Ad Hoc Networks, IETF, and Convergence in Solution Space

Charles E. Perkins
Nokia Research Center
charles.perkins@nokia.com

MobiArch 2006
December 1, 2006

NOKIA

Outline of Presentation

- Why I am here
- Ad Hoc Networks in general
- Recent results from *manet*
- Performance observations
- Flooding – a potential modular component
- Convergence – by creating parameterized modular components

Ad Hoc Network characteristics

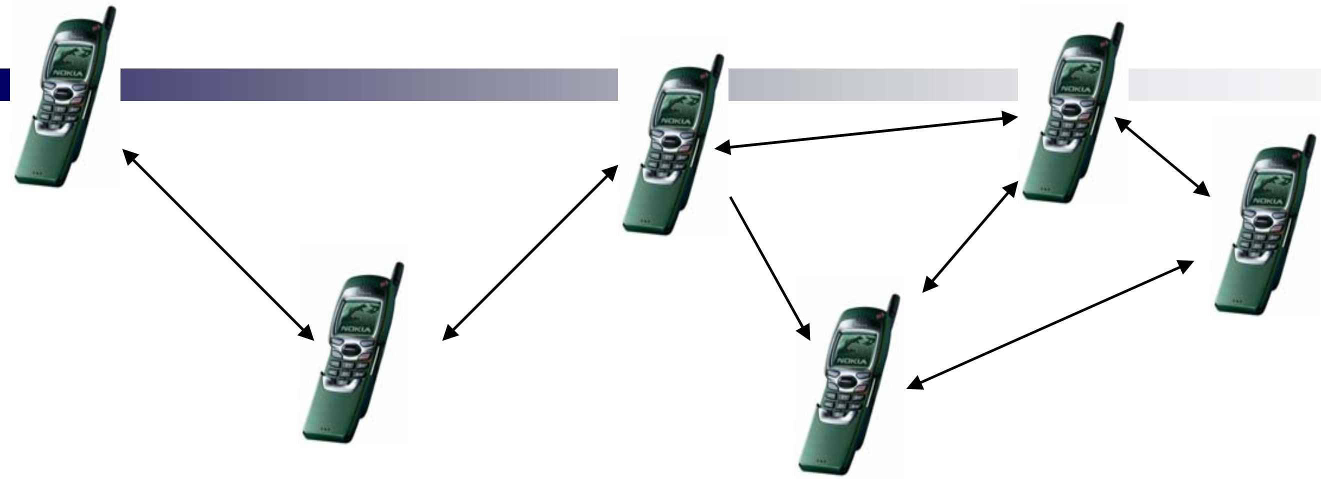

- peer-to-peer
- multihop
- dynamic
- *Really* "anytime, anywhere"
- zero-administration
- low power
- autonomous
- autoconfigured

But, most of these have exceptions!

Ad Hoc Research growing rapidly

- MobiHoc known as the premier ad-hoc forum
  - ~250 papers submitted for 2004 in Tokyo
  - Even more submitted for 2005 in Urbana-Champaign
  - See you in Florence for 2006!
- Numerous other conferences of interest
  - *Many* ad-hoc papers submitted to Mobicom
  - Globecom, Infocom, AdHoc-NOW, INSS, ...
  - Journal of Ad Hoc Networking
- Helps understand fundamentals of routing
- Major interest in sensor nets (e.g. NSF)
- Useful as a buzzword for paper acceptance

NOKIA

Ad Hoc Routing Projects

- Terminodes (EPFL)
- WINGs (JJ Garcia/UCSC)
- ROAM (JJ Garcia/UCSC)
- WAMIS (Gerla/UCLA)
- ODMRP (S.J. Lee/UCLA)
- TRAVLR (Kleinrock)
- Tora/IMEP (Park/UMD)
- Link Quality (Dube/UMD)
- LAR (Texas A&M)
- TBRPF/PacketHop (SRI)
- OLSR (Clausen/Jacquet)
- DSDV (Dest. Sequence #'s)
- AODV (refinement of DSDV)
- AOMDV (Multipath/Das et al.)
- LANMAR (Gerla et.al/UCLA)
- GPSR (Karp/Harvard)
- CBRP (Singapore)
- DSR (Dave Johnson, CMU)
- MMWN (Steenstrup/BBN)
- ABR (C.K. Toh)
- STAR (JJ Garcia/UCSC)
- ZRP (Zygmunt Haas/Cornell)
- Fisheye/Hierarchical (UCLA)
- CEDAR (Urbana-Champaign)

More Ad Hoc Routing Projects

- FRESH (latest encounter)
- ANTS(*swarm intelligence*)
- Ariadne
- Cryptographic Threshhold
- Insignia [QoS] (Columbia)
- AODV6
- FLR ["Feasible"] (UCSC)
- GPS/Geographic
- SHARP
- DMAC (Directional)
- Pulse
- TDR (Trigger based Distributive)
- DREAM
- SAODV (Guerrera)
- LDR (Mosko/Garcia …/Perkins)
- AODVjr(Chakeres/Klein-Berndt)
- WRP
- Minimum-energy approaches
- Compow
- Face Routing (GOAFR+,…)
- XTC (Topology Control)
- *Many* more…

NOKIA

Traditional Routing Methods

Single metric: number of hops to destination
- But this isn't really appropriate, esp. for 802.11

- Advantages of using routing protocols:
  - Self-Starting
  - Multi-Hop
  - Dynamic topology
- Link-State (*Dijkstra's* shortest-path algorithm)
  - Complete topology stored
  - OSPF (RFC 1583)
- Distance-Vector protocols (*Distributed Bellman-Ford*)

On-Demand Routing Protocols

- Eliminate route table updates for routes that are not used
- Fewer control packets:
  - → Better scalability
  - → Reduced congestion
  - → More robust protocol action
- Less frequent control packets → reduced processing requirement
- Even more localization for topology changes if distance vector
- Also can be made to work for link-state

On-Demand Routing, cont.

- Downsides:
  - ICMP Unreachable only after Route Discovery attempt
  - Latency → longer application launch times
  - Route Discovery broadcasts

What’s wrong with hopcount?

- For some applications, monetary cost or security or minimum delay is crucial
- For 802.11, hopcount is not the best!
  - minimal hop count maximizes the length of the links between nodes
  - longer links operate at reduced bandwidth
  - longer links are less robust against node movement
- Suggestion for metric: “link occupancy”
  - Reduces to hop count if each link is counted as 1.

Mobile Ad Hoc Networking (*manet*)

- AODV: *on-demand*, and *distance-vector*
  - Interoperability testing
  - Experimental RFC 3561
- Other *on-demand* protocol is (DSR)
- Two link-state, *table-driven* / *proactive* protocols
  - Optimized Link-State Routing (OLSR) is RFC 3626
  - Topology-Based Reverse Path Forwarding (TBRPF) is RFC 3684
- DSR should also be published as Experimental
- Many other protocols have been considered!
  - For instance, quite a few of the previous list

MANET status update

- IETF group rechartered to focus on engineering existing techniques into a Proposed Standard
  - Two routing protocols
  - One flooding/multicast protocol (is this possible???)
  - Internet Gateway operation
  - No “fancy” (usually, → unproven) approaches
- Recent discussion related to OSPF
  - Adjoining ad hoc domains to OSPF infrastructure?
  - Can OSPF become an ad hoc network routing protocol?
- IRTF group “Ad Hoc Network Systems” (ANS) has been formed [*needs help* : see www.irtf.org]

IETF structure

IETF has Areas and Area Directors (ADs)

IETF has over 100 working groups:

- General Area (AD is IETF chair)
- Applications Area
- Internet Area (most mobility groups here)
- Operations and Management Area
- Routing Area ([manet] is here!)
- Security Area
- Transport Area

NOKIA

IETF mantra

Rough consensus and running code

Consensus requires team building and persistence.
Running code requires, well, you know…
(but including interoperability too!)

