



# Peer-to-peer meets wireless: Opportunities and challenges

Antony Rowstron

Microsoft Research, Cambridge



# Talk outline


---

- Introduction (Overlays, underlays and wireless )
- Wireless routing inspired by structured overlays
  - Mesh networks
    - Design of Virtual Ring Routing (VRR)
    - Evaluation of VRR
  - Vehicular Networking
    - Design concepts of PVRP
    - Preliminary results



# Structured overlays/DHTs


- Experiences learnt with structured overlays
  - Self-organizing
  - Fault-tolerant
  - Scalable
  - Decentralized
  - Performance






# Structured overlays/DHTs

- Experiences learnt with structured overlays
  - Self-organizing
  - Fault-tolerant
  - Scalable
  - Decentralized
  - Performance





# Overlay versus underlay

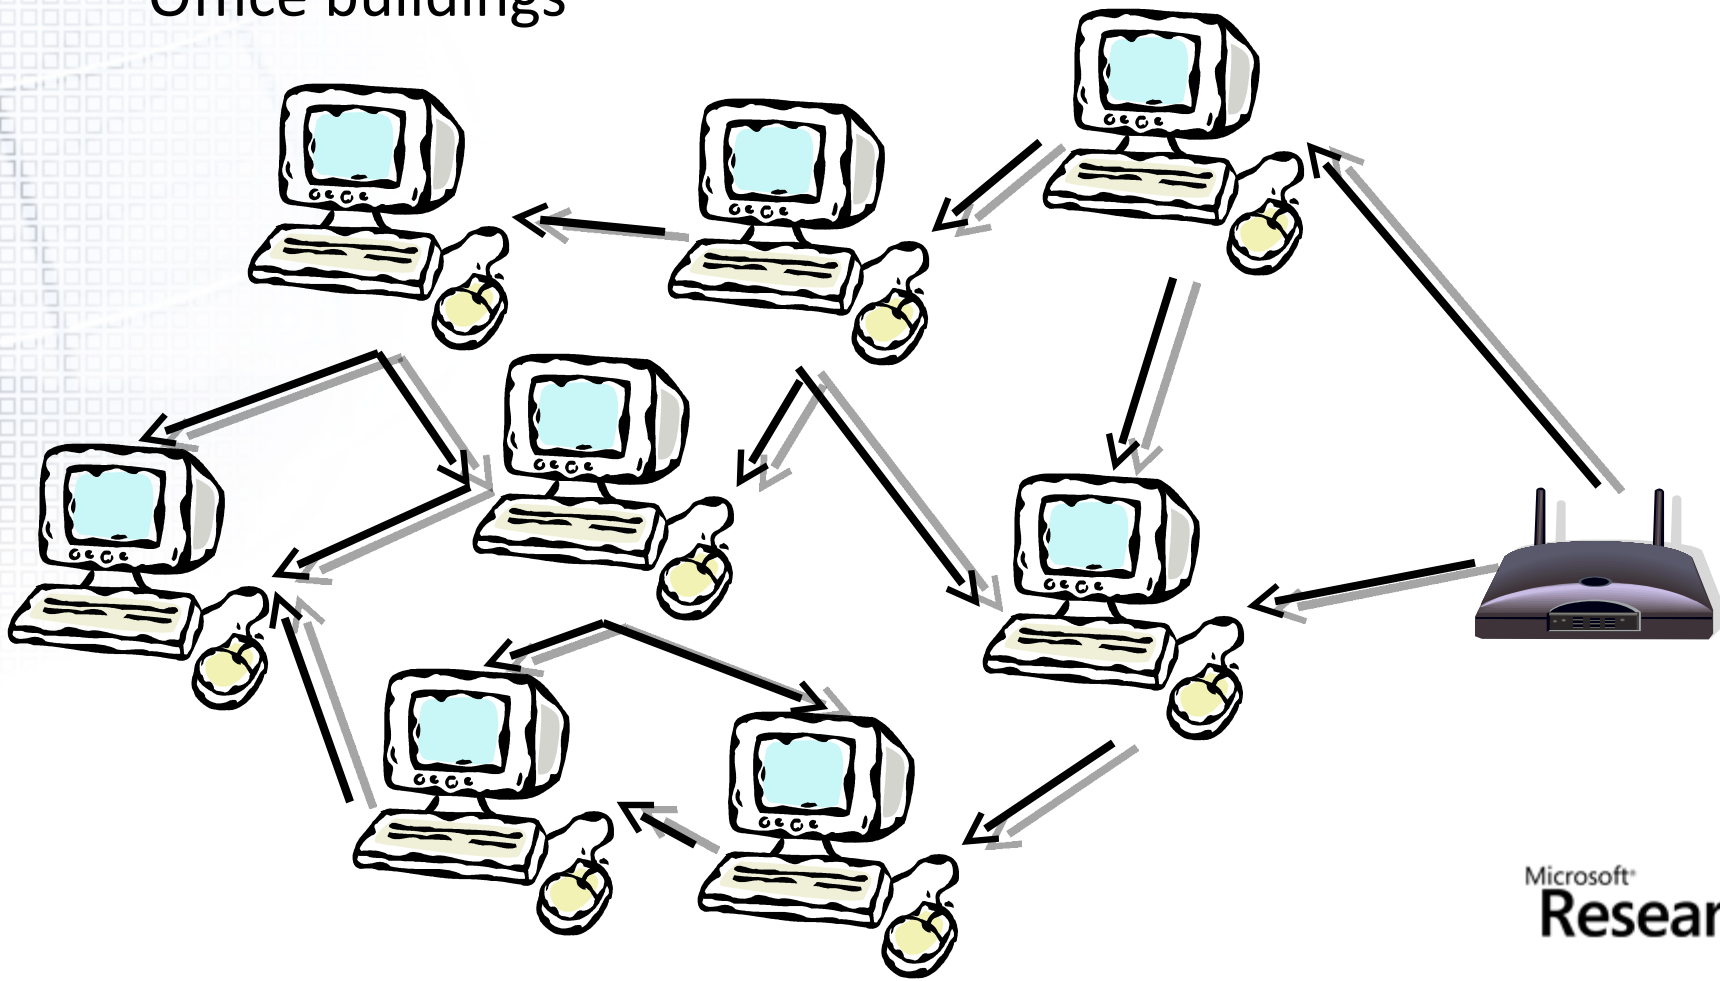


- Can we apply lessons to build routing protocols
  - Overlay routing at the network level



