



ROYAL INSTITUTE
OF TECHNOLOGY


Fast Failure Detection on Bifrost

Voravit Tanyingyong


Siraj Rathore

Background – Carenet Project

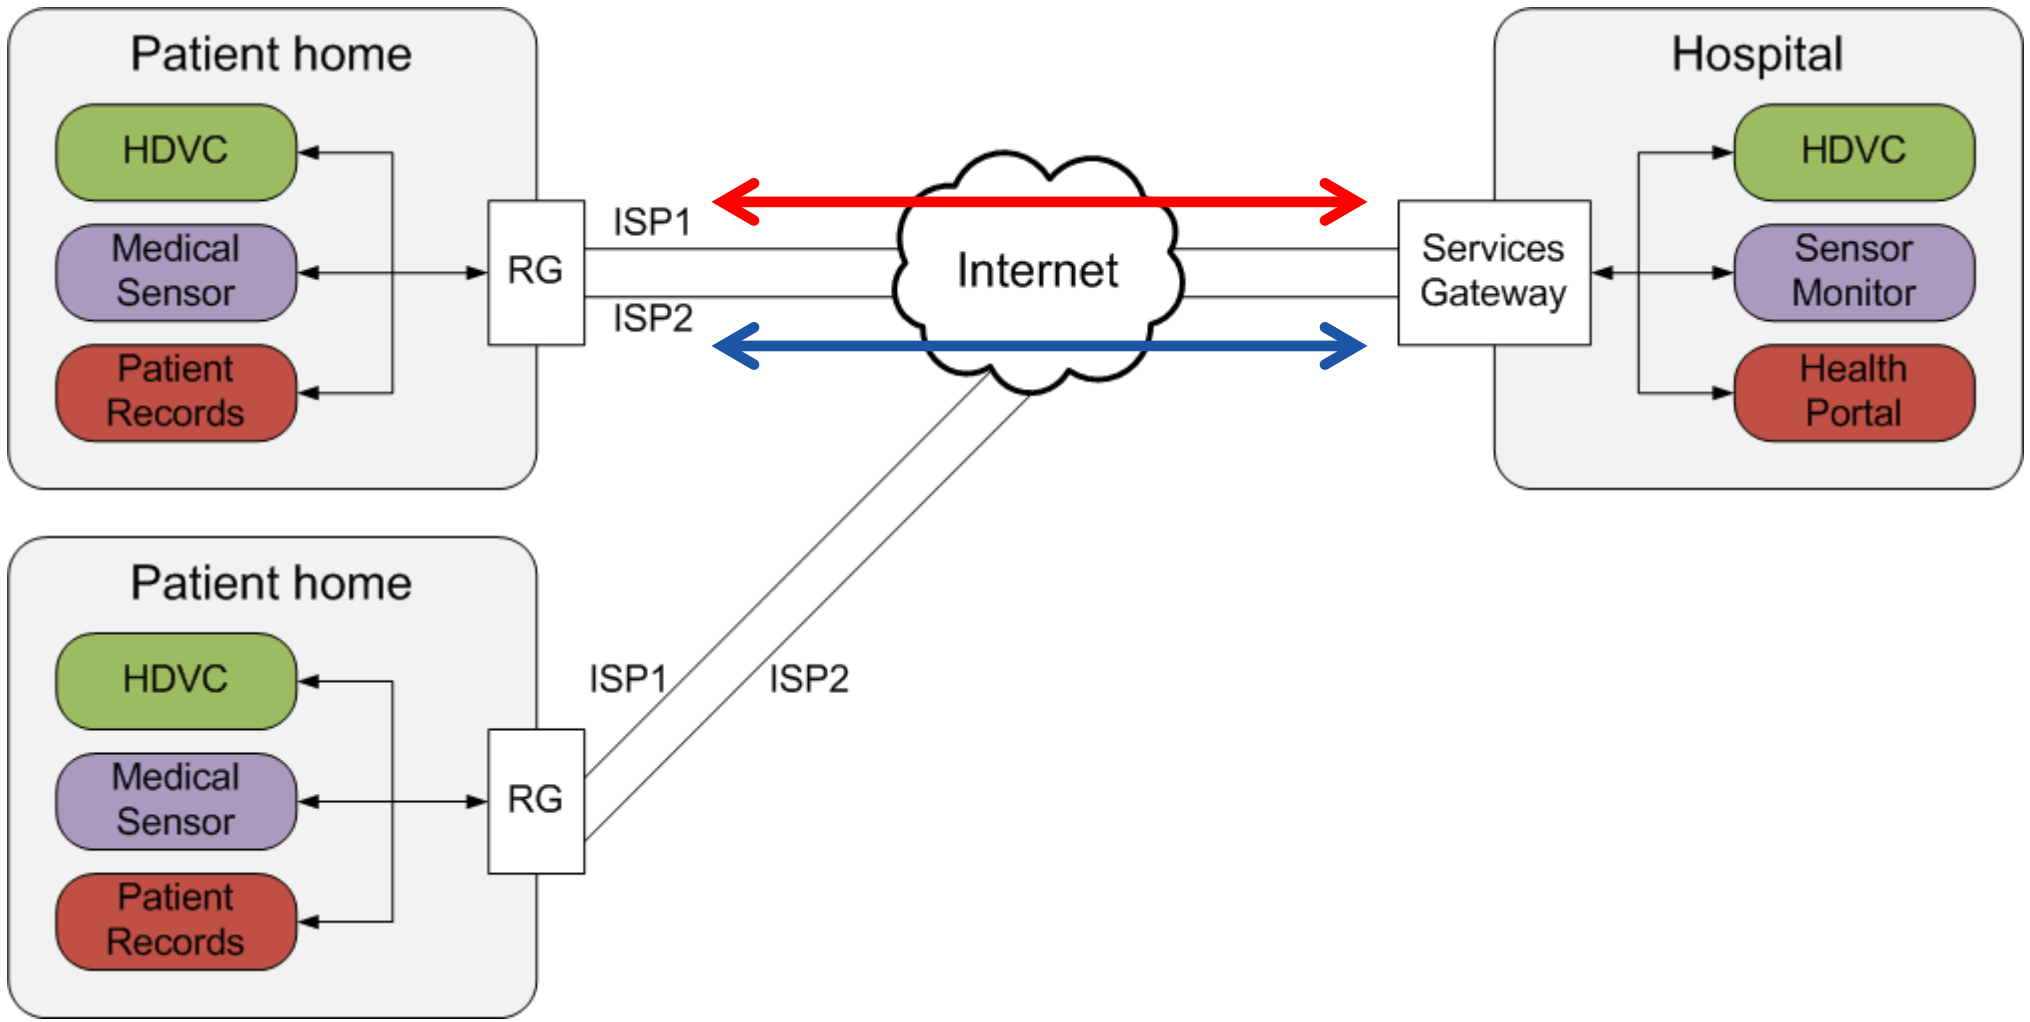
- Research project at Network Systems Lab
- Design of a robust and secure Residential Gateway (RG) that facilitates the interconnection of a dedicated private healthcare network into the homes of patients



