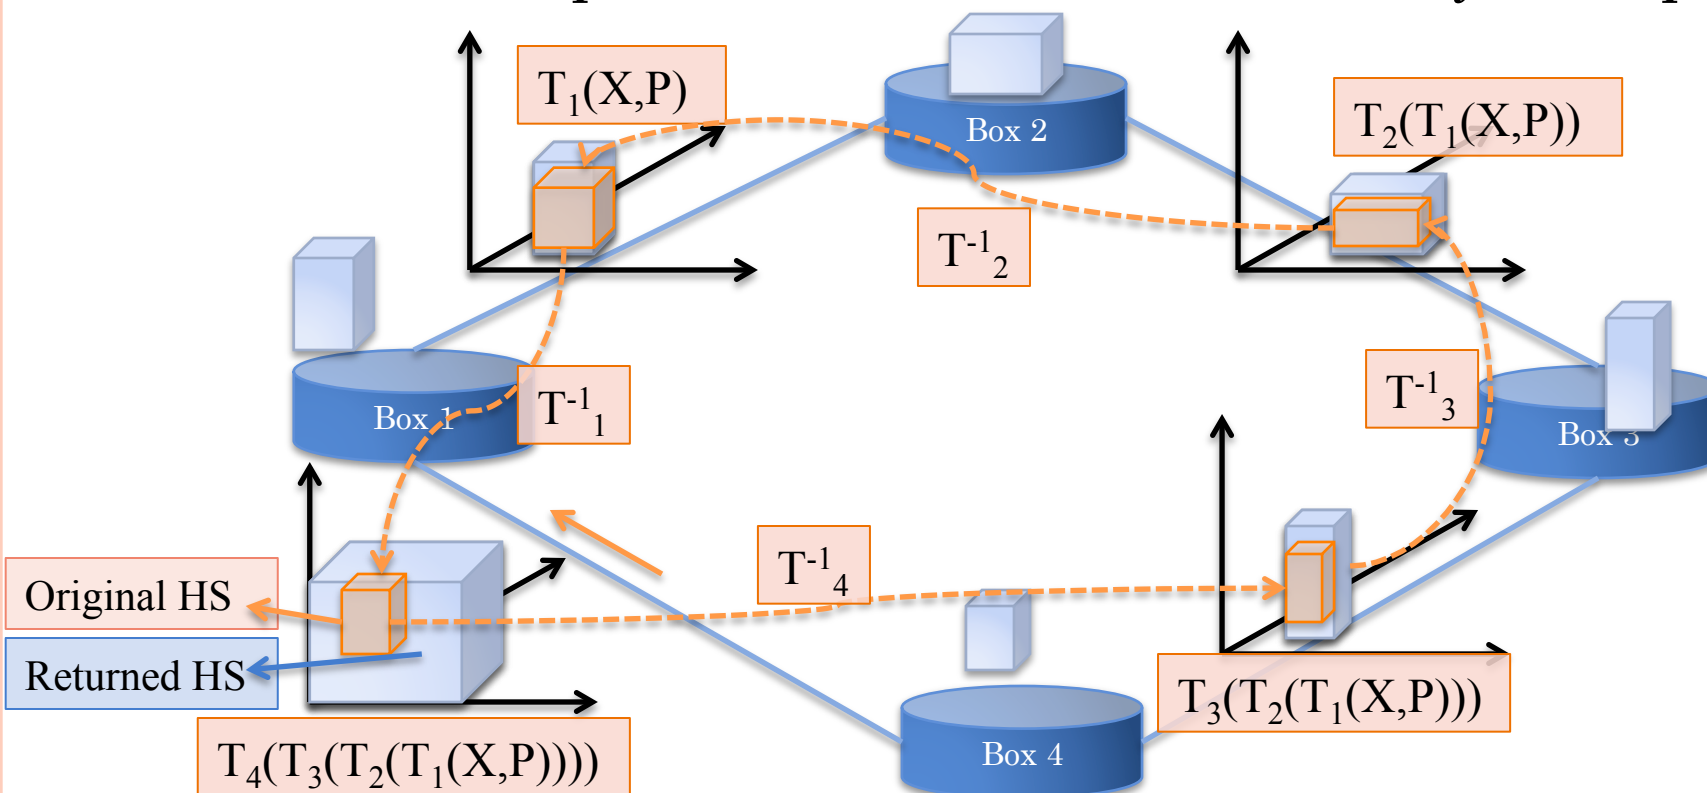
USE CASES

- Can host A talk to B?




USE CASES

- Is there a loop in the network?
 - Inject an all-x text packet from every switch-port
 - Follow the packet until it comes back to injection port




USE CASES


- Is the loop infinite?



