NRL Cross-Layer Workshop

Interlayer routing issues for wireless networks

June 2, 2004

Tom Henderson
Marcelo Albuquerque
Phil Spagnolo
Jae H. Kim

Boeing Phantom Works

1

	Form Approved OMB No. 0704-0188

Report Documentation Page

1. REPORT DATE **01 DEC 2007**	2. REPORT TYPE **N/A**	3. DATES COVERED

4. TITLE AND SUBTITLE **Interlayer routing issues for wireless networks**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Boeing**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT **UU**	18. NUMBER OF PAGES **35**	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**			

Motivation
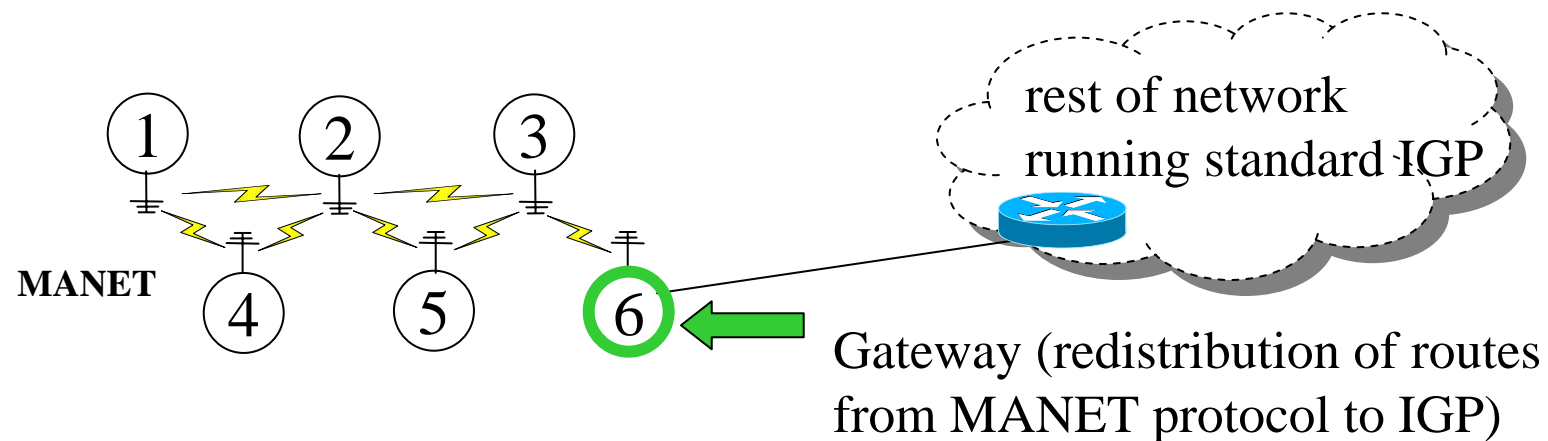
- The operation of conventional packet routing protocols over wireless networks has been repeatedly demonstrated to be suboptimal

   1. Reliability mechanisms in protocols not tuned to loss-prone environment

   2. Inefficient usage of transmissions (e.g., flooding)

   3. Underlying channel can be time-varying

Response to this problem: <u>Mobile Ad Hoc Routing Protocols</u>

MANETs and larger networks

How to interconnect the packet routing?

1 2 3

MANET

4 5 6

rest of network
running standard IGP

Gateway (redistribution of routes
from MANET protocol to IGP)

This is relatively easy if the MANET is a stub network

MANETs and larger networks
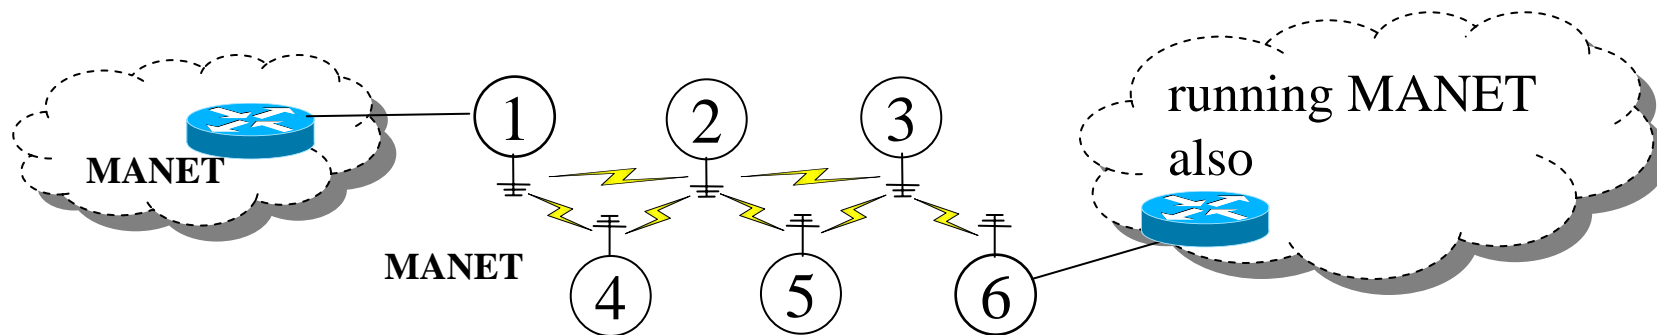

What if MANET provides transit routing?

MANET

rest of network
running standard IGP

- route redistribution between protocols is cumbersome to configure and manage
- redistribution can be lossy
- using multiple ASes interconnected by BGP is not attractive either

MANETs and larger networks

Why not run MANET as the system IGP?