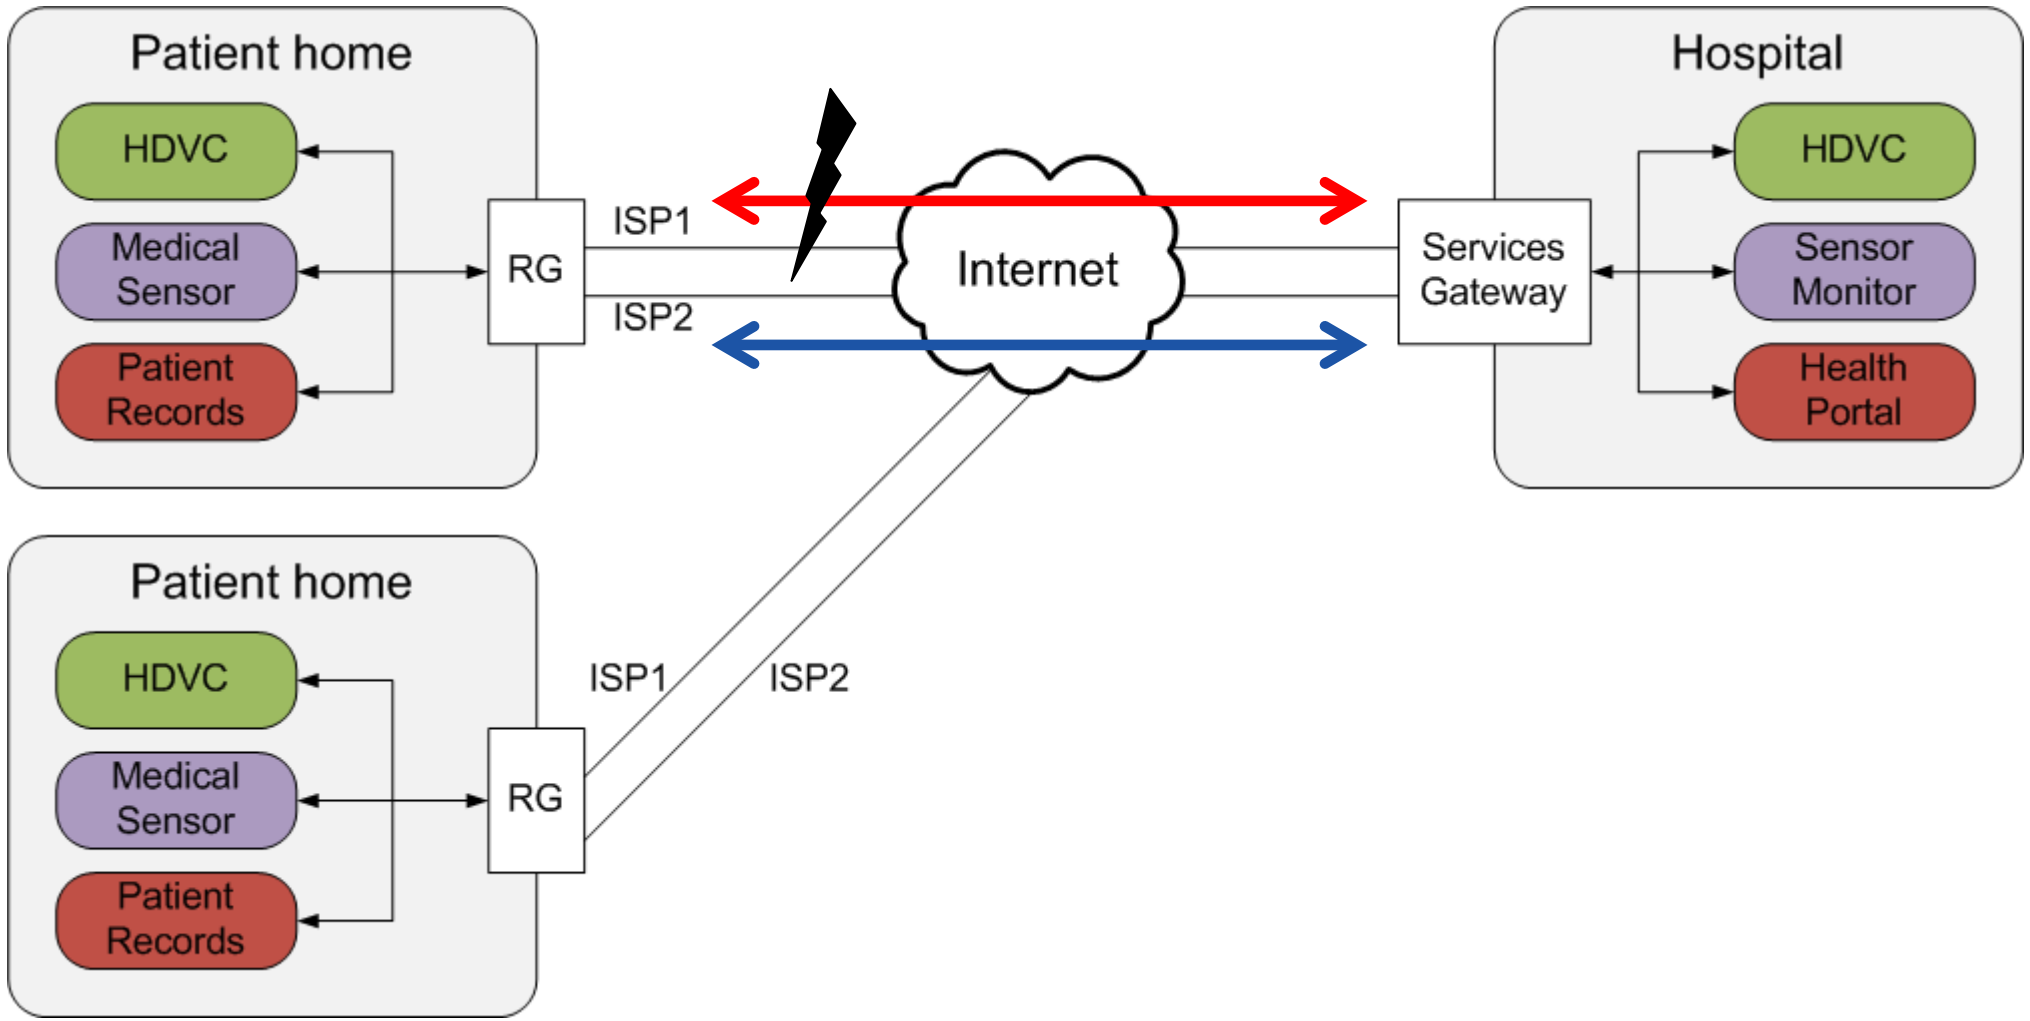
Carenet Topology



Carenet Topology



Carenet Topology



Fast Detection Alternatives

- Standard Routing protocol with modified hello interval
 - + Neighbor discovery
 - + Reacts to changes dynamically
 - Slow hello mechanism (in the order of seconds)
 - Higher overhead as hello interval decreases
- Bidirectional Forwarding Detection protocol (BFD)
 - + Protocol independent
 - + Low-overhead (24-bytes control packet)
 - + Short-duration detection of failures (subseconds)
 - No discovery mechanism
 - Failure detection only!

BFD Implementations

- Quagga
 - BGP with BFD
 - <http://sourceforge.net/projects/bfdd/>
- User-space Application
 - BFD daemon by Tom Phelan
 - <http://www.phelan-4.com/bfd/>
- Kernel module
 - Kbfd also with Quagga patch for zebra-0.95a
 - <http://kbfd.sourceforge.net/>

Implementations in consideration


- Standard Routing protocol with modified hello interval
 - Quagga OSPF with fast-hello
- Bidirectional Forwarding Detection protocol (BFD)
 - Quagga BGP with BFD support
 - BFD application with script to statically reconfigure routes



Failure Detection Time

- OSPF
 - Detection Time = Dead Interval (default 40s)
 - Minimum dead interval = 1s
 - Minimum hello = 100ms (default 10s)
- BFD
 - Detection Time = Hello Interval x Multiplier
 - Minimum hello Interval
 - Quagga+BGP/BFD: 200ms
 - BFD+static routes: 100ms
 - Minimum Multiplier
 - Quagga+BGP/BFD: 3
 - BFD+static routes: 2