IETF Organizational structure

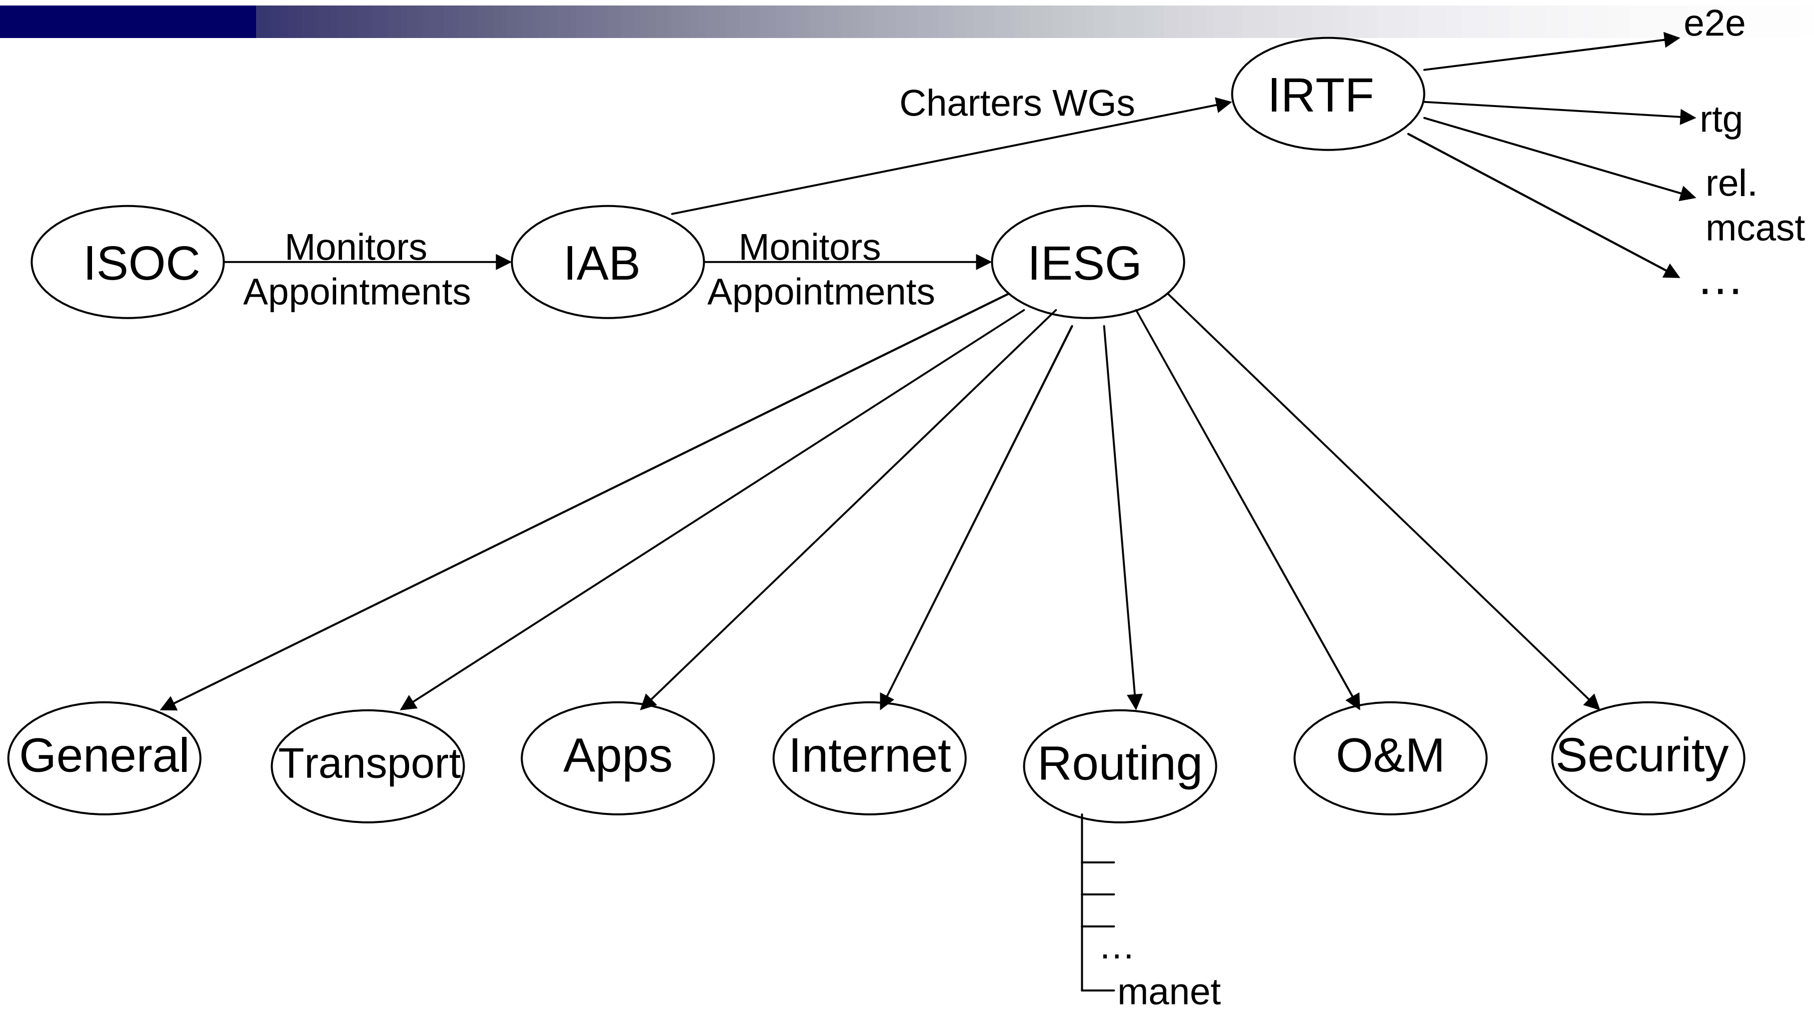

IETF process (theoretically…)