Finite Loop



Infinite Loop



?


USE CASES

- Are two slices isolated?
- What do we mean by slice?
 - Fixed Slices: VLAN slices
 - Programmable Slices: slices created by FlowVisor
- Why do we care about isolation?
 - Banks: for added security.
 - Healthcare: to comply with HIPAA.
 - GENI: to isolate different experiments running on the same network.

USE CASES

- Are two slices isolated?
 - 1) slice definitions don't intersect.
 - 2) packets do not leak.

Solution: Apply header space reservation of each slice to the slice's transfer function and check intersection of the result with other slices' reservations



COMPLEXITY

○ Run time

Reachability: $O(dR^2)$

Loop Detection: $O(dPR^2)$

- R: maximum number of rules per box.
- d: diameter of network.
- P: number of ports to be tested

Slice Isolation Test: $O(NW^2)$

- W: number of wildcard expressions in definition of a slice.
- N: number of slices in the network.

See paper for more details.

HEADER SPACE FRAMEWORK

- A Powerful General Foundation that gives us
 - A unified view of almost all type of boxes.
 - Transfer Function.
 - A powerful interface for answering different questions about the network.
 - $T(h,p)$ and $T^{-1}(h,p)$
 - Set operations on Header Space



IMPLEMENTATION AND EVALUATION


22

IMPLEMENTATION

○ Header Space Library (**Hassel**)