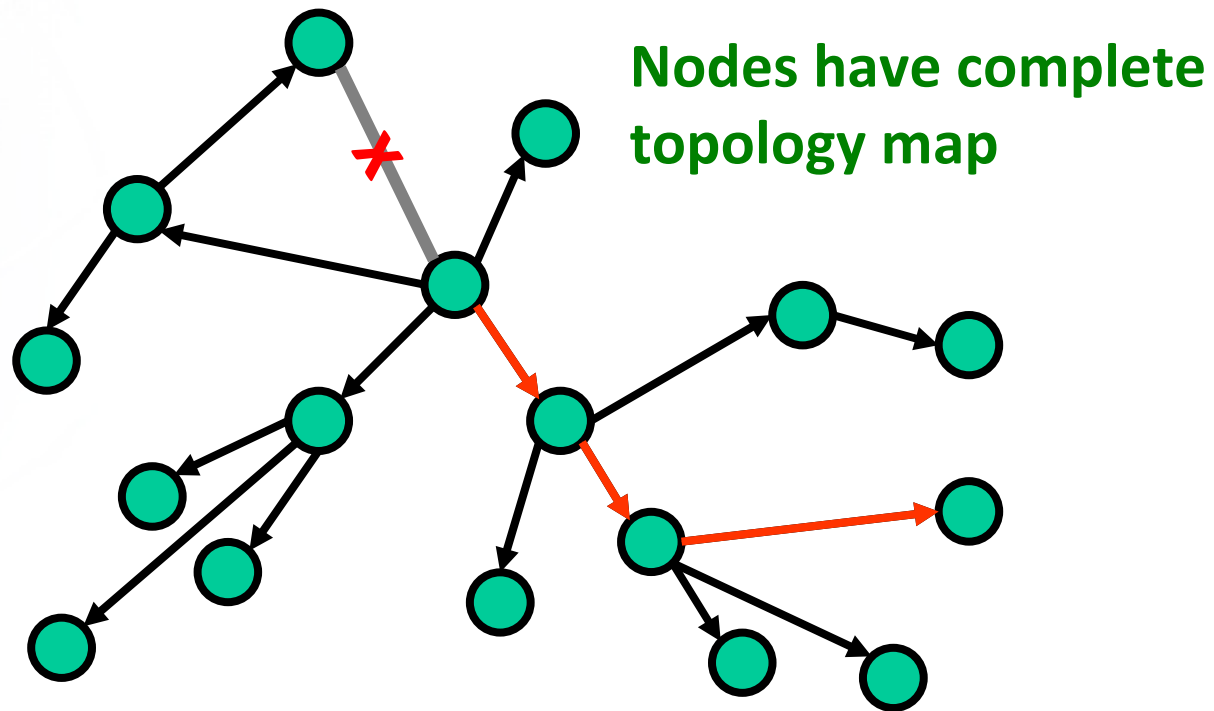
# Wireless routing

- Mesh networks beginning to be deployed:
  - Office buildings





# Proactive routing

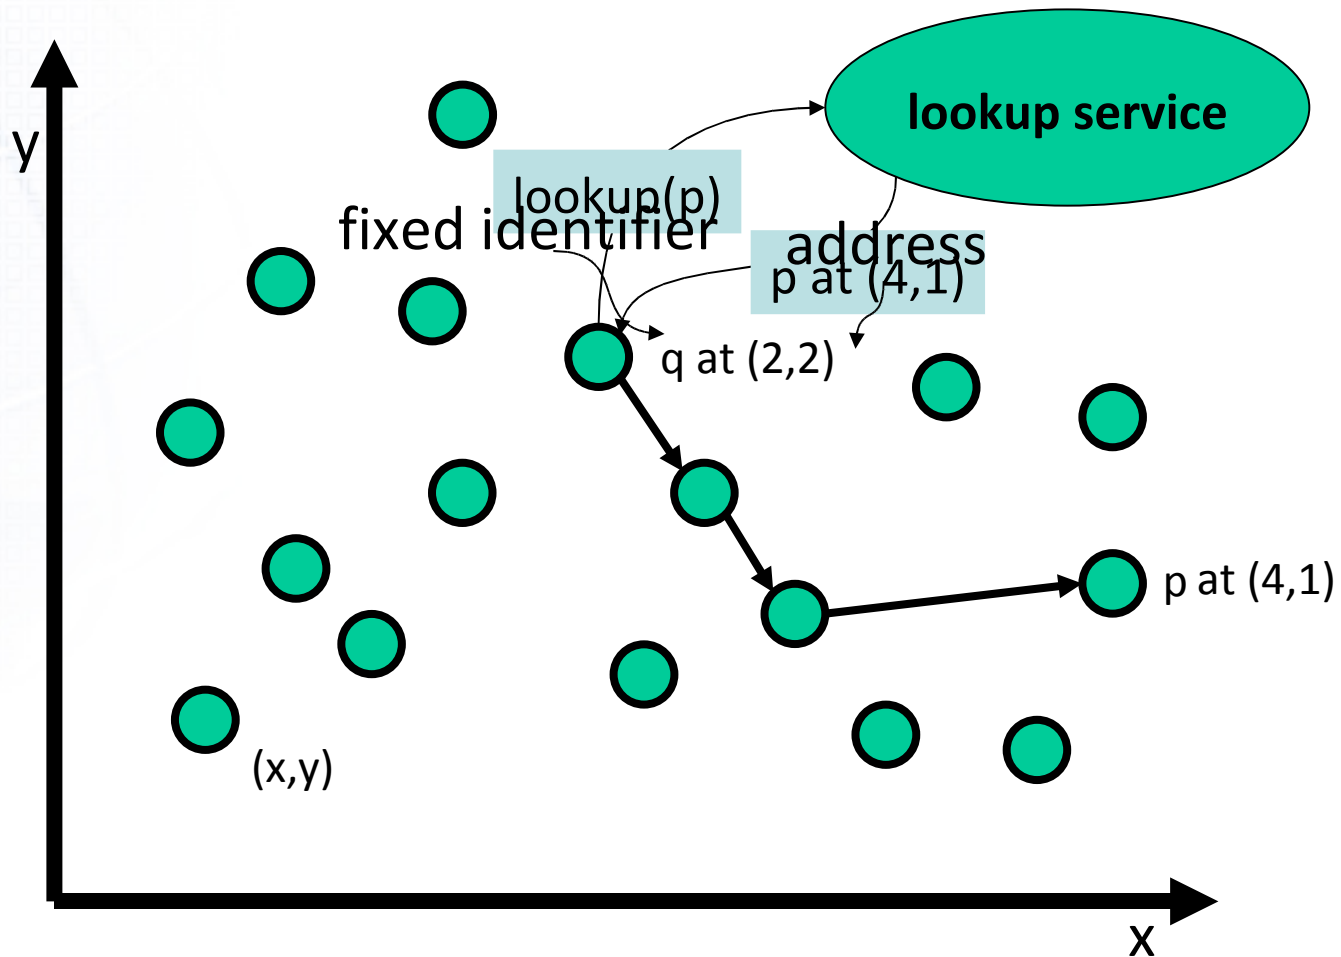


For example: OLSR and DSDV





# Coordinate-based routing



For example: GPSR and BVR



# Prior work on wireless routing

- Flooding based algorithms scale poorly
  - Proactive algorithms flood on topology changes
  - On demand algorithms flood to discover routes
- Geographical and landmark routing
  - Scale well but nodes have identifier and address
    - Some apps (sensor networks) may require just address
  - DHT-like structure to translate between the two
    - Route setup delays
    - Additional maintenance overhead
    - Another target for attacks