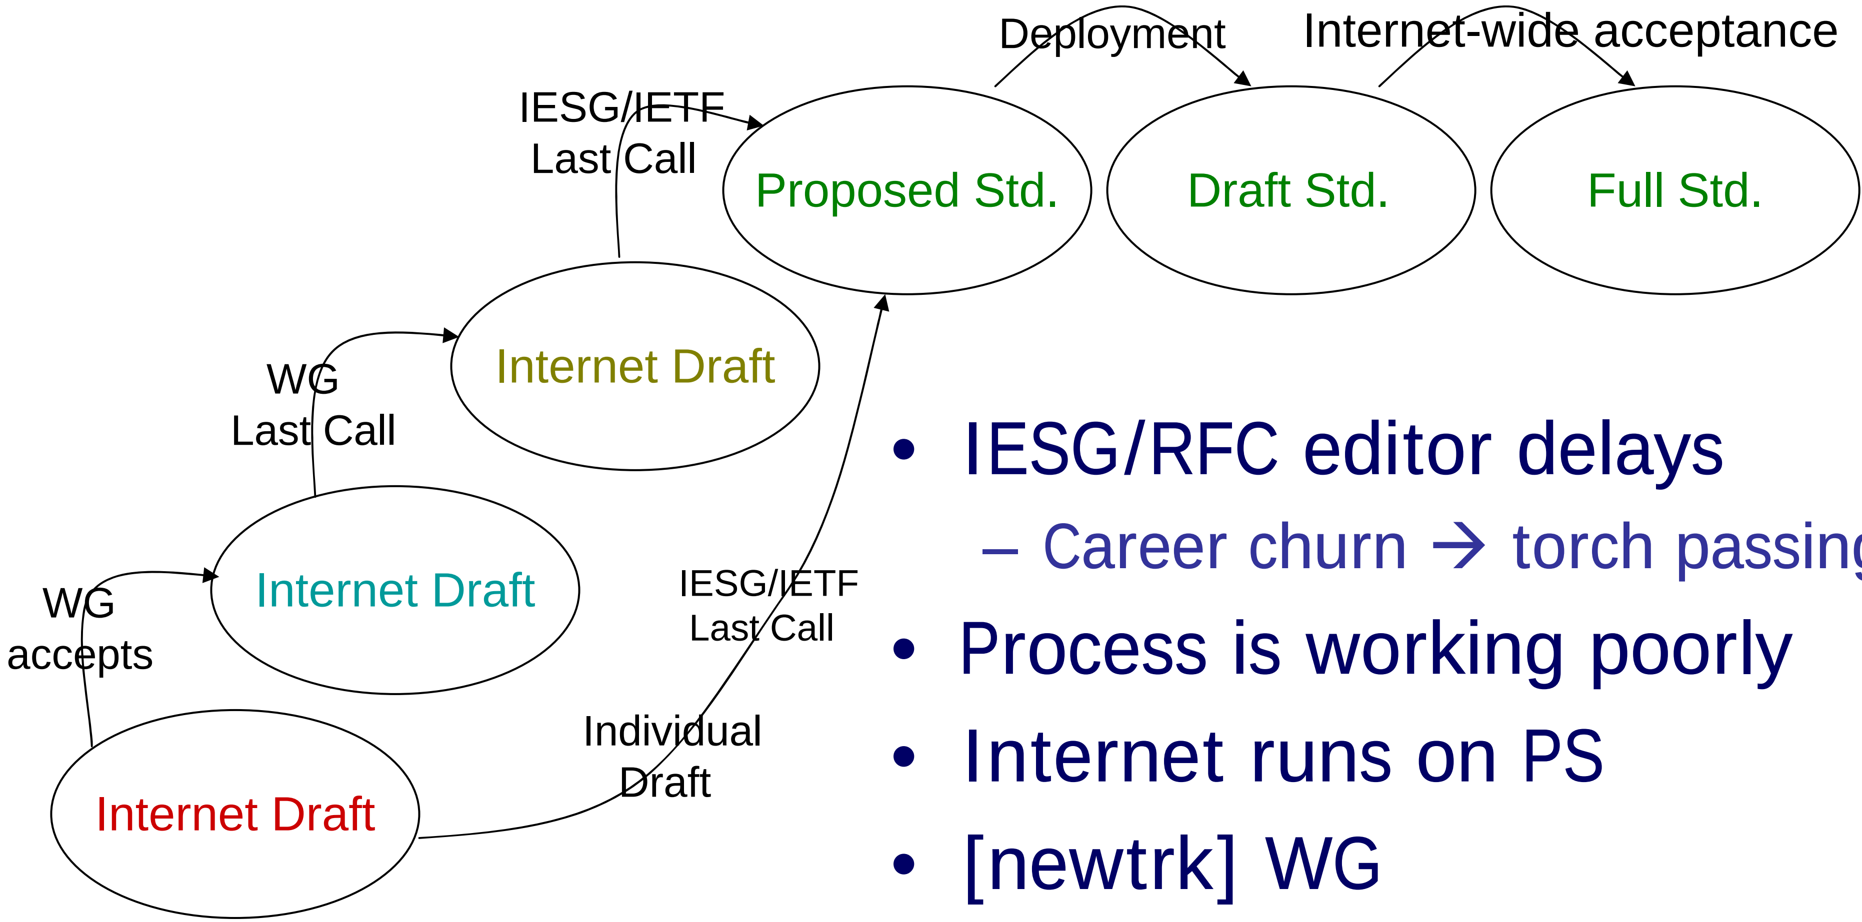

- IESG/RFC editor delays
  - Career churn → torch passing
- Process is working poorly
- Internet runs on PS
- [newtrk] WG

NOKIA

On-Demand Unicast Route Discovery Initiation

Route Request (RREQ) broadcast flood

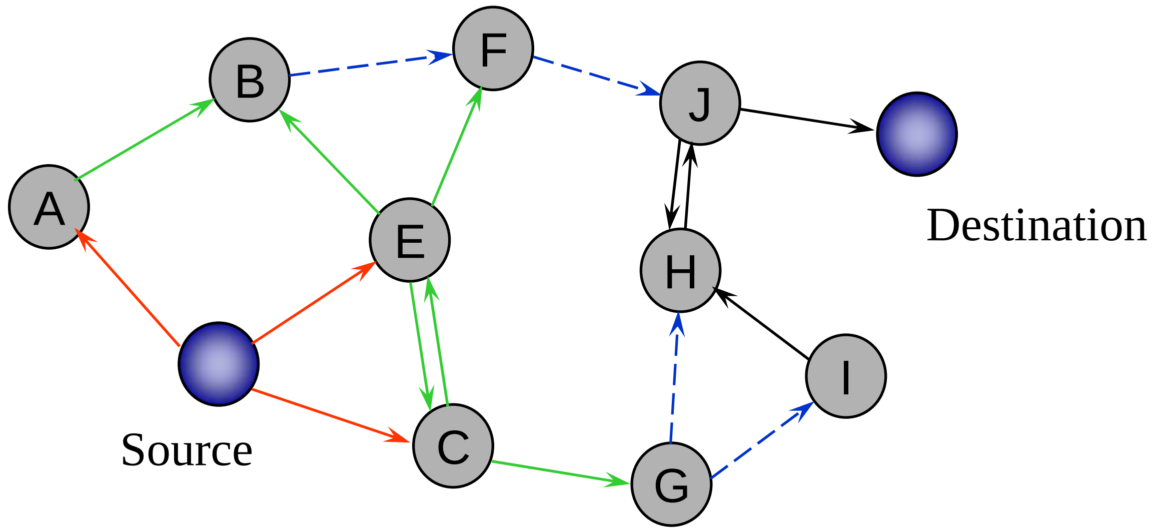

On-Demand Unicast Route Discovery Completion

Route Reply (RREP) propagation

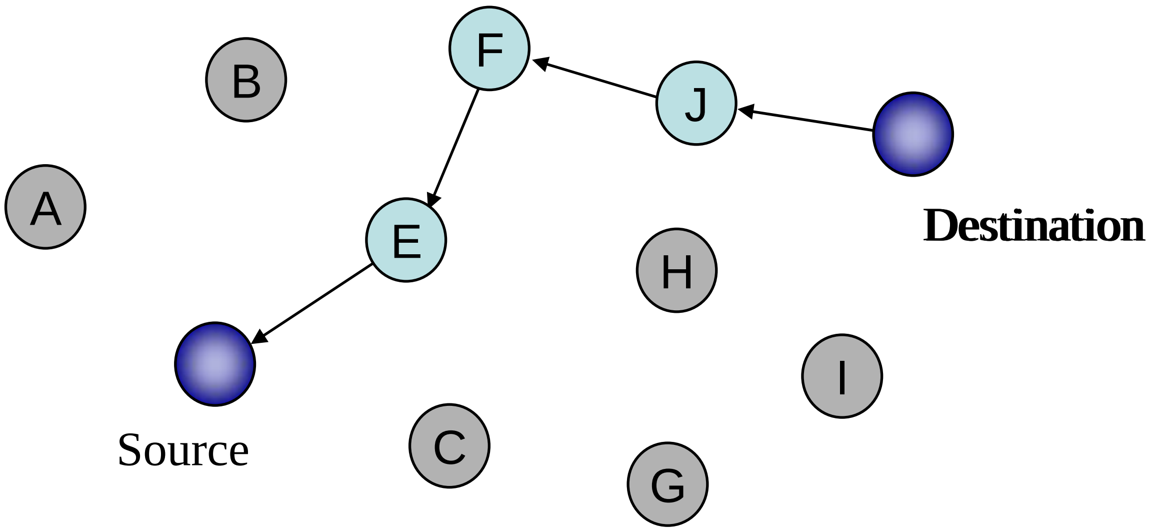

Some general performance observations

- When two protocols both lose almost all packets, maybe it doesn't matter which one is "better"
- Flooding → congestion, *and* flooding is unreliable
  - Problematic for creating OSPF extensions!
- At low node populations, what choices matter?
- High hop count increases fragility, latency
- N.B.: minimum hop count can be a *lousy* metric
- On-demand increases startup latency
- Table-driven tends to increase congestion
- Simulation times grow quadr. w/node population

Simulation performance results

- Old AODV at 10,000 nodes performs poorly
  - 25% packet deliveries in the best of circumstances
  - Even worse without local repair and expanding-ring
- AODV vs. DSR with limited node populations
  - DSR works better under conditions of low mobility
  - Node movement favors AODV's route management
- MAODV has been tested under ns-2, and shows performance difficulties even at low populations
- Gün Shirer at Cornell offers the *Staged Network Simulator* (SNS) using ns-2 for big simulations

More performance results

- # RREQs ~ linearly with the node population
- Line's slope changes depending on strategy
- At 10,000 nodes, most packets are control traffic (in one case, ratio was 5000 to 1)
- End-to-end delay wasn't outrageously terrible (150ms) even at high node populations
- AODV w/expanding ring has the longest latency
- Query localization seems not to work (?why?)
- Should be similar for other on-demand protocols

NOKIA

Is Distance Vector *better* than Link-State?