MANET

running MANET also

MANET

① ② ③

④ ⑤ ⑥

MANET protocols not as mature for operation in heterogeneous networks

– Not optimized for supporting heterogeneous subnet technologies in transit configurations

Note: this is not to say that MANET protocols couldn't evolve to be a full-fledged IGP

Alternatives to redistribution

Closer interaction with IGP routing protocol

- **Solution 1:** Run a MANET protocol as a "subnet" routing protocol, and overlay the IGP

- **Solution 2:** Modify the IGP to perform more like a MANET

- **Solution 3:** Solution 1, but with "cross-layer" integration

Interestingly, all three approaches are currently being developed by standards bodies (IEEE, IETF) and the government (JTRS)!

"Layers" and routing

- Hop-by-hop, shortest path routing is typically done at "layer 3" (internetworking layer) in the Internet

- Routing or bridging is also done at the subnet layer in many cases
  - e.g., Ethernet spanning tree
  - we will call this "layer 2" routing in this presentation

- Note: ISO terminology sometimes refers to this as "layer 3c" (subnet independent convergence) vs. "layer 3a" (subnet access) routing

Problem statement

What are the qualitative and relative performance tradeoffs between these routing architectures?

- Layer-2 and Layer-3 (operating independently)
  - run unmodified layer-3 protocol on a layer-2 MANET protocol
- Only Layer-3
  - modified layer-3 protocol that is MANET capable
- Layer-2 and Layer-3 (cross-layer)
  - run (modified?) layer-3 protocol on a layer-2 MANET protocol, with "cross-layer" interaction

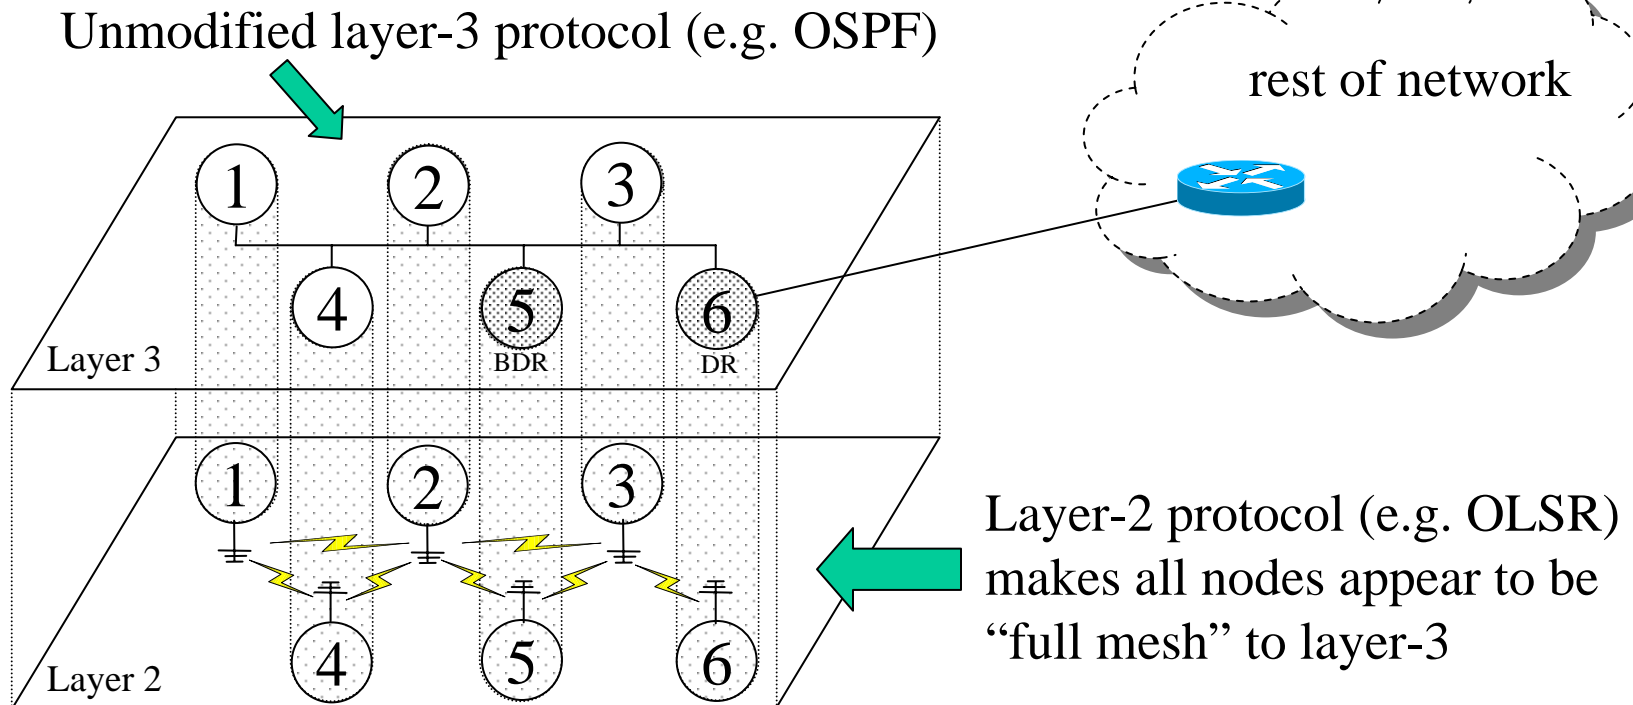# Scope of this presentation

- Focus on best-effort, hop-by-hop routing
  - no direct focus on energy-efficiency
  - no QoS or delay issues

- Limited amount of cross-layer interactions examined
  - not aggressive feedback of physical layer SNR, for example

- Routing technique studied: Link-state (OSPF and OLSR)