# Virtual Ring Routing

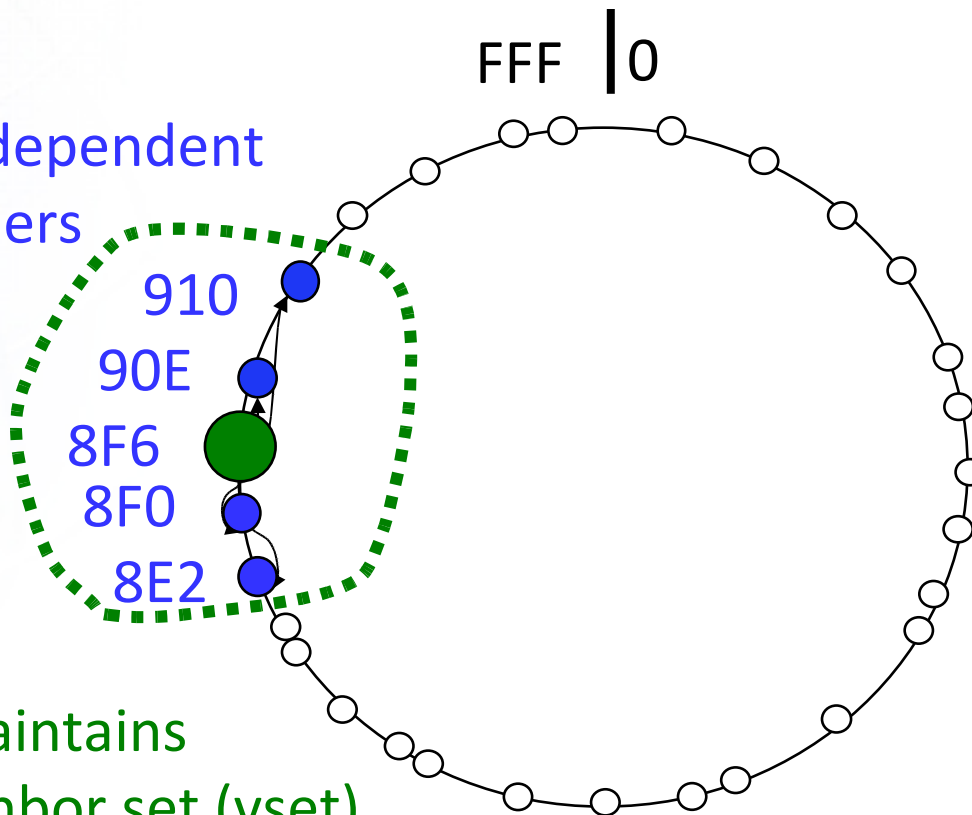
Joint work with Miguel Castro, Greg O'Shea plus interns

- Virtual Ring Routing – protocol inspired by DHT design
  - Single fixed identifier
  - No flooding
  - Provides DHT for free
    - Strict layering on wireless routing is inefficient
    - Poor interaction between DHT and routing layer



# VRR: The virtual ring

Topology-independent  
node identifiers




Each node maintains  
a virtual neighbor set (vset)

Nodes organized into virtual ring  
by increasing identifier value



# VRR: Routing paths



Nodes only maintain routing paths to virtual neighbors:

- Paths maintained proactively
- Paths are bidirectional and typically multi-hop




# VRR: Forwarding table

endpointA	endpointB	nextA	nextB	pathId
8F6	90E	me	F42	31
910	8F6	10E	me	10
14A	140	F42	10E	2
8F6	F42	me	F42	FF

forwarding table for node 8F6

- Paths recorded in forwarding tables along path
- Forwarding table contains
  - Paths between node and vset members
  - Paths between other nodes that go through node
  - Paths to physical neighbors







# VRR: Routing

- Forward message destined to  $x$  by
  - Picking endpoint  $e$  numerically closest to  $x$
  - Forwarding message to next hop towards  $e$
- Deliver message to node with id closest to  $x$