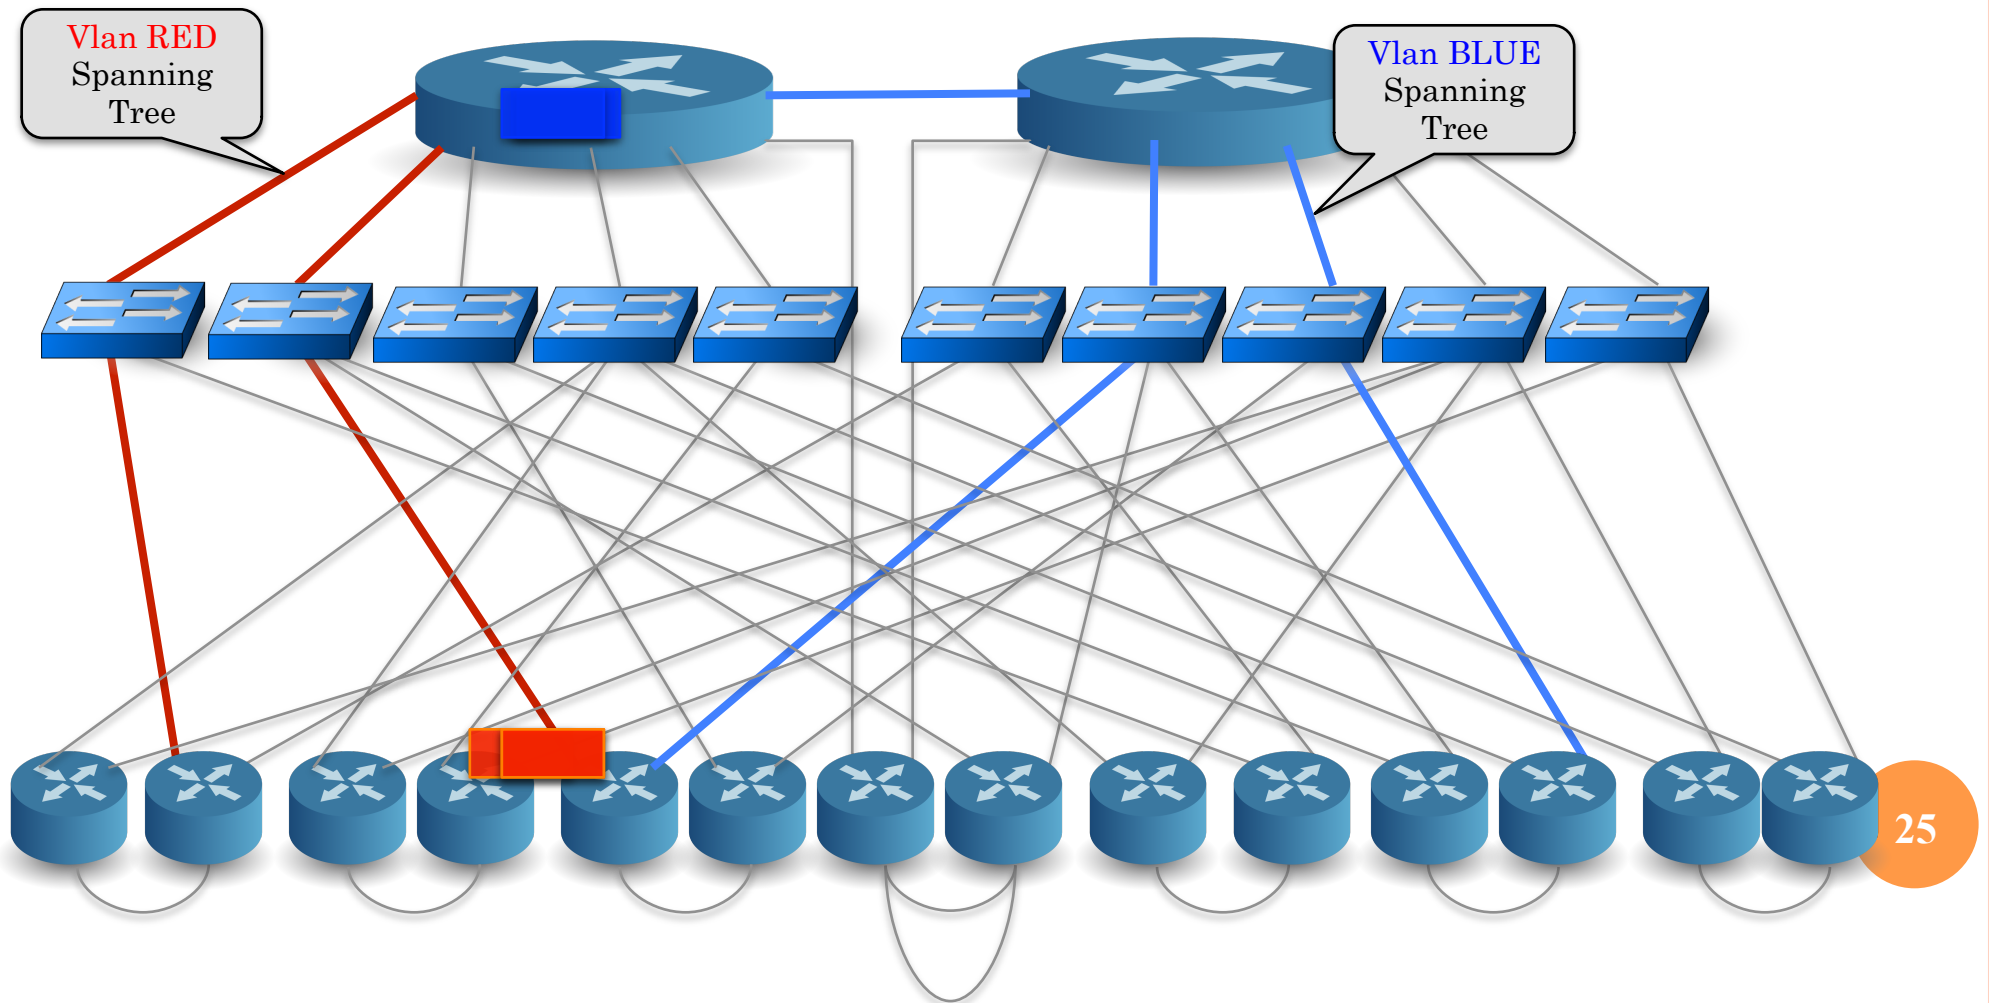
- Written in Python
- Implements **Header Space** Class
 - Set operations
- Implements **Transfer Function** Class
 - T and T^{-1}
- Implements Reachability, Loop Detection and Slice Isolation checks.
 - < 50 lines of code
- Includes a Cisco IOS parser
 - Generates transfer function from output of IOS commands and config file.
 - Keeps the mapping from Transfer function rule to line number in config file.
- Publicly available: git clone <https://bitbucket.org/peymank/hassel-public.git>

STANFORD BACKBONE NETWORK



STANFORD BACKBONE NETWORK

- Loop detection test – run time < 10 minutes on a single laptop.



PERFORMANCE

Performance result for Stanford Backbone Network on a single machine: 4 core, 4GB RAM.

Generating TF Rules	~150 sec
Loop Detection Test (30 ports)	~560 sec
Average Per Port	~18 sec
Min Per Port	~ 8 sec
Max Per Port	~ 135 sec
Reachability Test (Avg)	~13 sec

SUMMARY

- We Introduced Header Space Analysis:

A Powerful General Foundation that gives us

- A unified view of almost all type of boxes.
- A powerful interface for answering different questions about the network.

- We showed that our initial Python-based implementation can scale to enterprise-size networks on a single laptop.

Thank You!

Questions?