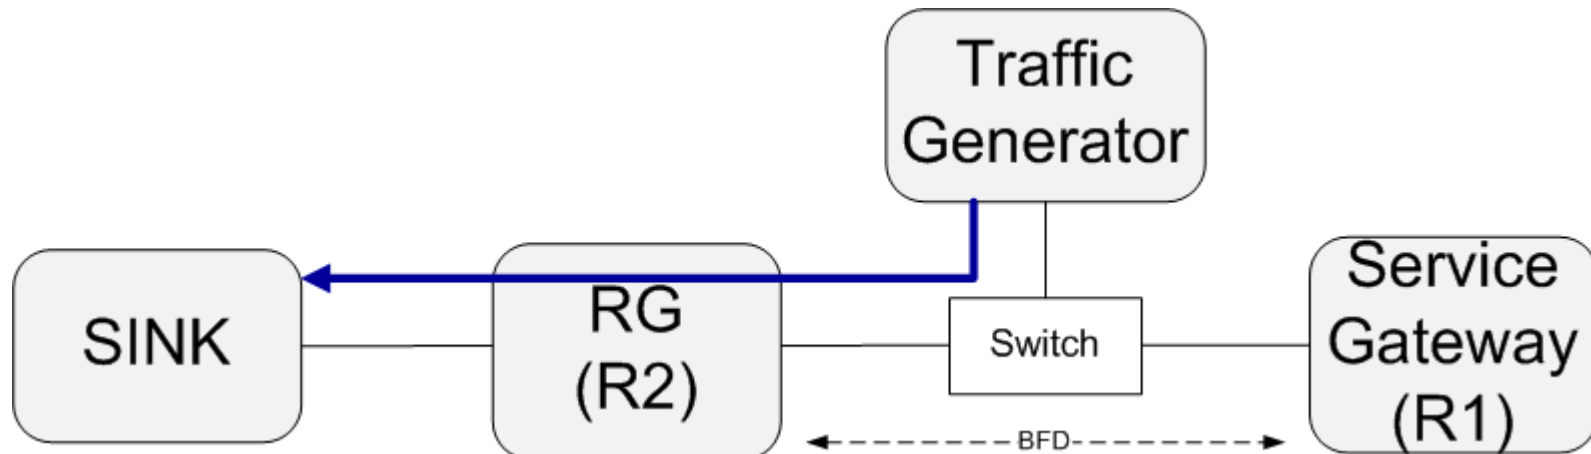
Failover Time Comparison



Detection method	Hello (ms)	Mult	Dead Interval (ms)	Failover (ms)
OSPF fast hello	100	10	1000	988.5162
Quagga/BGP+BFD	200	3	600	588.3378
BFD+Static routes	200	3	600	537.6276

Closer look at BFD + Static routes

- How traffic load affects BFD?
- Fixed detection time = 600ms
- Vary multiplier and hello interval
 - Multiplier/Hello: 6/100, 3/200, 2/300
- Traffic load duration: 1 mintues



Saturation Test (Link + CPU)

- False-Alarm: reported downtime in milliseconds
- Status Flapping: BFD status switching up-down-up
- Average from 5 iterations
- Load: 64-byte packet size at line-rate (1Gbps)