- Channel/MAC: 802.11b

1. Layer-3/Layer-2 routing

- ## Motivation:
  - let a specialized layer-2 protocol operate in the wireless subnet
  - layer-2 protocol builds and maintains full-mesh connectivity between nodes
  - operate layer-3 protocol with a broadcast-based interface
  - simple example: Ethernet spanning-tree bridging, with OSPFv2 (broadcast interface) on top

- ## Related approach in development:
  - IEEE 802.11 Extended Service Set (ESS) Mesh

1. Layer-3/Layer-2 routing

Unmodified layer-3 protocol (e.g. OSPF)

rest of network

Layer 3

BDR   DR

Layer 2

Layer-2 protocol (e.g. OLSR) makes all nodes appear to be "full mesh" to layer-3

Note: OLSR could be operated as a layer-3 protocol as well-- here, we just use it as layer-2
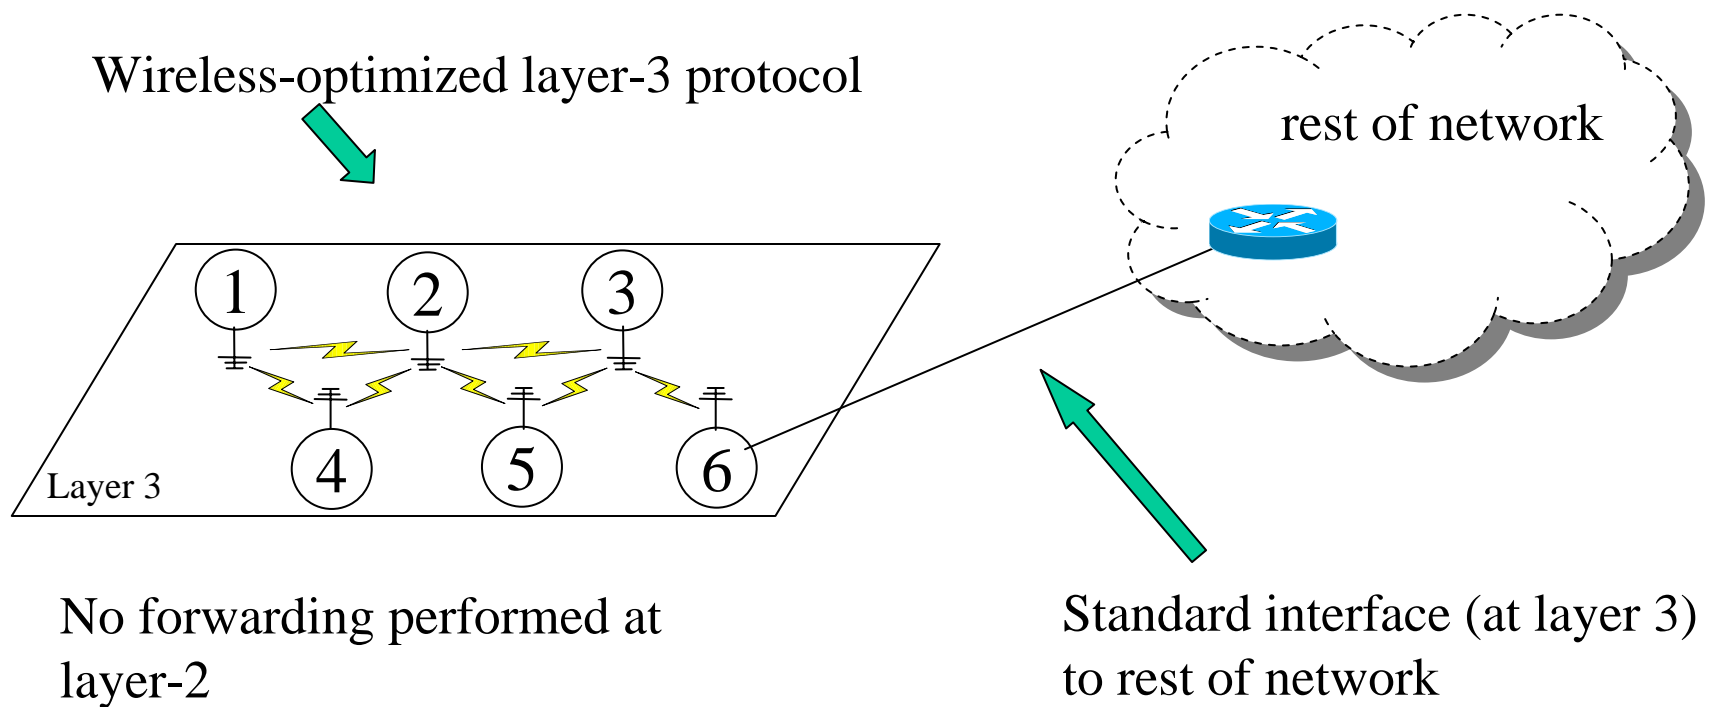
1. Pros and cons

- ## Advantages
  - no modifications to layer-3 protocol
  - layer-2 protocol can be tailored to the specific subnet radio technology
  - IP broadcast address handled naturally
  - wireless topology disruptions can be hidden from layer-3 (and rest of network!)
- ## Disadvantages
  - must operate two routing protocols on top of each other (operationally more complex)
  - generates redundant overhead (e.g., neighbor discovery)

2. Layer-3 routing only

- ## Motivation:
  - Modify layer-3 IGP routing protocol to operate more efficiently in wireless environment
  - Note: most MANET protocols are not full-fledged IGPs, designed for operation across diverse subnets