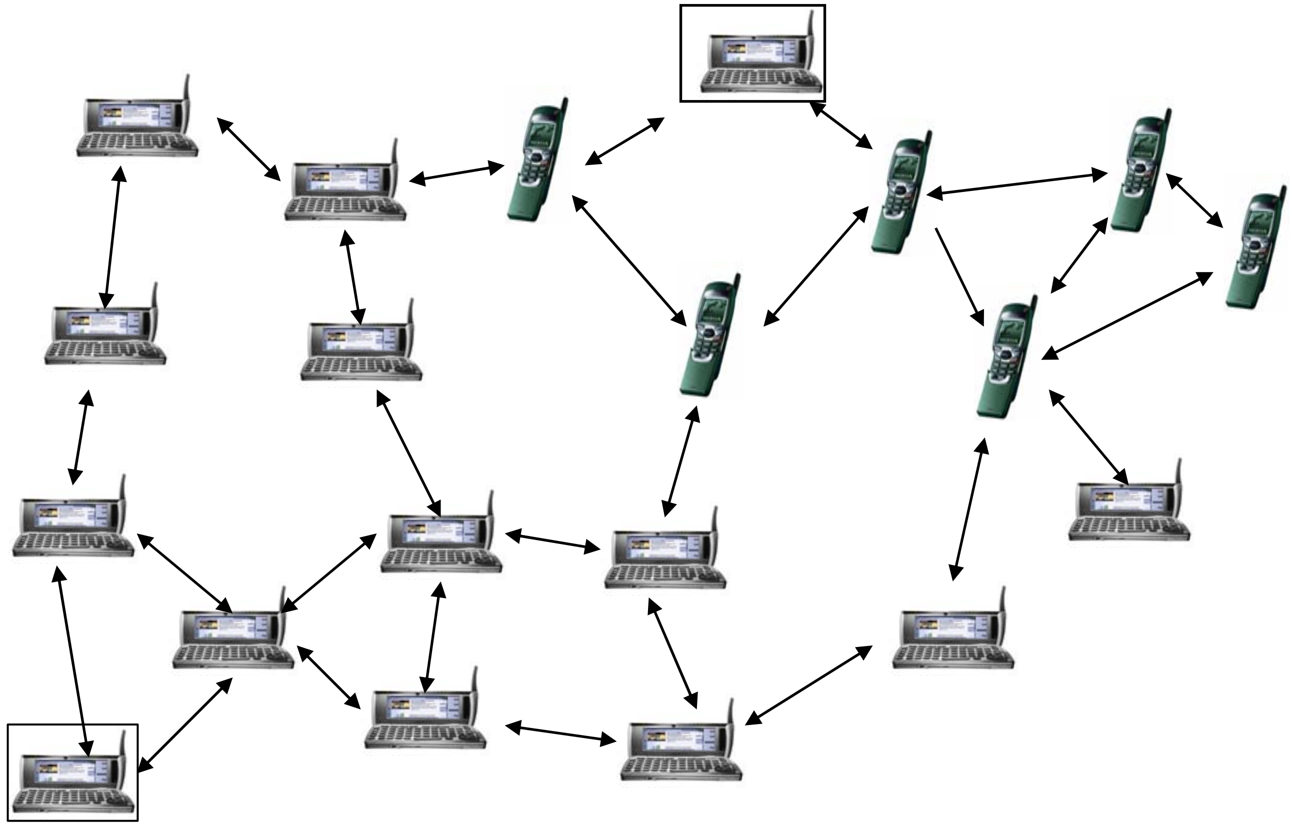

Distance Vector Characteristics

- Very suitable for *on-demand* operation
- Remote movement less likely to propagate
  - i.e., mobility has more localized effects
- Natural fit for IP route table operation
  - e.g., OLSR and TBRPF use a shortest-path algorithm to fill route table with distance-vector entries
- To handle multipath, sort by metric

DYMO: DSR + “AODV--” converged

- Co-authors: Ian Chakeres, Elizabeth Belding-Royer
- Fewer mandates (using experience from AODVjr)
- Only destination answers RREP (in “base protocol”)
- Route caching and route lifetime extension
- Distance-vector, on-demand
- New header formats, easy to program
- Future extensibility by way of flexible error handling
- Optional path accumulation for route discovery
- Improved sequence # algorithm
  - fewer incrementations, distance taken into account

Incrementing a sequence number

- A sequence number is ***only*** updated by the node that owns it (in contrast to DSDV and to AODV), on generation of RREQ or RREP
- A destination (e.g., TargetAddress) keeps its sequence number the same if THopCnt is greater than the REHopCnt of the source's REBlock (i.e., the first one).
  - In this case, the destination has not moved farther away from the source
  - This algorithm maintains proof of correctness

Sequence number algorithm

- Nodes accept a route table entry update if:
  1. The sequence number of the destination node is greater, <u>or</u>
  2. The sequence number is the same, and the metric (e.g., hopcount) is smaller
- If the sequence number in the update is smaller, the node MUST NOT accept the update.
- If the metric is the same, the node MAY accept the update
- Recent RREP observations → slight modification!

Use Route Discovery for Topology

- Why *waste* all that capacity?
- *path accumulation* on RREQ and RREP
- One step along the way towards link-state approach (but must avoid stale link info...)
- Only relevant for *on-demand*

Path accumulation, schematic

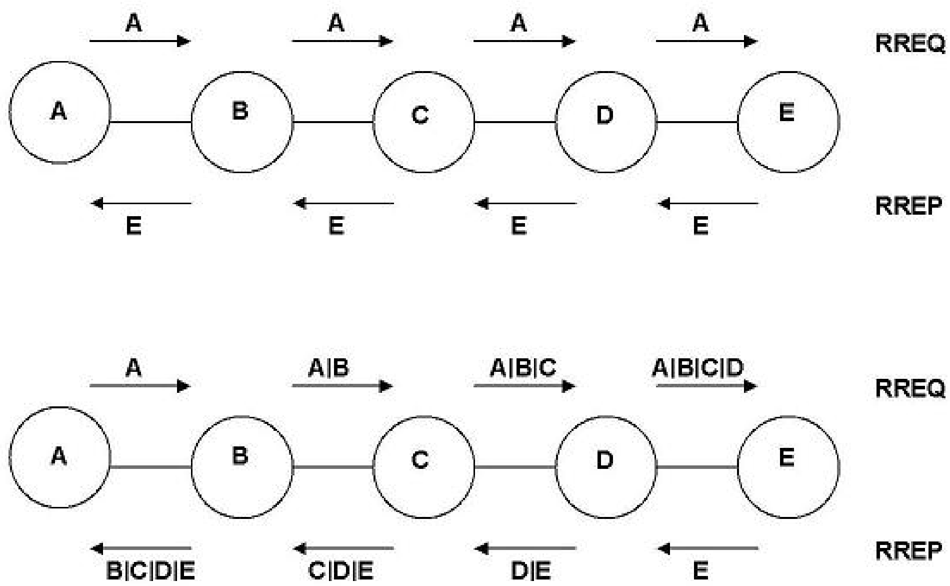

- Put more topology data in RREQ and RREP
- Longer routes allow acquisition of more data