# VRR: Example routing



physical network topology



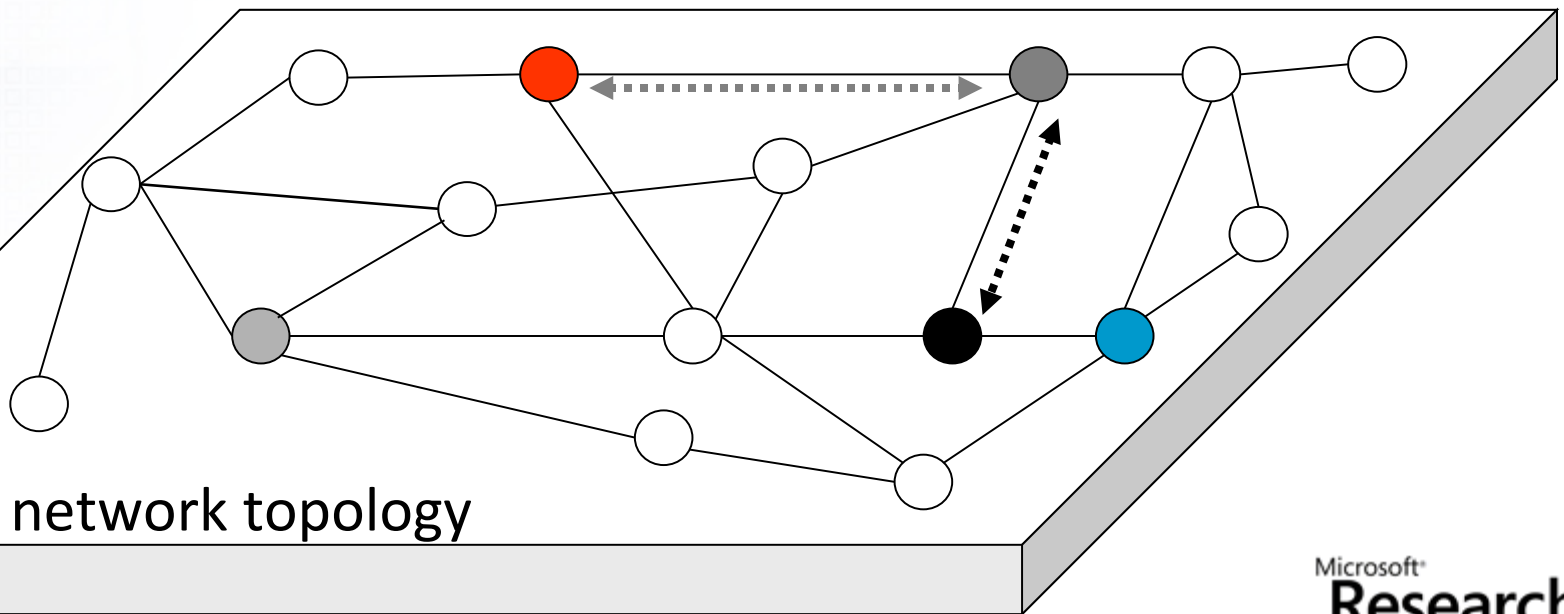

# VRR: Example routing



physical network topology

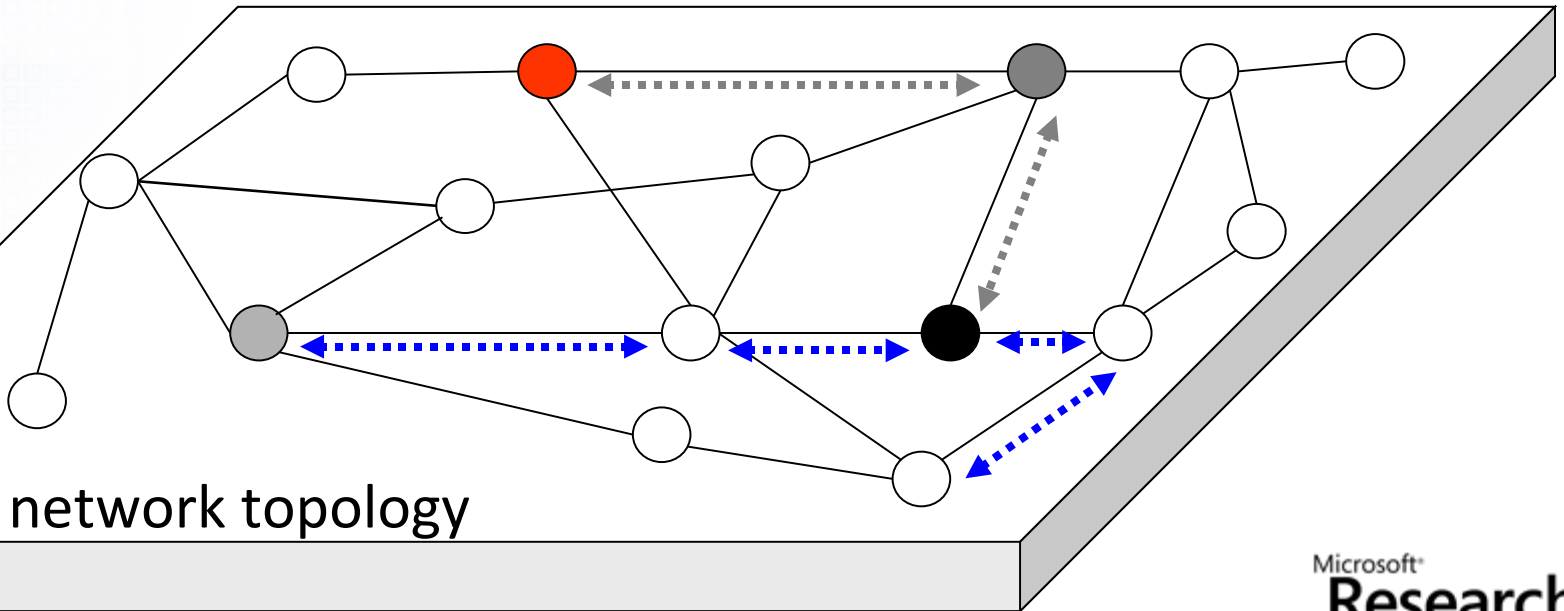



# VRR: Example routing



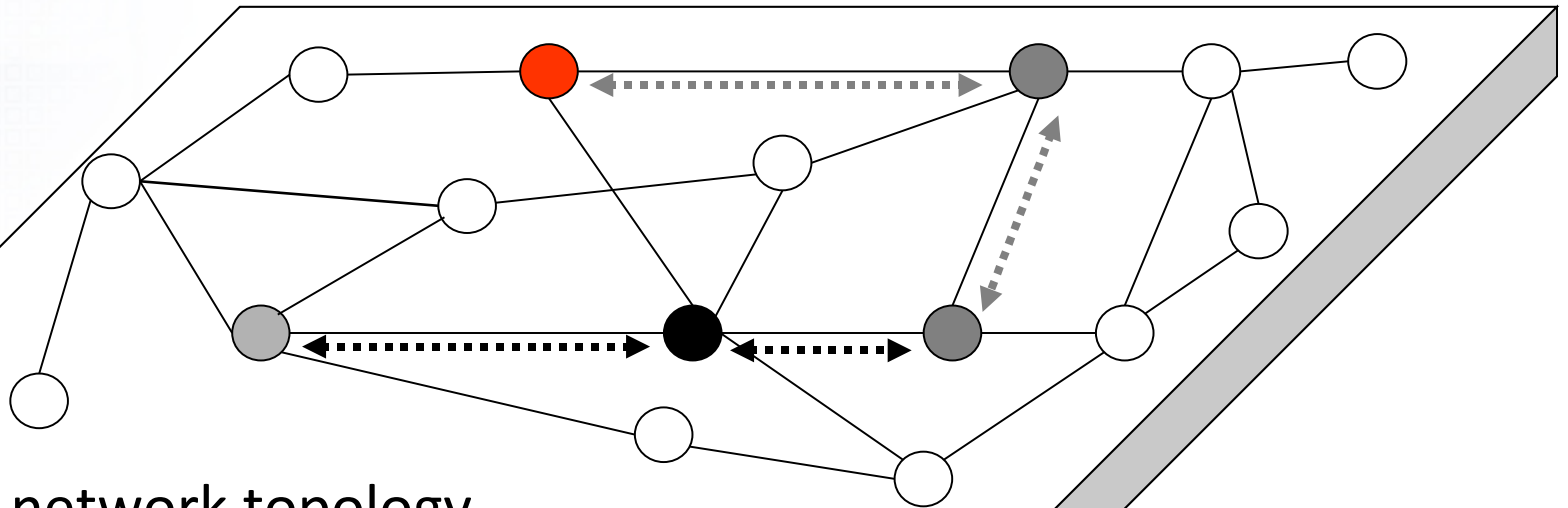



# VRR: Example routing





# VRR: Example routing



physical network topology



# VRR: Routing summary

---

- Paths to virtual neighbors ensure correctness
- Stretch empirically small
- Many alternate paths to route around failures