- ## Related approaches in development:
  - Boeing/INRIA "wireless" OSPF interface type
    - draft-spagnolo-manet-ospf-wireless-interface-01.txt
  - Cisco MANET OSPF
    - draft-chandra-ospf-manet-ext-00.txt
  - IETF OSPF WG forming a design team to work on this topic

2. Layer-3 routing only

Wireless-optimized layer-3 protocol

rest of network

1 2 3

4 5 6

Layer 3

No forwarding performed at
layer-2

Standard interface (at layer 3)
to rest of network
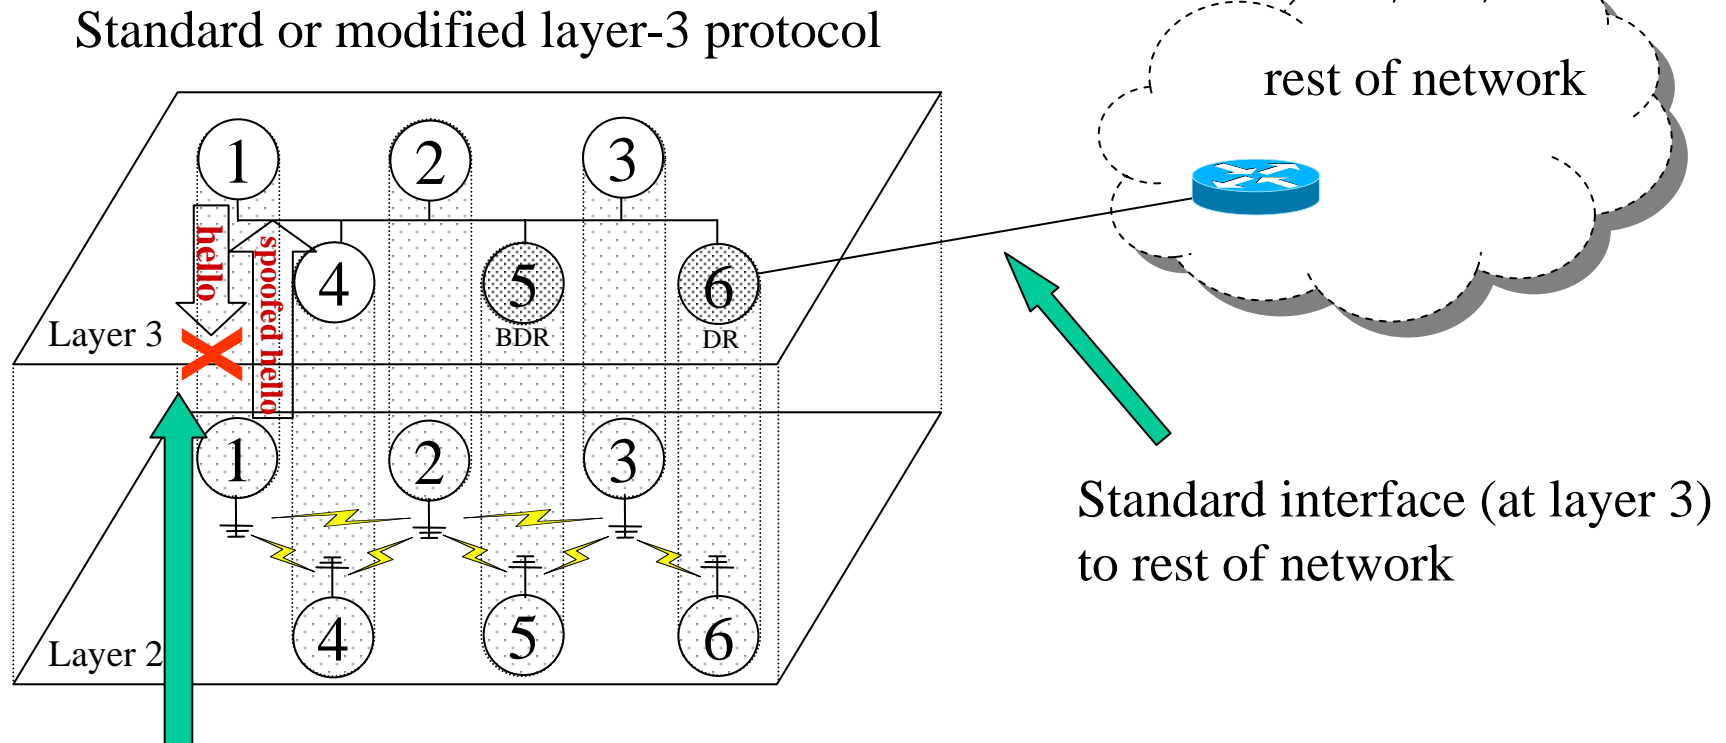
2. Pros and cons

- Advantages
  - eliminate redundant overhead (e.g. neighbor discovery)
  - reduces risk to instability if the layer-2 protocol does not converge fast enough to reliably deliver a full-mesh topology to layer-3
  - multicast routing protocols may operate without resorting to a mapping to an underlying broadcast mechanism
  - exposes the inner topology so that external routers can find good entry points into the network (good for avoiding routing stretch)

- Disadvantages
  - may not be optimized for a given subnet radio technology
  - exposes topology changes to rest of network (bad for outside network)

3. Integrated layer-3/layer-2

- Motivation:
  - Obtain efficient operation by sharing information between layers
  - Avoid negative interactions between the protocols
  - Design to achieve best of both worlds
- Examples
  - proprietary radios

- This approach also known as "cross-layer"

3. Integrated layer-3/layer-2

Standard or modified layer-3 protocol

rest of network

Standard interface (at layer 3) to rest of network

- Layer-3 neighbor discovery can be suppressed
- Layer-2 triggers may improve layer-3 responsiveness
- Physical layer metrics (link performance) could affect routing choices
- ...