NOKIA

Added-value Signaling

Route Request (RREQ) broadcast flood

– Each retransmission carries more link info

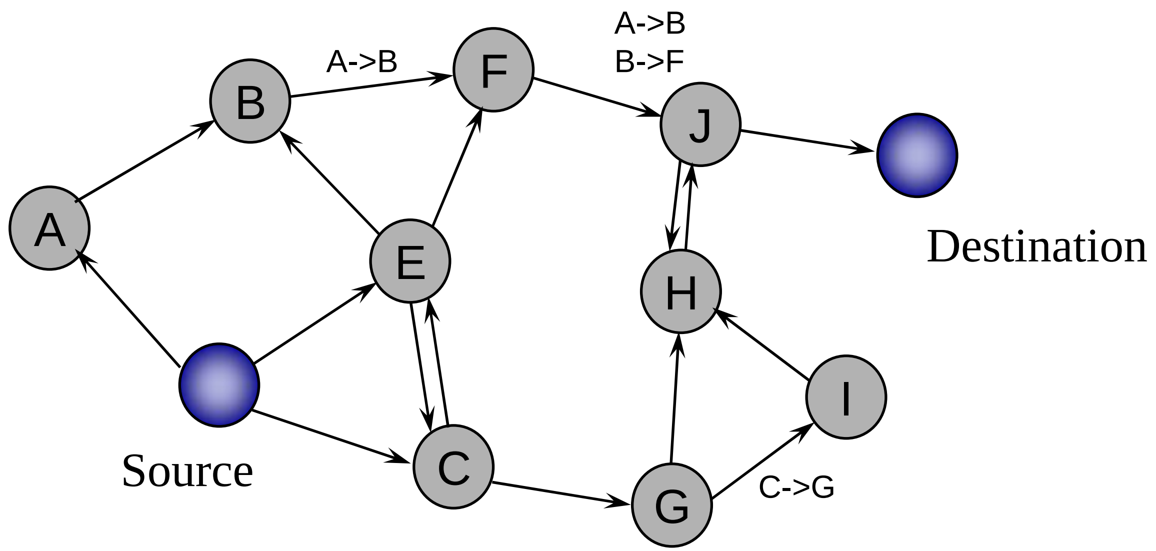

NOKIA

Path accumulation

- DYMO specifies an extension for it
- Sometimes improvements, sometimes no change, and sometimes slight deterioration
- When basic signaling gives very high PDR, then path accumulation will not improve it
- Reducing RREQ will allow higher node density without producing congestion

Path Accumulation problem!

- *Strangely,* early results show performance loss in some cases (50 nodes, 20 m/sec, 15 UDP sources)
- Previous publication showed performance improvement
- We're now learning why

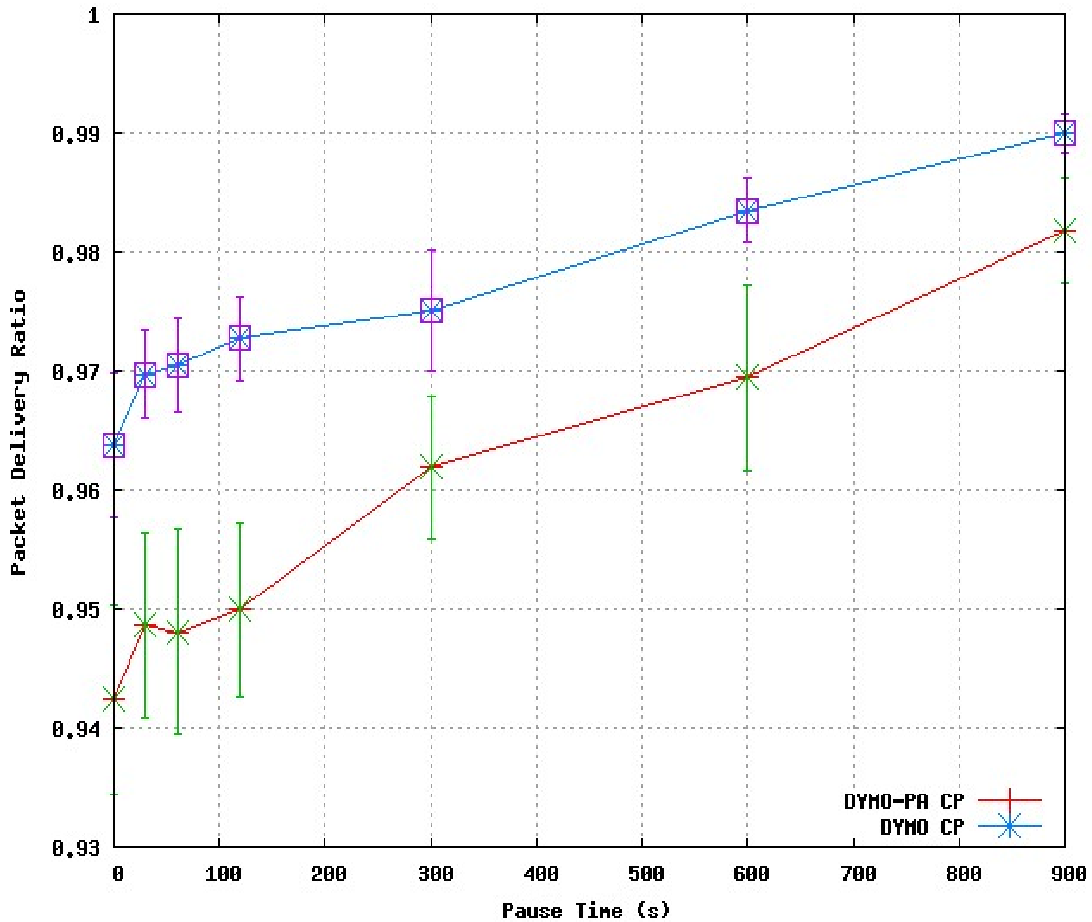

Evaluating path accumulation (results by Karim Seada et al.)

- How well does path accumulation help to discover the network topology?
- First analysis has been for static case
- How many RREQs are needed to discover the entire network topology?

NOKIA

Pre-empting Route Discovery (analytical result)

- Analysis by iteratively applying an expression for incremental effectiveness of each RREQ

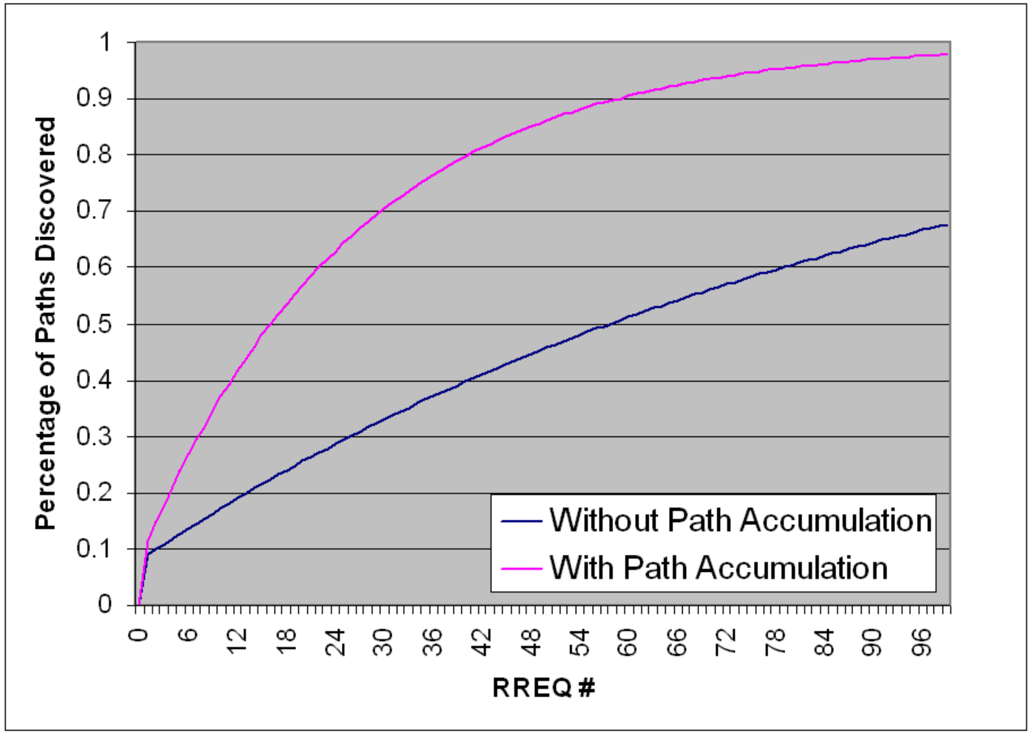

NOKIA

Pre-empting route discovery (simulation)