# VRR: Ring maintenance

---

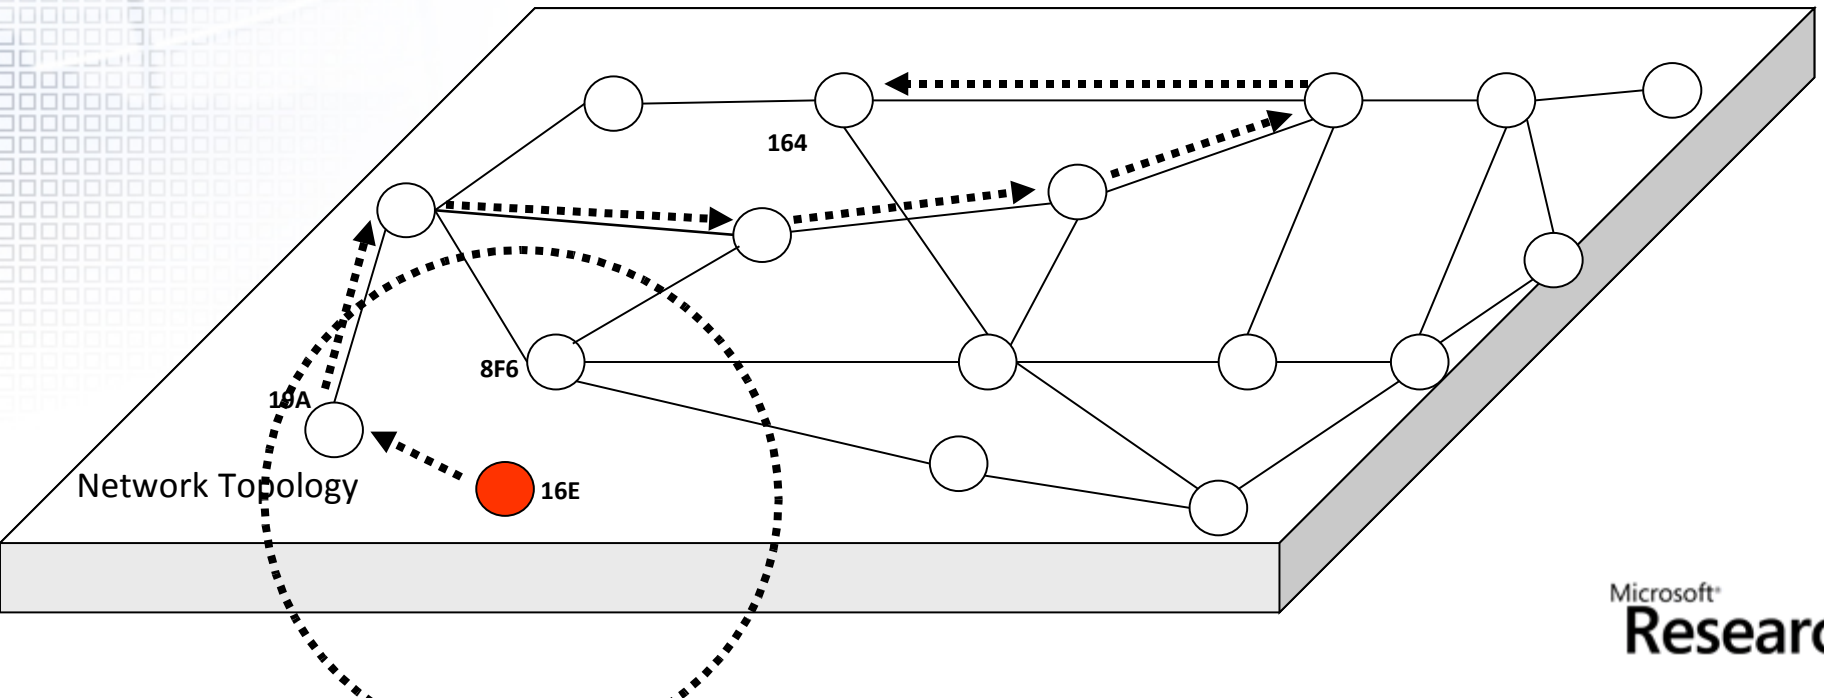
- No flooding
  - All messages routed as described
- Single topology independent identifier
- Five message types
  - hellos maintain physical neighbor set
  - setups update forwarding table state along path
  - setup requests ask another node to send setup
  - setup replies refusal to send setup
  - teardowns remove forwarding table state



# VRR: Node joining




broadcast hellos  
Send setup request to 16E  
to find physical neighbors

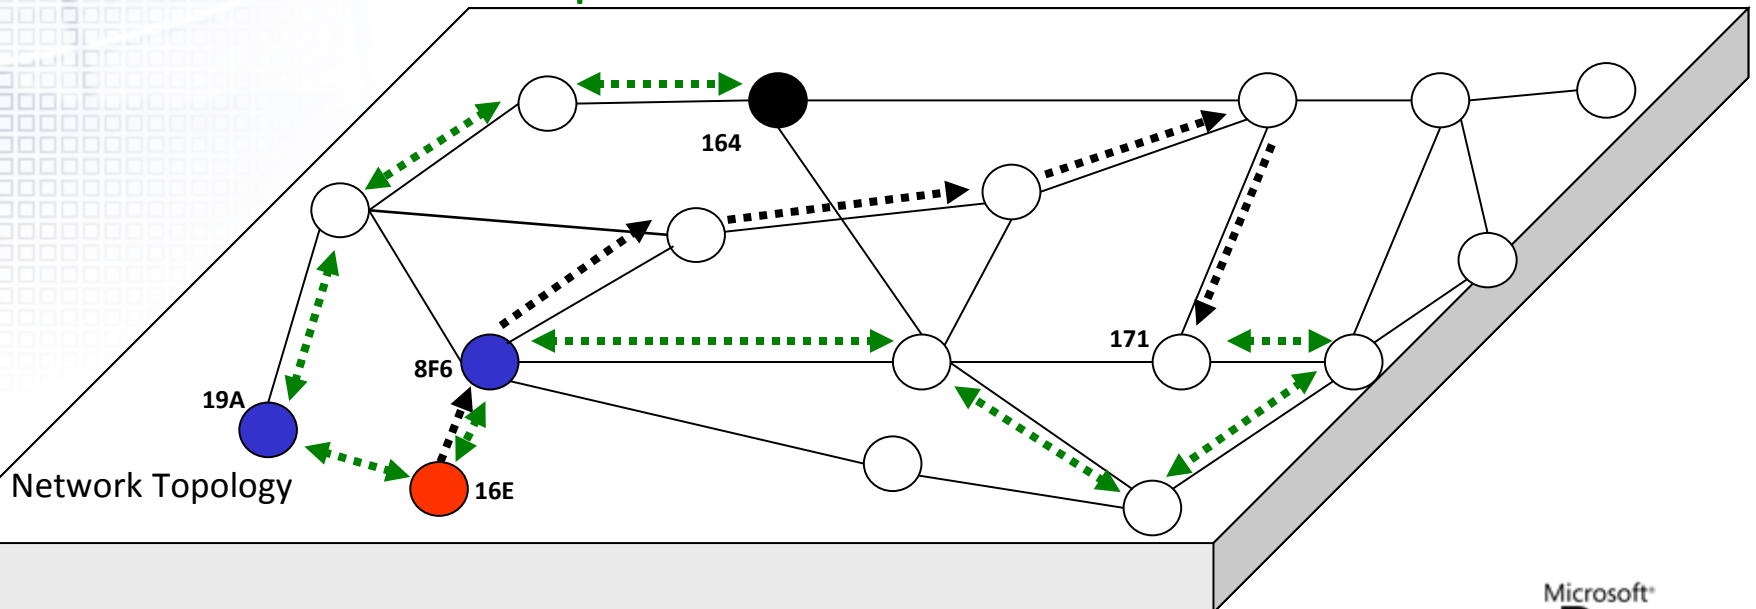




# VRR: Node joining



16E sends setup requests  
to nodes in received vset  
with node ID  
16E adds node to vset  
when it receives setup