17

3. Pros and cons

- ## Advantages
  - can potentially obtain benefits of both approaches

- ## Disadvantages
  - not currently as likely to be a standard approach (however, an API might be specified)
  - perhaps not as general of a solution

Methodology

- Discrete-event networks simulation (based on QualNet 3.5 models)
  - 2.4 GHz, 802.11b radios
  - OSPFv2 (point-to-multipoint, and broadcast)
  - Boeing/INRIA "wireless" OSPF (defined later)
  - Optimized Link State Routing (OLSR) version 7
  - random waypoint mobility
- Key performance statistics:
  - Overhead: measured at the IP layer
  - Packet delivery ratio: delivery of user's UDP data in the network (not measuring routing stretch)
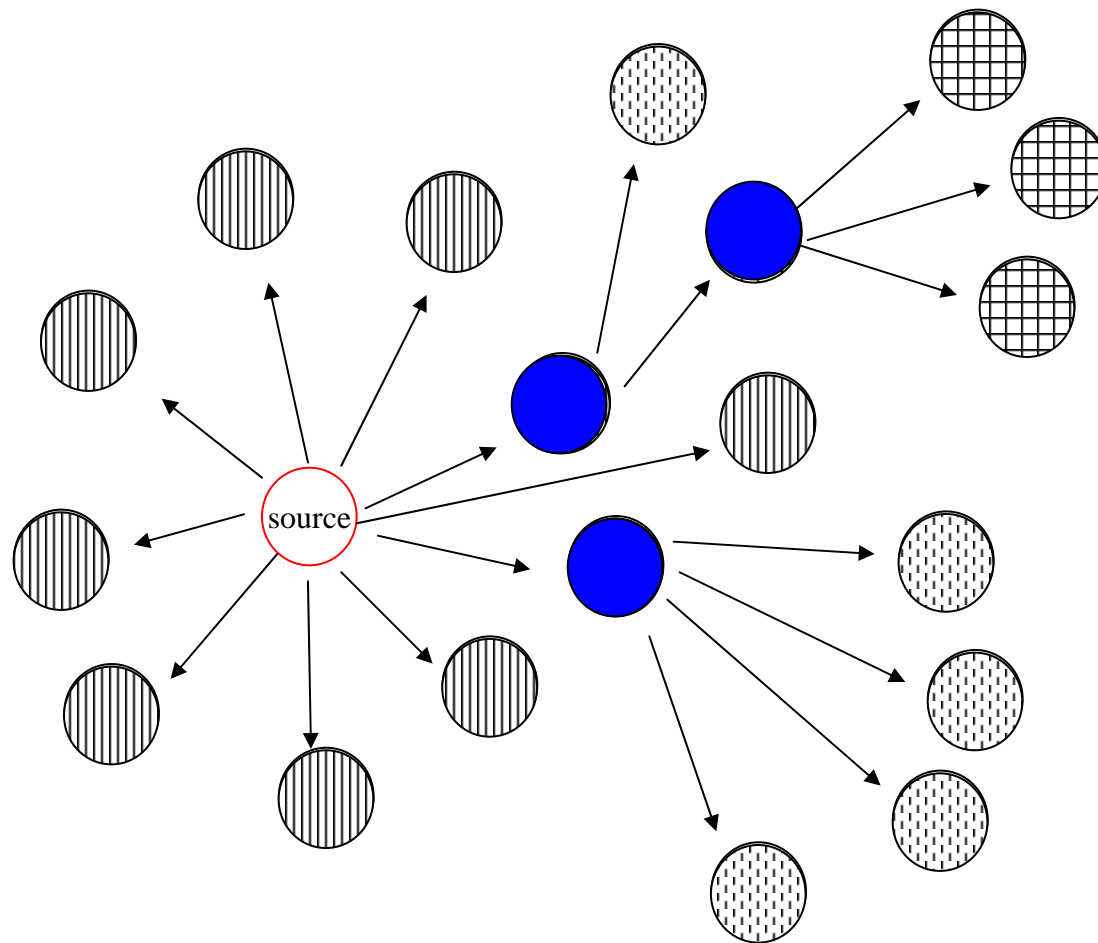
OSPFv2 description

- Link-state (proactive) routing protocol for IPv4

- <u>Point-to-multipoint</u> interface type
  - Forms routing adjacency with each other router
  - Can operate with layer-2 full-mesh or partial-mesh

- <u>Broadcast</u> interface type
  - Reduces adjacencies in network from $O(N^2)$ to $O(N)$
  - Requires a full-mesh connectivity between nodes

OLSR details

- Optimized Link State Routing
  - developed by INRIA, France
  - categorized as "proactive" protocol
- Most like OSPF
  - Shortest Path First (SPF)-based algorithm
  - Unreliable flooding algorithm
- <u>Strategy for scaling:</u> Subset of neighbors (Multi Point Relays) responsible for routing information dissemination

OLSR example

- OLSR implements a heuristic, distributed solution to the minimum connected dominating set problem

- Only blue nodes are responsible for reflooding messages from source

- In standard OSPF flooding, all nodes reflood received messages, leading to lots of redundant transmissions and interference

Wireless interface specifics

Wireless OSPF-- brings OLSR concepts to OSPFv2

- Main change is <u>topology dissemination</u>
  - OSPFv2 uses reliable flooding and database exchanges
  - wireless OSPFv2 uses periodic, optimized flooding, and no exchanges
- <u>Periodic flooding:</u> flooding is unreliable (does not need ACKed)
- <u>Optimized flooding:</u> uses concept of OLSR Multi Point Relays (MPRs) for efficient flooding