R1	6	3	2
False-Alarm	0	433	5421.2
Status flapping	0	0.2	4.4

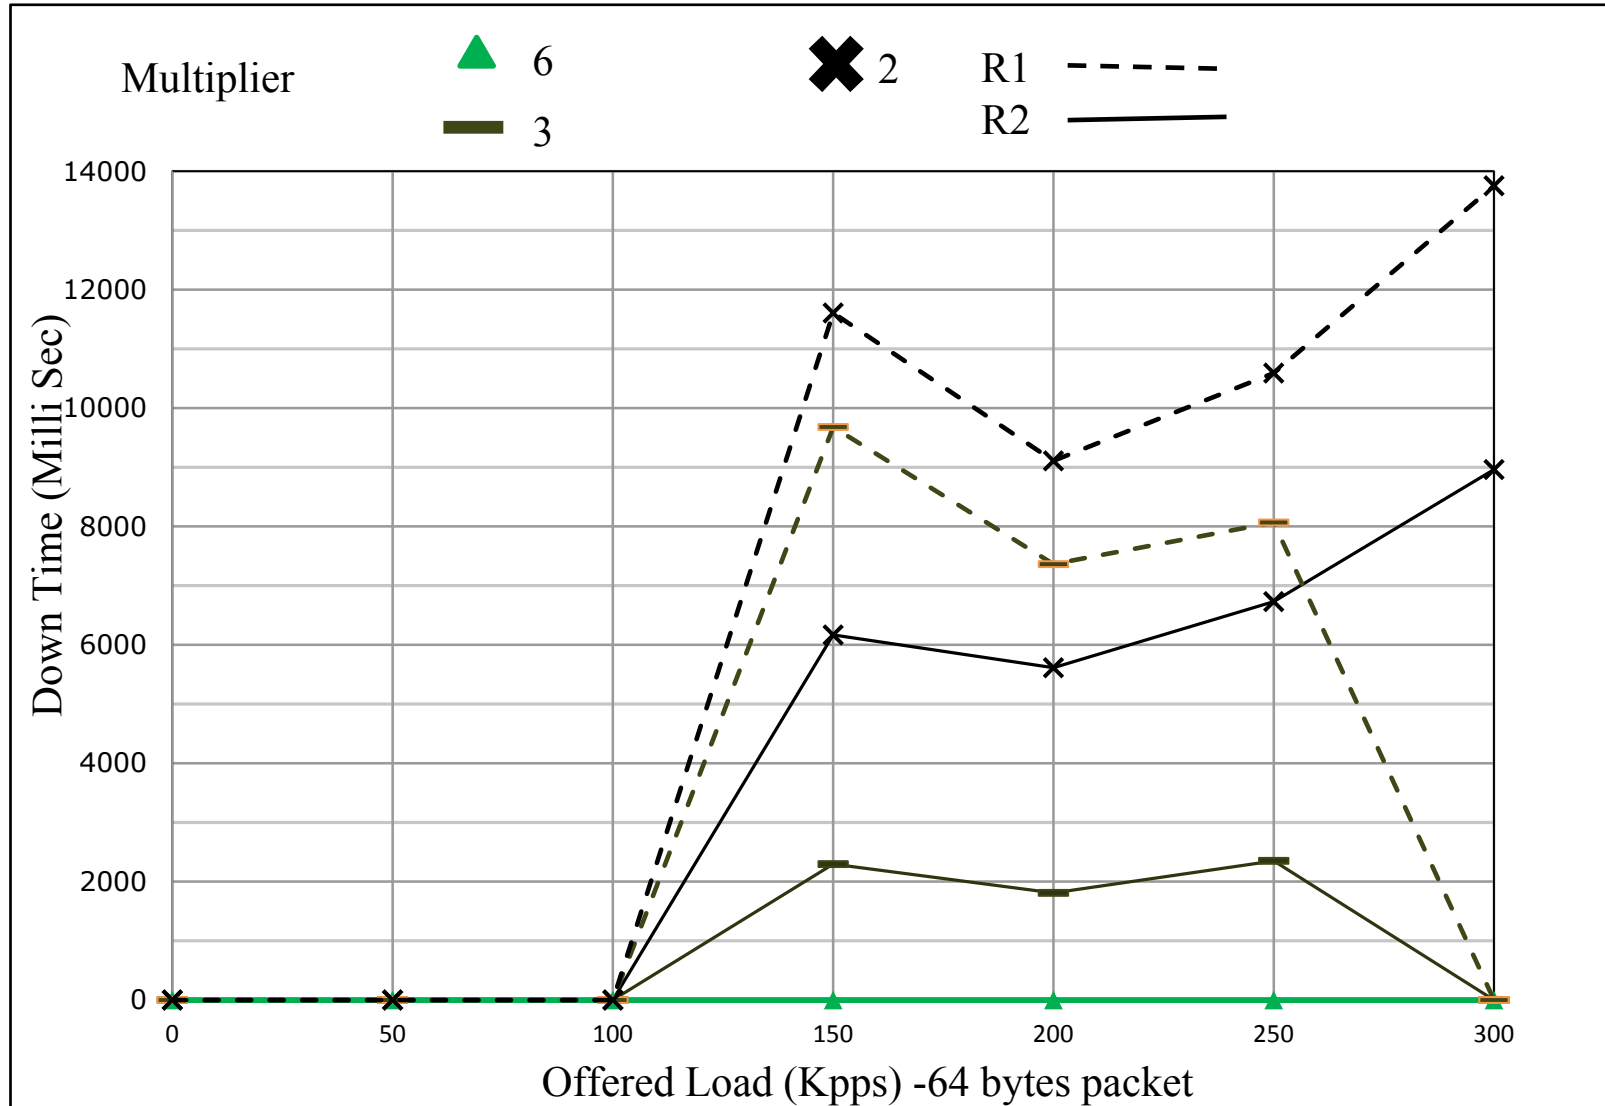
R2	6	3	2
False-Alarm	0	186.2	3649.8
Status flapping	0	0.2	4.4