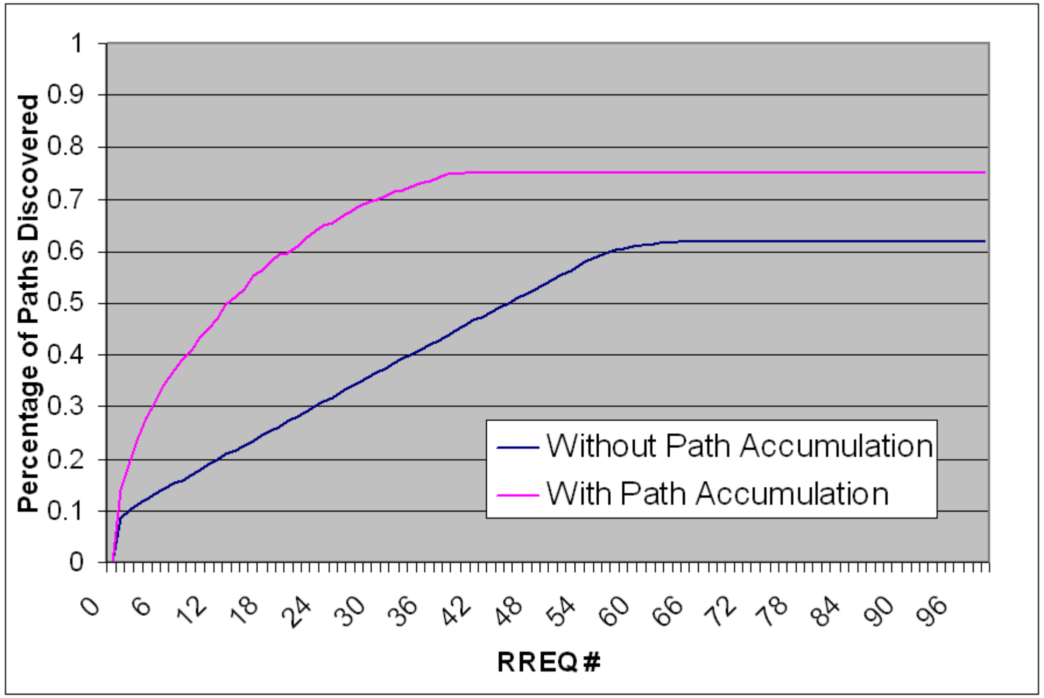

- TODO: eliminate simulation artifact causing nonconvergence to 1.0

NOKIA

Route length reduced

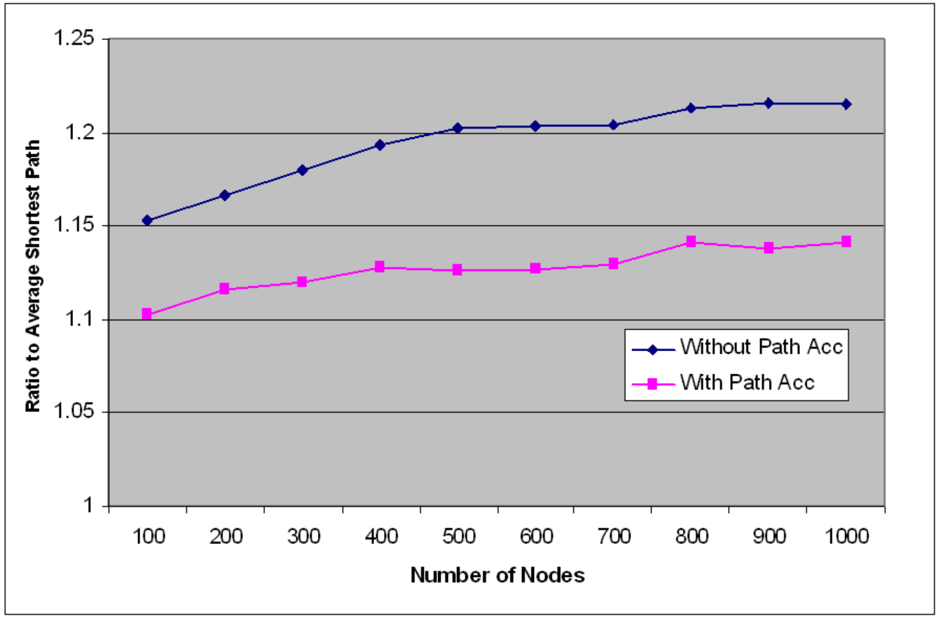

- More benefit from intermediate node RREP

What about overall performance?

- Path accumulation definitely reduces the number of RREQs
- However, it also increases the packet size
- Benefit is reduced if newly discovered routes are not used before timeout
- Packet size is often a burden that negates some of the benefit of path accumulation
  - Heed this as a warning against packet bloat!!

Ways to produce convergence

- Modularize features, new and old (easier said than done!)
  - Flooding – example given, SMURF
  - Expanding rings search/fisheye routing
  - QoS routing
  - Pulsar/clusterhead/hierarchical/...
  - Internet Gateway operation
  - Multipath, address allocation, etc., etc., ...
- Apply new advances to each routing protocol...
- Eventually, common part may dominate!

Flooding: Needed for *discovery*

- “Application” flooding vs. “IP-level” flooding
  - TTL = 1 vs. TTL = network-diameter++ vs. ...
- Multicast vs. Broadcast vs. ???
  - No multicast tree needed
  - 255.255.255.255 isn’t right
  - No subnet broadcast
  - Wanted: *manet-local* flooding
- Our goal: ***Many*** fewer packet retransmissions
- Technique: Fewer nodes retransmitting
  - E.g., by picking a set of multipoint relays (MPRs)
- Needed: unique identification for flooded packets

NOKIA

Connected dominating set (CDS)

- A dominating set covers the whole network
- A *connected* dominating set simplifies forwarding
- The set of all nodes in a network is a (big) CDS (assuming the network is connected)

Multipoint Relay (MPR) Flooding

- Taken from OLSR: so, it's known to "work"
- Module meant for use by <u>any</u> flooding application (e.g., all four of the current experimental protocols)
- One kind of *dominating set*

Neighborhoods and advertisements

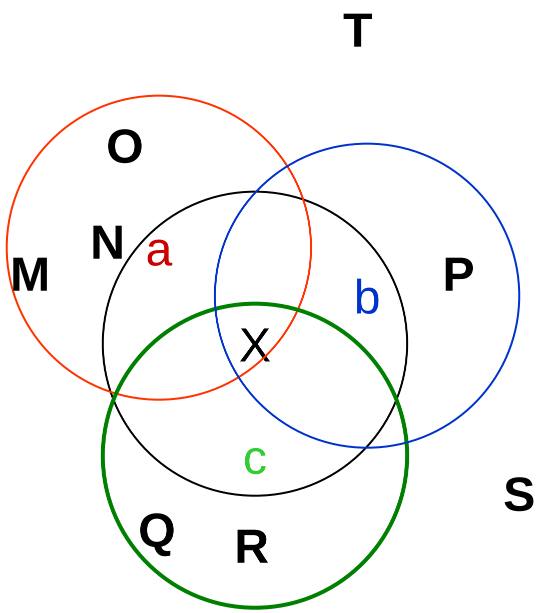

- X's one-hop neighborhood: $N_1(X) = \{a, b, c\}$
- X's two-hop neighborhood: $N_2(X)=\{a, b, c, M, N, O, P, Q, R\}$

MPR selection algorithm

- A node X combines local advertisement data to tabulate its 2-hop neighborhood $N_2(X)$
- The idea is to pick some neighbors that will “cover” the two hop neighborhood $N_2(X)$
- It first picks those neighbors that MUST be included in order to cover $N_2(X)$
- It then uses a greedy algorithm to cover the rest of the two-hop neighborhood
  - Namely, preferentially pick the neighbor that covers the most second-hop neighbors, and iterate

MPR Flooding Issues

- Use of *all-manet-nodes* multicast address
- *Bundling* for multiple simultaneous messages?
- MPR dependence on last hop
  - Else, how do receivers detect sender's identity?
- ICMP vs. UDP vs. IP vs. ??
- Redundant coverage (> 1 seems advisable)
  - Broadcast minimized ➔ process is *fragile*
- Only MPRs can be relays → non-optimal routing!
- MPR nodes in all routes → reduced MPR life!
- Unneeded for uncongested or transient networks

Flooding comparisons (results 2005)

- We can show nice pictures for the nodes that become part of the broadcast skeleton
- Minimal broadcast *does* often reduce PDR
- At 1,000 nodes, TBRPF took all weekend to simulate 3 seconds
- At 1,000 nodes, AODV plus MPR flooding took 30 minutes to simulate 900 seconds
- We also have ideas for further improving the simulator (SNS)
- MUCH work remained to be done!!