Wireless interface specifics

- LSAs and route computations are unchanged
    - every node has "full" adjacency with every other node, without the DB exchange
    - difference from OLSR: every router originates an LSA (full link state)
- New message type Link State Flood (LSF) replaces Link State Update (LSU)
- Extensions to Hello message to select MPRs for flooding

- No changes to multicast OSPF (M-OSPF)

Interaction with legacy

- Hybrid router can contain wireless and other traditional interfaces
  - from outside, wireless subnet looks like a Point-to-Multipoint subnet
- Outside LSAs can be flooded into the wireless domain
  - we developed some heuristics for how to efficiently do this
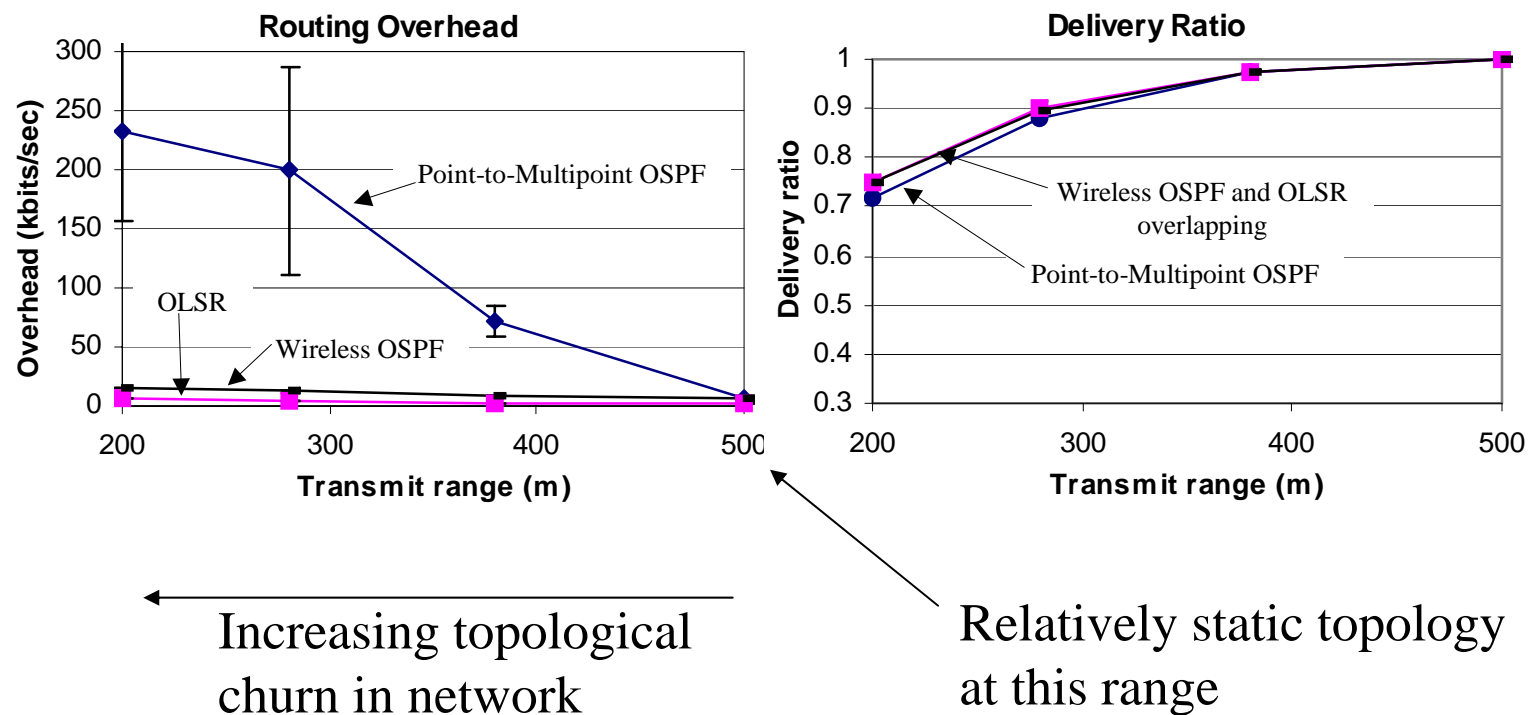 - more work is probably needed here

Protocols used

- Layer-3 (separate): Wireless OSPFv2
- Layer-2/Layer-3 (independent): OSPFv2 (broadcast) at layer 3, running over OLSR at layer 2
  - OSPF broadcasts were encapsulated and flooded as OLSR broadcasts
- Layer-2/Layer-3 (cross-layer): Same as above, but Hello messages suppressed from Layer-3
  - spoofed from local node's OLSR routing table
  - did not implement other cross-layer notifications

Background (20 nodes)

- Previous simulation results (have been validated with implementation testing as well):