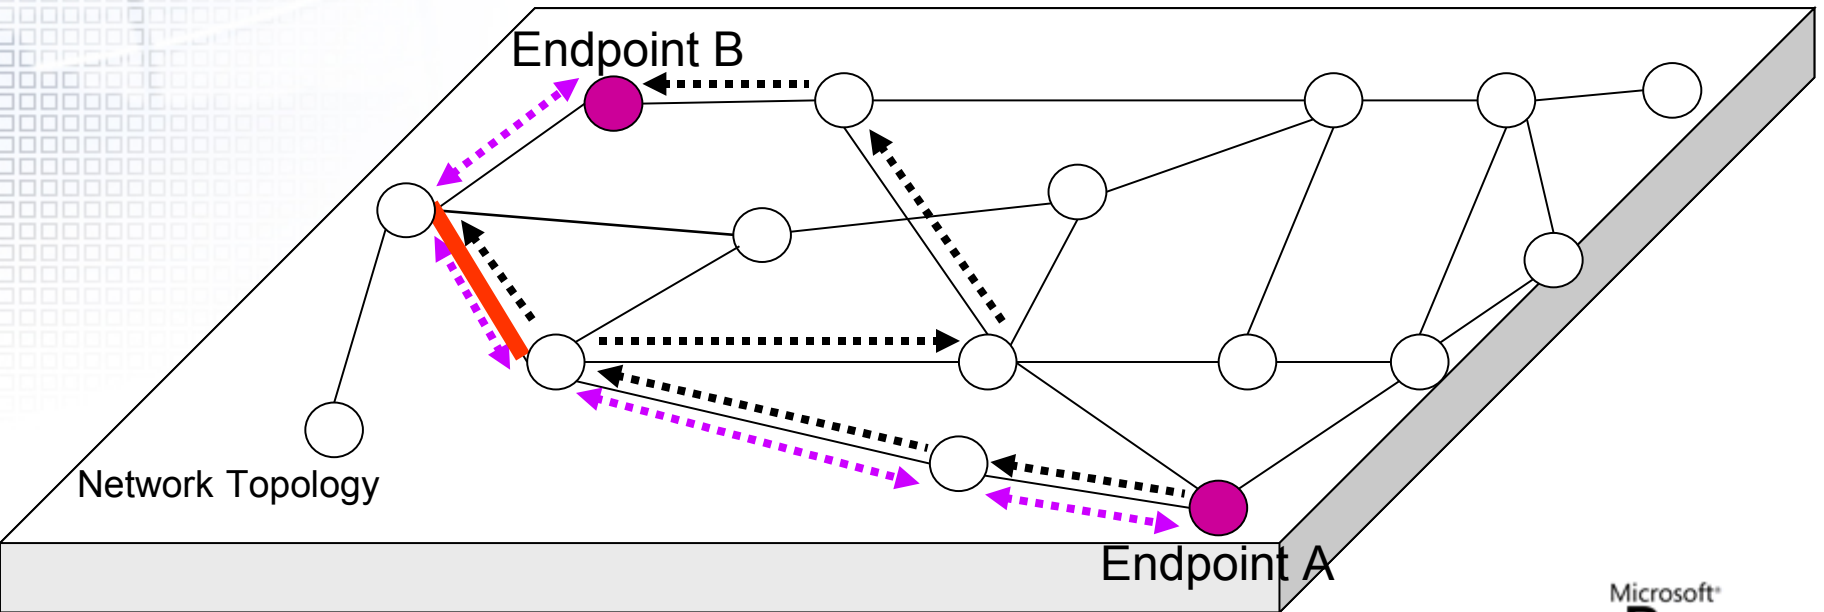



# VRR: Handling failures

- **Routing state is hard**
  - No end-to-end heartbeats
  - Failures detected on missing acks or hellos
  - Local repair attempted first;
  - Otherwise, teardowns sent along all affected paths
- **Two techniques to ensure consistency**
  - Symmetric failure detection and acks on teardowns
    - If  $x$  marks  $y$  faulty,  $y$  is guaranteed to mark  $x$  faulty
  - Lightweight optimistic transactions
    - If in doubt abort (teardown)

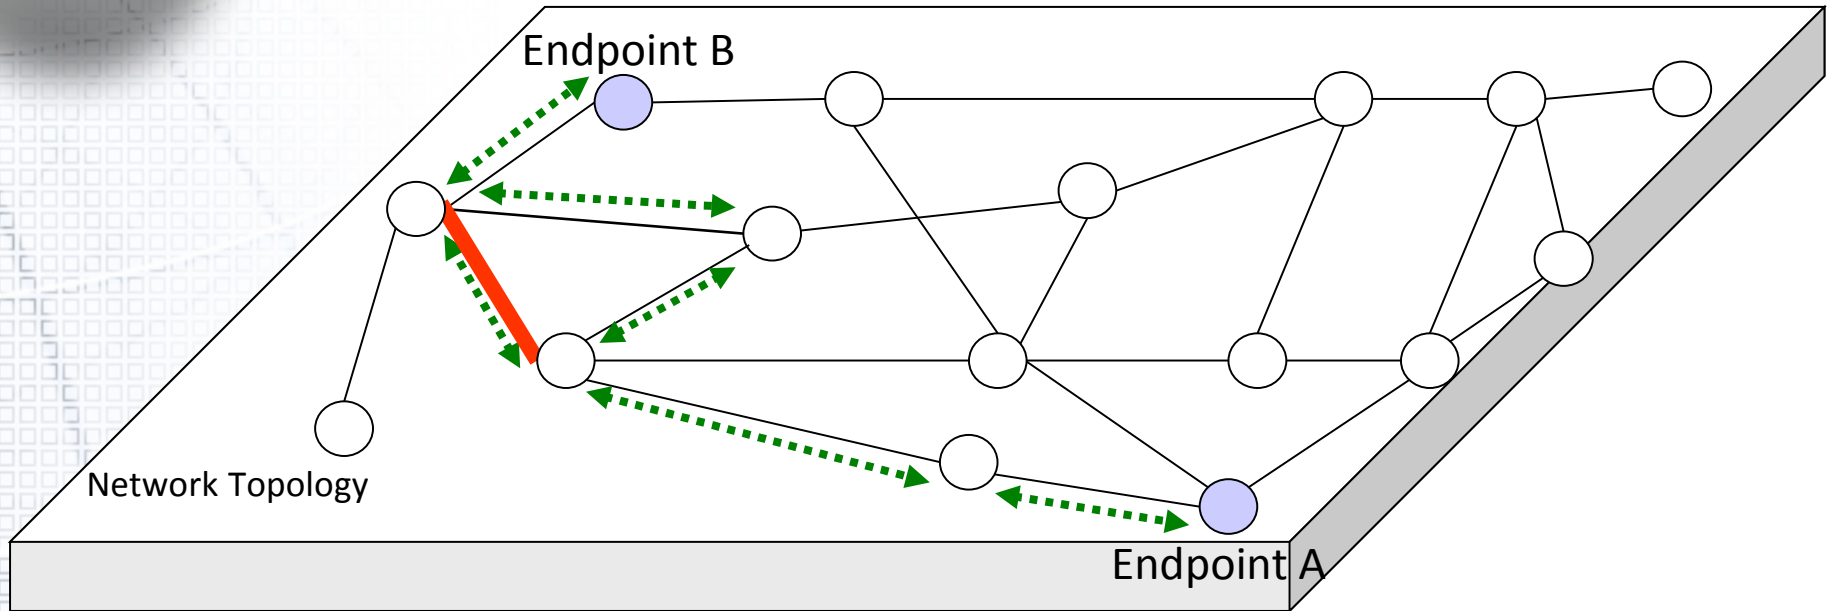


# Routing during node or link failure





# A link failure example



- **Repair is truly local**
  - Only involves nodes near failed link or node
  - No end-to-end path metrics
- Repair **aborted** if local consistency checks fail




# Simulation experiments in ns-2

- Experiments with 802.11b MAC
- Varied network size, mobility, session lifetime
- Compared with DSDV, DSR, and AODV
  
- VRR performed well in all experiments
  - high delivery ratios even with fast movement
  - significantly lower delays with route instability




# Delivery ratio: fast movement





# Delay: fast movement






# Sensor network

---

- Sensor network testbed
  - 67 mica2dot motes in UCB building
- Comparison with BVR (coordinate-based protocol)
- Delivery ratio with mote failures




# Sensor network: mote failures






# Wireless office testbed

- 30 machines running windows
- Communicate using 802.11a
- Throughput comparison with LQSR using tcp





# Wireless office testbed: throughput






# Virtual Ring Routing Summary

- Routing protocol inspired by structured overlays
- Unique (new point in the design space):
  - Single identifier per node
  - No flooding
- Provides DHT for free
- For more information see:
  - M. Caesar, M. Castro, E. Nightingale, G. O'Shea and A. Rowstron, "**Virtual Ring Routing: Network routing inspired by DHTs**", Sigcomm 2006, Pisa, Italy, September 2006.