BFD under varies traffic load

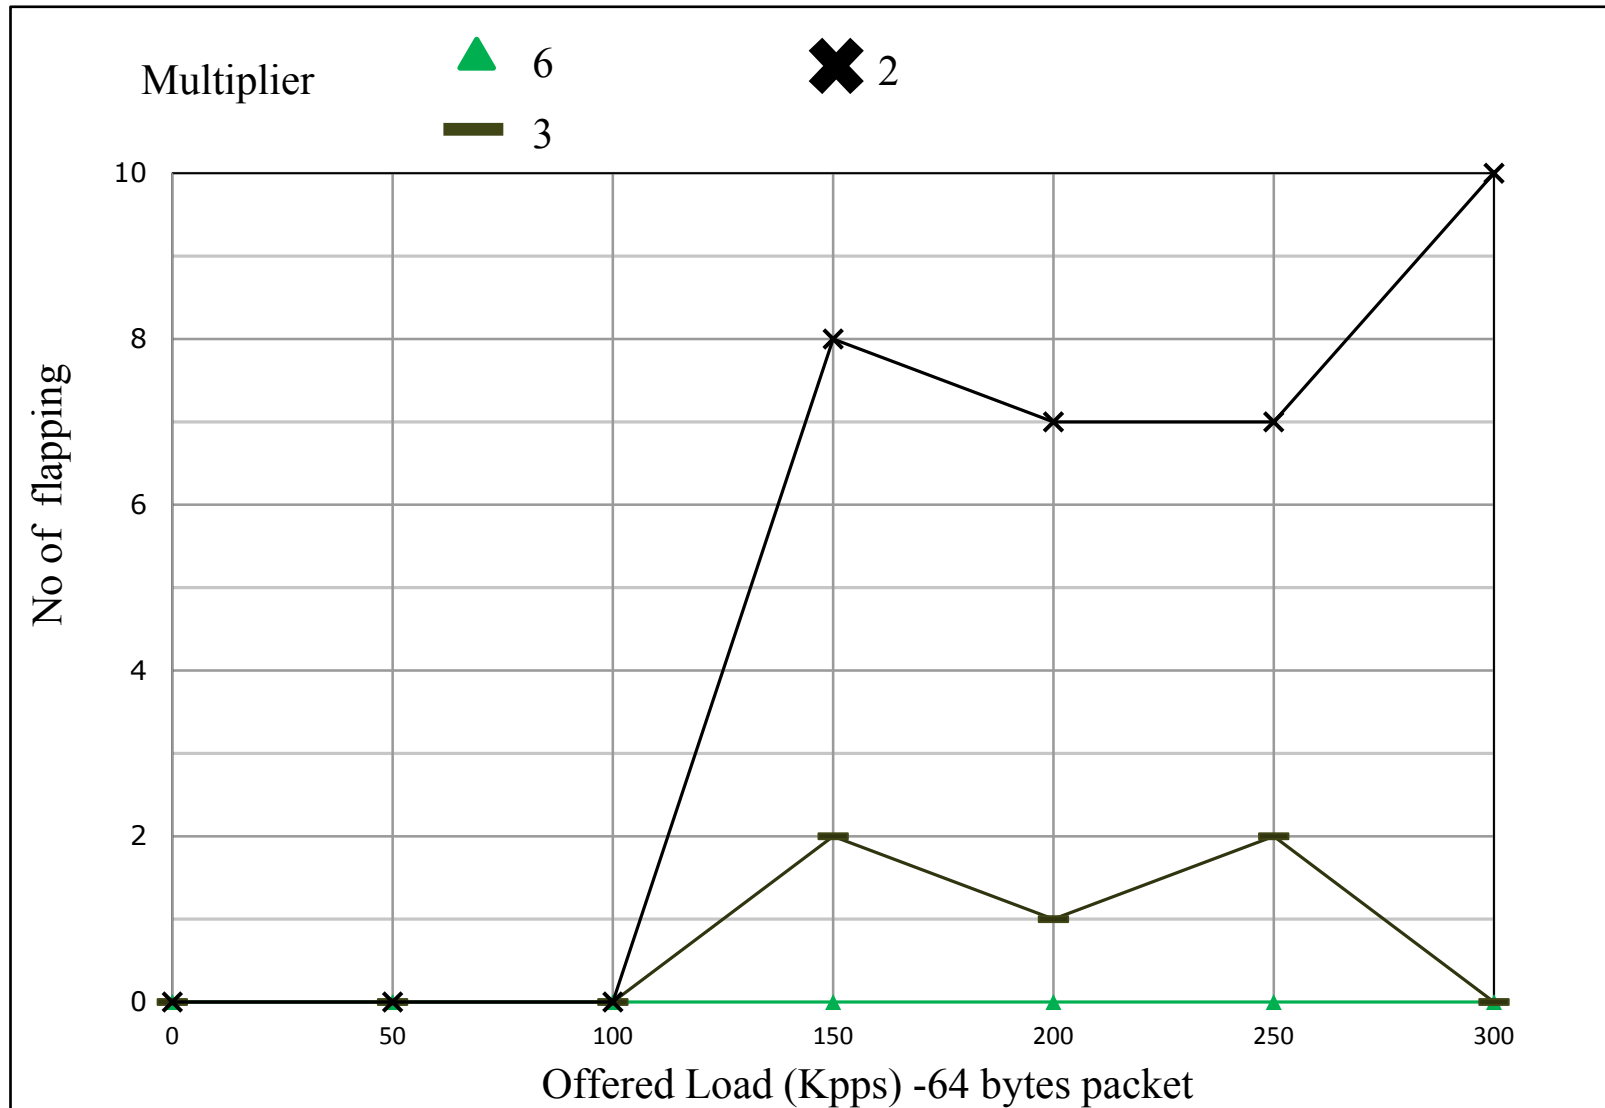
- False-Alarm: reported downtime in milliseconds
- Status Flapping: BFD status switching up-down-up
- Load: 64-byte packet size at different rates
- Test duration: 1 minute

Packet Rate	6	3	2
50000	51.25	51.43	51.17
100000	87.77	87.87	87.87
150000	100	100	100
200000	100	100	100
250000	99.97	99.97	100
300000	100	99.97	99.97

False-Alarm (Reported Downtime)



Number of Status Flapping



Future Direction

- Performance with tunnel settings
- Influences from other link characteristics such as delay, and loss
- How to implement BFD on a software router in a way that make it more robust and less prone to errors and false-positives



**ROYAL INSTITUTE
OF TECHNOLOGY**

Thank You for Listening!