Increasing topological churn in network

Relatively static topology at this range

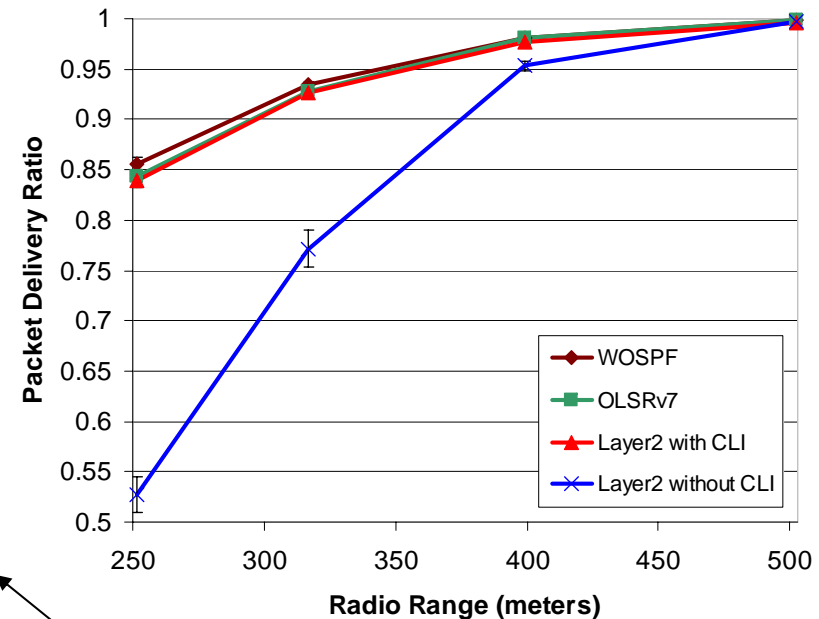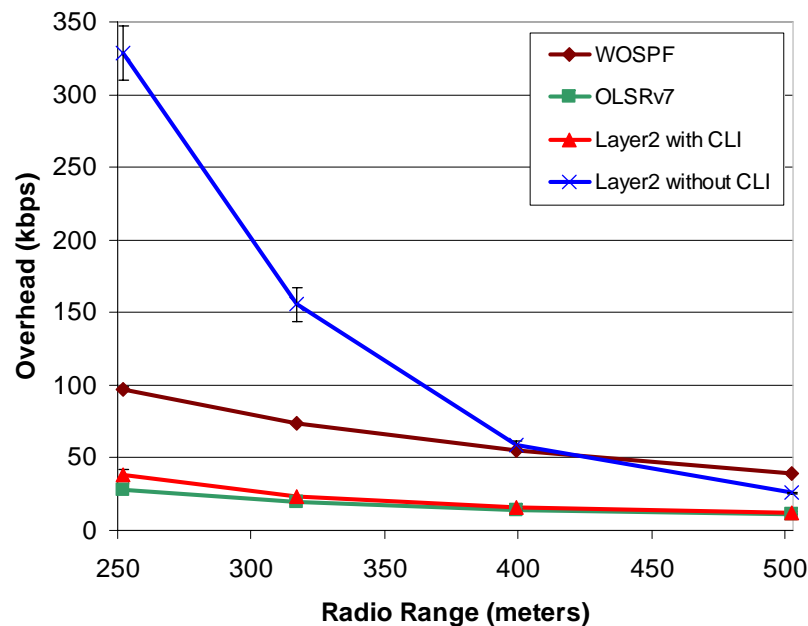* Henderson, Spagnolo, and Kim, "A Wireless Interface Type for OSPF," IEEE MILCOM 2003 Proceedings.

Background (performance)

- OLSR and wireless OSPF exhibit good scaling properties as the network topology changes more frequently

- OLSR outperforms OSPF because OLSR only uses MPRs to source LSAs

- OSPF in point-to-multipoint mode has very bad scaling properties as mobility increases

- Main observation: Layer2 with cross-layer integration (red triangle line) provides superior performance
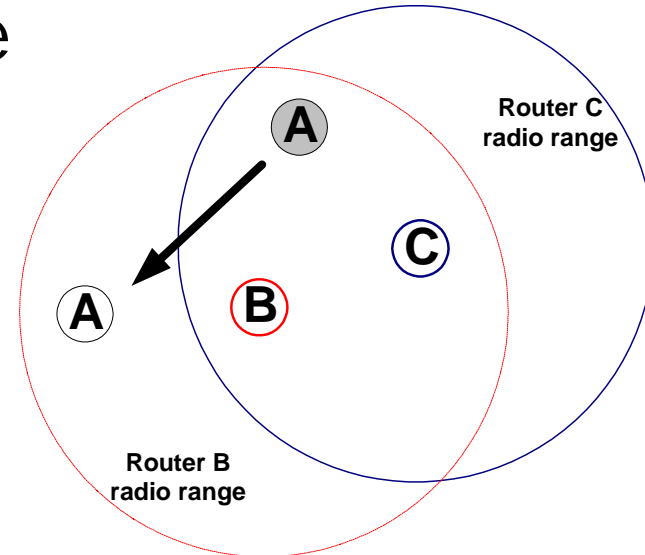