# Vehicular networking

- New challenges for scalability
  - Very different characteristics






# Vehicle-to-Infrastructure (V2I)





# Vehicle-to-Vehicle (V2V)



- Car manufacturers interested and “driving”:
  - 802.11p (Dedicated short range communications)
  - Intelligent Transport Systems
    - “What is the car in front doing?”
- But then.....



# The disruptive device.....

- SatNav (TomTom)
  - Dual connectivity model:
    - Bluetooth to mobile phone
    - USB-to-computer (new map > 1Gb)
  - Download data for premium services (Home service)
    - TomTom Map Share (Web 2.0 app)
    - TomTom QuickFix (Assisted GPS)
    - TomTom Updates
  - “Connectivity key”





# Personal Navigation Devices

“Hand held” or In-dash!






# PVRP: Practical Vehicular Routing Protocol


## Joint work with G. Pau and P. Lutterotti (UCLA)

- Multi-hop vehicle-to-infrastructure
  - Route to fixed access points
- Multi-hop vehicle-to-vehicle
  - Route to specific vehicle or a service
  
- Why do we need another routing protocol?


# Greedy Routing



# Discovery: Dense



# Discovery: Sparse




Sparse more common than dense  $\rightarrow$  "delay" tolerant protocol



# The PVRP approach

- Assume nodes have:
  - Digital maps (e.g. NavTech digital maps)
  - 802.11a/b/g WiFi (or equivalent)
  - GPS system
  
- Perform routing and discovery in map space not physical topology
  - Opposite to VRR

# Assume nodes maintain one-hop topology information




ID	Location	Service
	12,32	A,B,F,G
	34,23	Q,R,F, S
	18,38	A,B
	45,34	A,B,G




# Discovery: Exploit map

- Avoid pure flood in physical topology
- Exploit map to ensure good exploration





# Map-based source routing



Packet: <Src> <Dest> <etc>

<Path:

Victoria Road -> Garden Walk Go W  
Victoria Road -> Harvey Avenue Go W  
Victoria Road -> Frenchs Road Go SW  
>

Packet: <Src> <Dest> <etc>

<Path:

Victoria Road -> Garden Walk Go W  
Victoria Road -> Harvey Avenue Go W  
Victoria Road -> Frenchs Road Go N  
>

**Note:** Between any two Junctions the packet may traverse several network hops



# Preliminary results

- Currently understanding the design choices and their impact:
  - Path selection metrics:
    - Average density,
    - Max lowest density, etc
  - Junction selection metrics:
    - Select the important junctions
    - Comparing against delay tolerant greedy algorithm