Get a smaller CDS from MPRs

- The set of MPRs emanating from a node is a (relatively big) CDS
- Algorithm for getting a smaller CDS from MPRs:
  - Node joins the CDS if it has lowest ID in the its neighborhood
  - Node joins the CDS if it is an MPR of its neighbor that has the lowest ID.
  - See INRIA Research Report 4597 (C. Adjih et al.)
- This CDS is _not_ “source specific” (see proof)

Recent results for SMURF

- SMURF – Simplified Multicast Routing and Forwarding
  - A modular flooding component
- Uses the CDS algorithm just described
- Shows increased PDR under recent tests
- Has a component for reliable flooding
  - BUT – making broadcasts reliable increases congestion (undesirable!)

Greatly reduced # of RREQs

- SMURF backbone does its job very well!

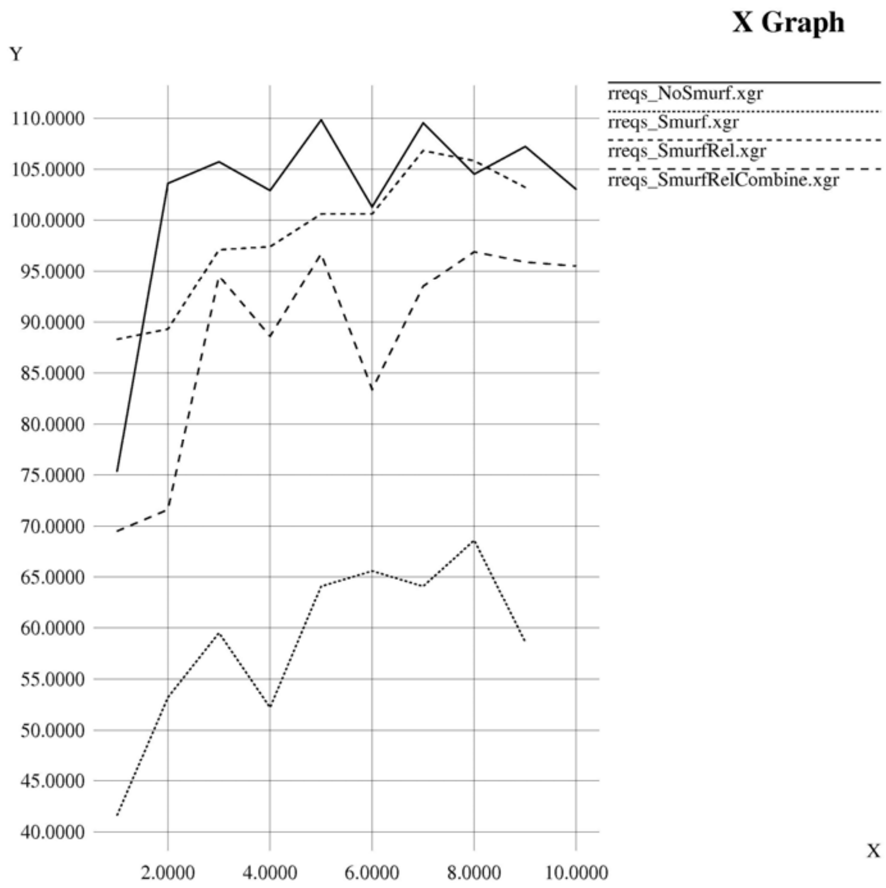

Greatly improved PDR

SMURF without any reliability signaling enables AODV to perform quite a bit better than base AODV.

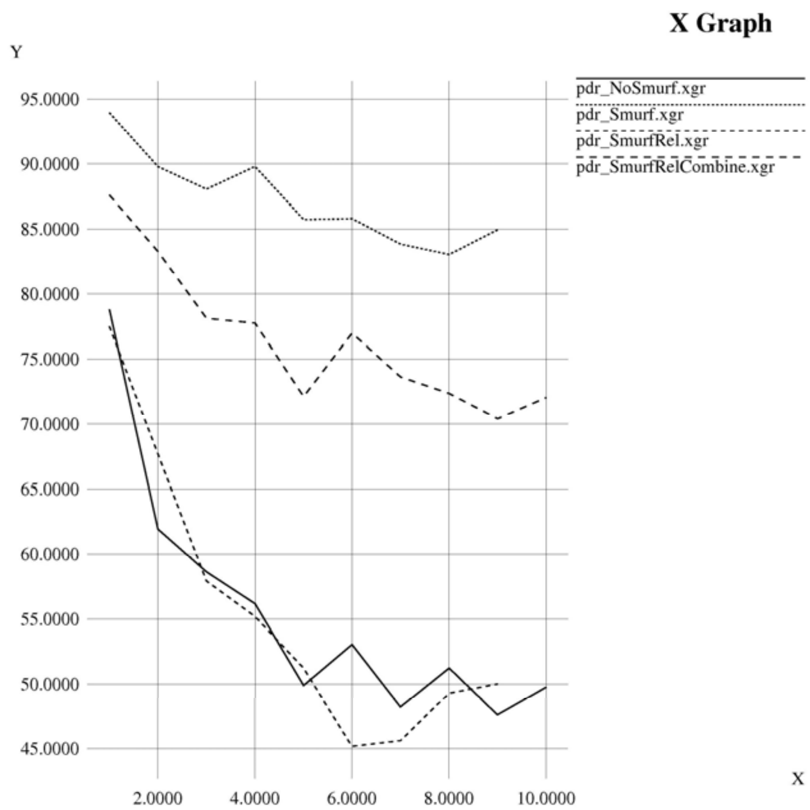

Next steps

- Why doesn't reliability help??!
- Improve signaling overhead of backbone flooding, without reducing effectiveness
  - Needed: bundling with nbhd discovery
- Comparison of OLSR with AODV
  - OLSR is better for very dense communications
  - AODV is better for sparse communications
  - Where is the crossover? [measured against percentage of N^2 possible communications]

Convergence ideas

- DYMO := AODV + DSR; OLSR with TBRPF
- Distance Vector with Link State
- On Demand with Proactive
- Modular, Constructible approach
- Adaptive/Hybrid approach
- Path accumulation for all on-demand protocols
- Simulation Results
  - http://lsewww.epfl.ch/Documents/acrobat/CSA02b.pdf
  - "Simplified Simulation Models for Indoor MANET Evaluation Are Not Robust" (Secon 2004)

NOKIA

OLSR converged with TBRPF

- Protocols are both link-state routing protocols
- Both report only restricted topology information
- Story about the role of patents
- OLSR uses "Multi-Point Relays", as described
- A TBRPF node relays broadcasts from a neighbor *j* only if it belongs to *j*'s *reported nodes* set.
  - *Reported nodes* are those which are next hops towards farther destinations (i.e., not leaf nodes in the *source tree*)

NOKIA

Merging Proactive and On-Demand

- Key parameter: *ACTIVE_ROUTE_TIMEOUT*
- If *ACTIVE_ROUTE_TIMEOUT* >> 0, route repair will maintain routes
  - Example: Internet Gateways
- Special case solution: multi-hop Route Advertisement
- Helpful: frequent topology updates
  - potentially via “rich” Route Discovery

Service Discovery

- Needs same sort of “flooding” operation
- But, instead of an “IP address”, a service is needed which meets some desired service criteria (name & attributes)
- Allow a service to be identified by the application *port number*
- Alternatively, use SLP *service descriptors*
  - Others exist

Ad Hoc Quality of Service

- Add QoS constraint to link descriptor
  - RREQ for on-demand
  - Topology updates for proactive
- Nodes only forward RREQ if they can possibly meet constraint
- Need ICMP for links that “fail”
- NP complete problems abound, due to congestion management, scheduling

Challenges for the Future

- Getting to Standard!
- Multicast/Anycast/Geocast/Mobicast
- Security (e.g., route repair!)
- Scalability: the 1/sqrt(N) capacity limit
  - Backbone formation and maintenance
- QoS – and don't forget layer 2!
- Multipath routing "vs." route caching
- Route Repair vs. multihop context transfer
- Re-examine the "client-server" paradigm
- Using positional hints (for sensors, worth it!)

Summary and Conclusions

- IETF *manet* working group working to converge
- Distance Vector can be made loop free, and localizes the effect of topology changes
- On-demand protocols offer many advantages
- Flooding, e.g., is very basic but also complex
- Creating modular components aids convergence
- Convergence aids getting to standard
- Ad Hoc Networking is a great research area
  - Can be applied whenever *infrastructureless*
  - Related fields: sensor networks, graph theory, ...