Increasing topological churn in network

Relatively static topology at this range

Note: OLSRv7 as separate layer-3 protocol is plotted, for comparison
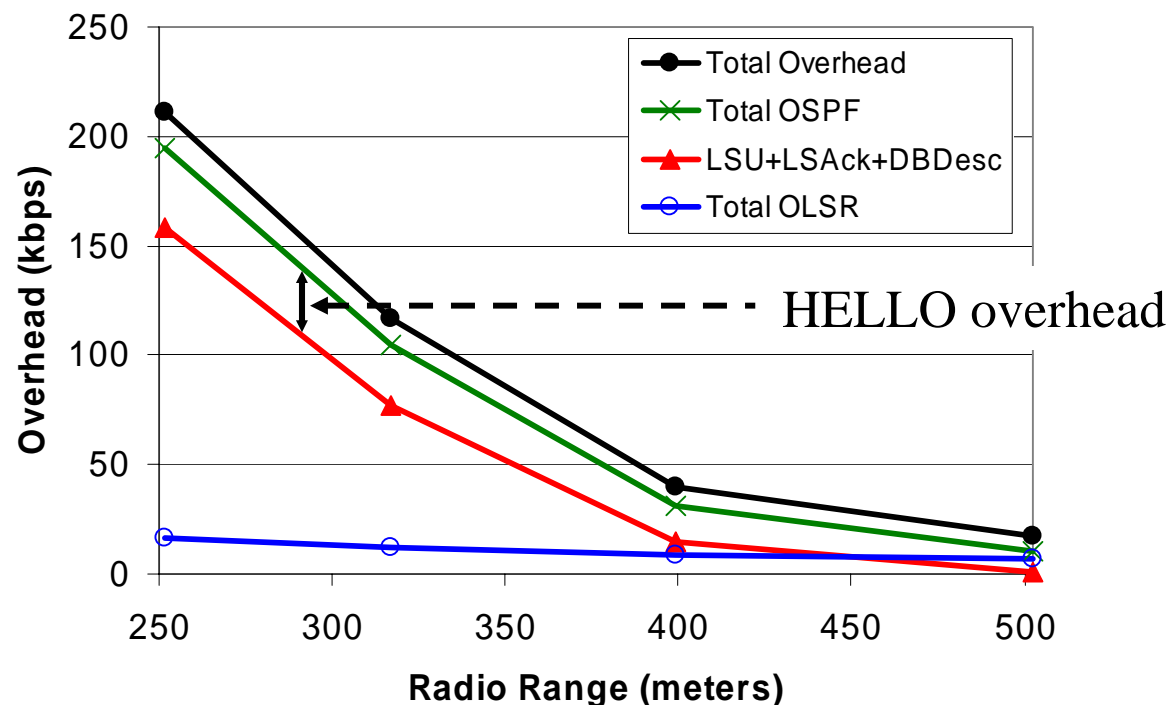
29

Why does cross-layer help so much?

- relaying Hellos via multiple hops is fragile, especially with 802.11 MAC

- in figure at right, while layer-2 (OLSR) may heal the route between A and C, enough Hellos may be lost to cause LSA regeneration

- with cross-layer integration, layer-2 routing table is used as layer-3 neighbor table, and layer-3 adjacencies are more stable
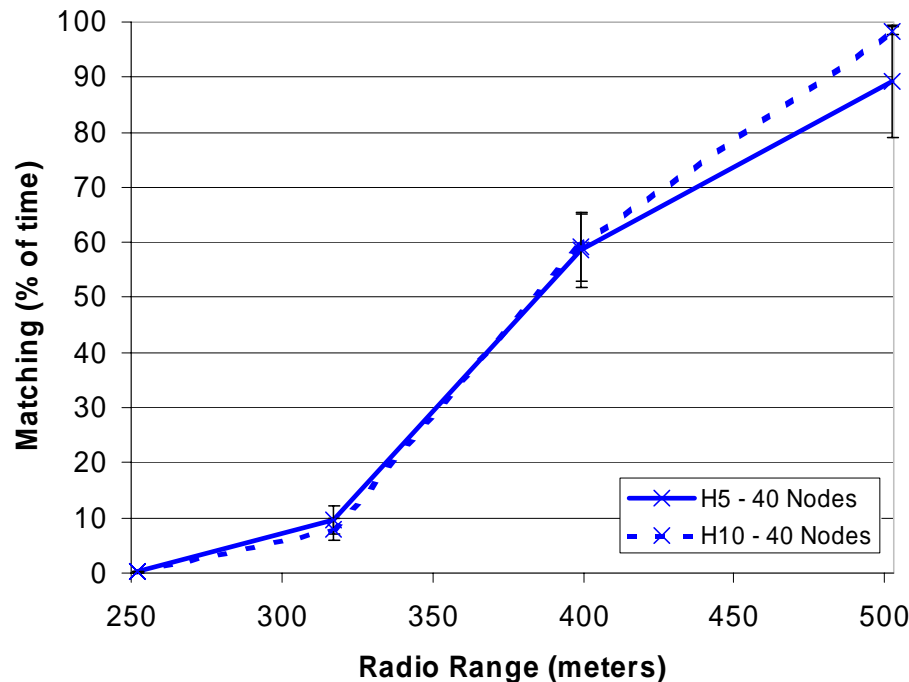
OSPF overhead components

- ## Why non-integrated routing performs poorly:

  1. duplicate HELLOs between layer-2 and layer-3 are only a small part of the relatively poor overhead performance -- it is instead a large amount of other OSPF packet types that contribute to the overhead.

31

Routing table misalignment

- ## Why non-integrated routing performs poorly:

  2. there is a considerable lag between the discovery of topology change at layer-2 and its subsequent discovery at layer-3. The mismatch between layer-2 and layer-3 leads to excessive overhead as OSPF adjacencies are broken and formed.

% of time that OLSR routes match OSPF neighbors

Summary

- Non-integrated routing, working independently at different layers, could perform more poorly than expected
  - Inconsistencies between topology map at layer-2 and layer-3 manifest themselves in overhead and churn

It appears dangerous to casually assume that a MANET protocol operated at layer-2 can provide a robust full-mesh "illusion" to layer-3

Summary

- Cross-layer routing integration has the following potential benefits:

  – i) provides a stable view of the layer-2 topology to layer-3, thereby reducing overhead to just a bit more than the layer-2 protocol itself;

  – ii) is potentially better suited to multicast routing approaches;

  – iii) causes a stable picture of the wireless subnet to be propagated to the outside network, while being robust to subnet partitioning.

Summary

- What about a modified layer-3 IGP?
  - i) provides better ingress points to a wireless topology
  - ii) also good for multicast, for the same reasons
  - iii) may eventually be more standards-based

Cross-layer benefits:
- more scalable for larger networks, due to hiding of topology changes from outside network

MANET-capable IGP benefits:
- better ingress routing may be better for resource constrained networks (less routing stretch)
- may be a more general, standards-based approach