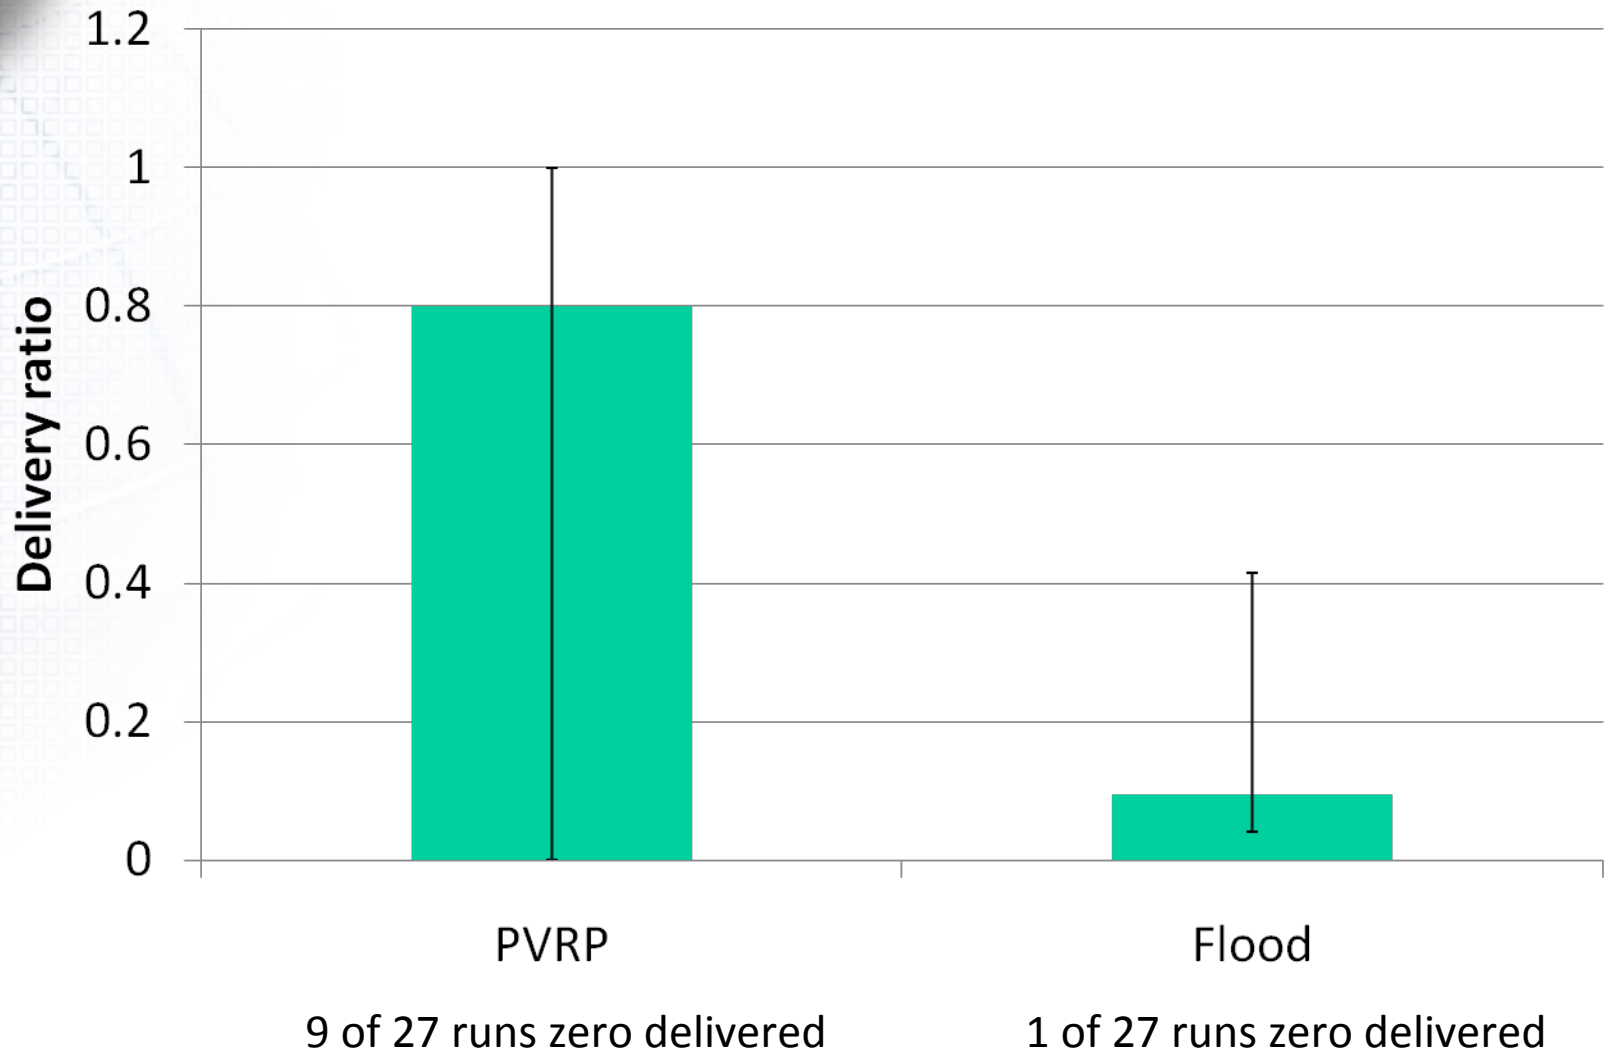


# Preliminary results

- Currently evaluating PVRP using:
  - “Simple” simulator (versus Flood)
  - QualNet simulator (versus AODV/DSR/GPSR)
  - Stand alone implementation (keeping us honest!)
- Using mobility traces for Portland
- Results today generated using simple simulator
  - Comparing PVRP with “Flood”
    - Flood represents best that AODV/DSR/GPSR could do
  - Results using “realistic” mobility trace
    - 59 runs randomly selected static end-points with distance between 250m and 750m
    - Across all runs only 27 experiments delivered packets

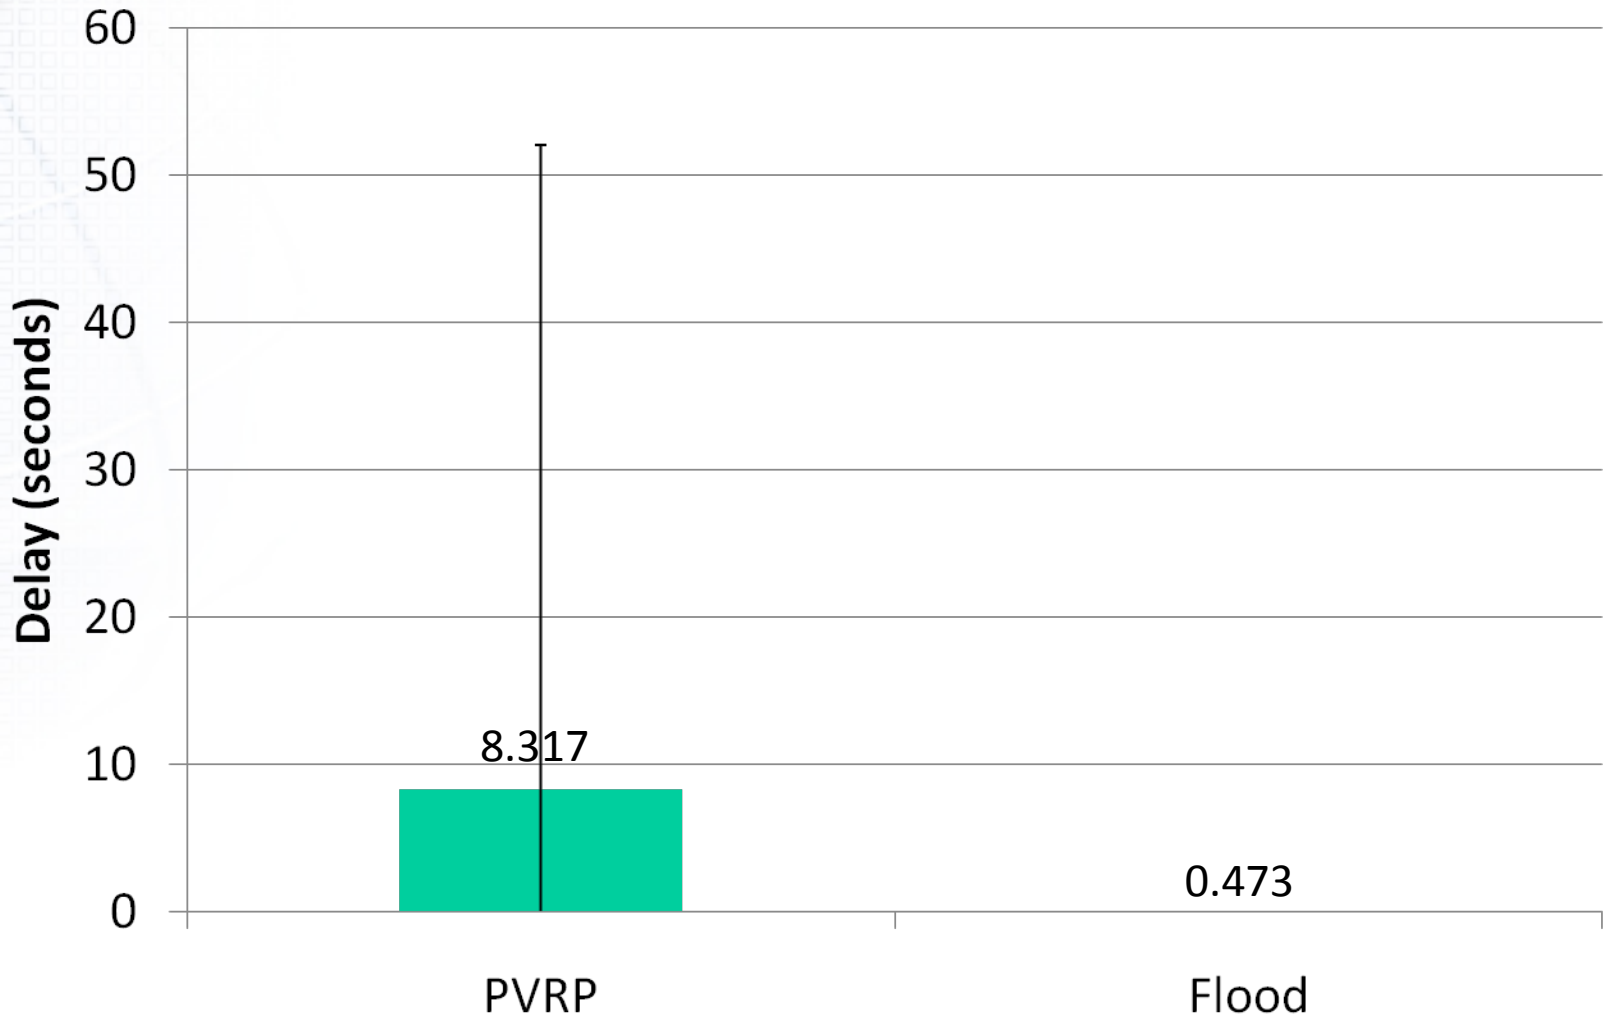


# Delivery ratio: Preliminary results





# Delay: Preliminary results





# PVRP Summary

---

- Practical multi-hop routing protocol for vehicular networks
- Exploits digital maps rather than just the physical network topology
  - To work over partitioned networks



# Thank you

---

<http://research.microsoft.com/~antr>