NOKIA

Backup slides…

Commercial Opportunities

- Conferencing
- Home networking / Community (mesh) networking
- Emergency services
  - Ambulance, Police
  - Disasters (natural or man-made)
- Hospitals
- Embedded computing applications
  - Ubiquitous computers with short-range interactions
  - Automotive/PC interaction
  - What if wireless computers are *everywhere*?

Other Envisioned Applications

- Digital Battlefield Communications
  - Including sensor networks
- Movable base stations
  - Many military applications
- Campus wireless access from quadrangles
- Immediate, interpersonal communications
- Range extension for cellular telephones
- Enable computing where subnets do not exist

What is ***networking*** good for?

Sensor Network Characteristics

- Less dynamic than other ad hoc networks
- Large network sizes
- Battery power truly at a premium
- Congestion less of an issue
- What about latency?!
- Identity of individual nodes less important
  – Affects even concepts of addressability
  – Increases need for multicast/anycast/geocast?

Generic DYMO element structure

```
0                   1                   2                   3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|M| H |     Type      |          Len            |     TTL     |I|Reserved |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:          NotifyAddress (Only Types with M-bit set)            :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
: TargetAddress (for non-DYMOcastAddress IPDestinationAddresses):
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                              Data                             :
:                      Type-Specific Payload                    :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- ‘M’ bit: error requires UERR
- ‘I’ bit: element has been ignored
- ‘H’ bits: How an unrecognized element is to be handled
  - Possibilities: drop, remove, skip, skip but set ‘I’

NOKIA

DYMO “Routing Element” format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |              Len            |      TTL    |I|A| Rsvd  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                          TargetAddress                        :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          TargetSeqNum                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  THopCnt  |Res|                                               :
+-+-+-+-+-+-+-+-+                                               :
:               Routing Element Blocks (1 or more)              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- ‘A’ bit: TargetAddress MUST acknowledge
- THopCnt: how far away the Target is, according to the source’s most current information.
  - tells the destination if incrementation is needed

NOKIA

Routing Element Block (REBlock)

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                                |G| Prefix  |Reservd| REHopCnt  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                          RENodeAddress                        :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          RENodeSeqNum                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- “Path accumulation”: each node along the way of a RREQ or RREP should append a REBlock (the originator MUST do so)
- Then, all nodes receiving the Route Element (including the destination) can see how far they are from previous nodes and get sequence number data

NOKIA

Neighborhoods and advertisements

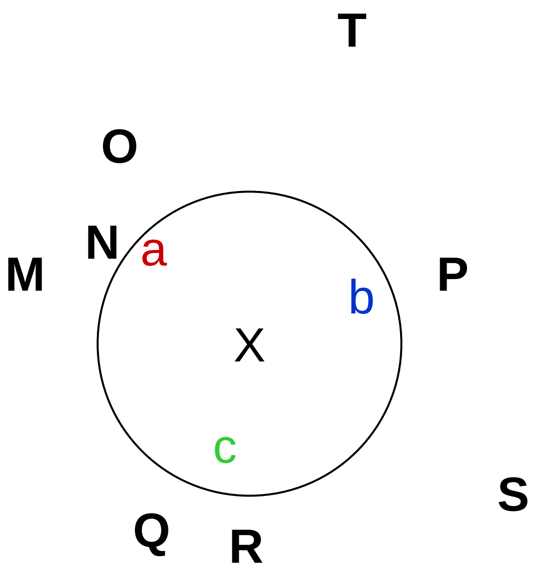

- X's one-hop neighborhood: $N_1(X) = \{a, b, c\}$

Neighborhoods and advertisements

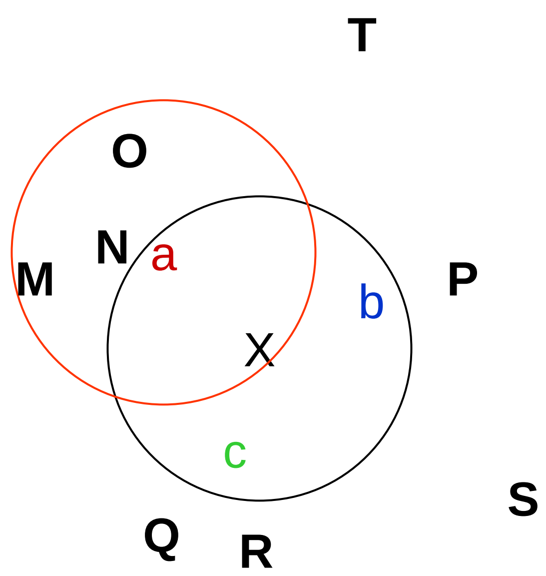

- X's one-hop neighborhood: $N_1(X) = \{a, b, c\}$
- a advertises one-hop neighborhood: {M, N, O, X}

Neighborhoods and advertisements

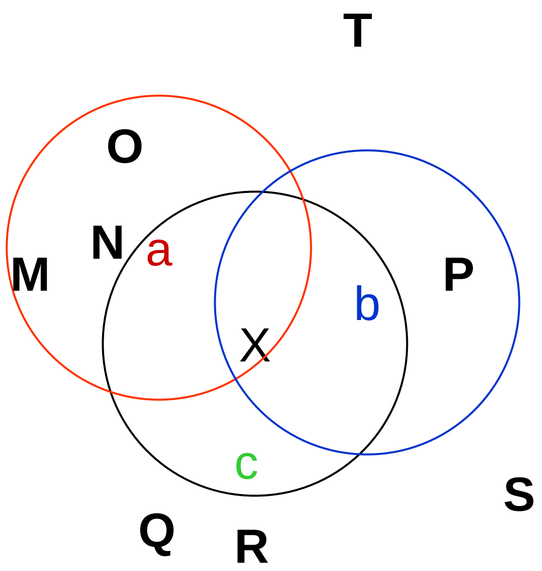

- X's one-hop neighborhood: $N_1(X) = \{a, b, c\}$
- a advertises one-hop neighborhood: {M, N, O, X}
- b advertises one-hop neighborhood: {P, X}

Neighborhoods and advertisements

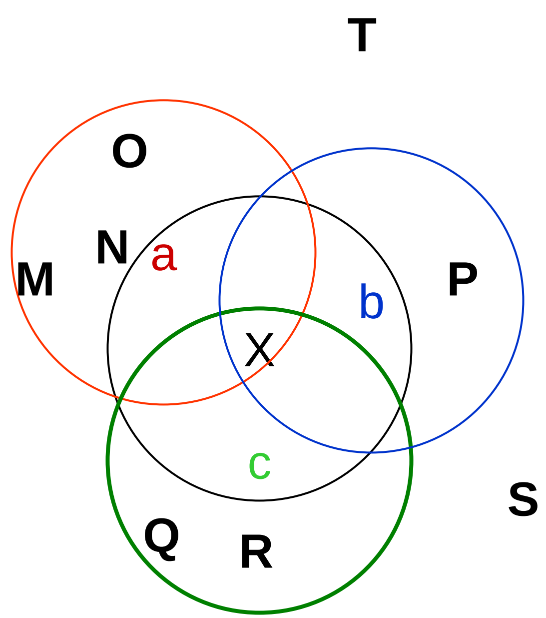

- X's one-hop neighborhood: $N_1(X) = \{a, b, c\}$
- a advertises one-hop neighborhood: {M, N, O, X}
- b advertises one-hop neighborhood: {P, X}
- c advertises one-hop neighborhood: {Q, R, X}

DYMO: DSR + “AODV--” converged

- AODVbis took a major step for this, namely path accumulation during route discovery
- DSR source routes are not always beneficial
  – Distance-vector more robustly enables route repair
- AODV route caching is beneficial
  – Inverse dependence on *relative* mobility
- AODV and DSR can use the same tricks
  – And offer the same extensions
- Multipath: DSR vs. AOMDV

Packet formats, yecch…

Type		Reserved (16 bits)
Target Address		
Target SeqNum		
Options (e.g., Binding Authorization Data)		

	A	H	S	D	
Sequence # (16 bits)	Reserved (16 bits)				
Lifetime (32 bits)					
Home Address (128 bits)					
Options (e.g., Binding Authorization Data)					

Two methods of protocol design (IETF vs. pure researcher)

1. Combine competing proposals
2. Look for ways to apply well-known ideas
3. Attend to potential dislocation incurred by pre-existing uses
4. Acknowledge and solicit contributions and criticisms

1. Maintain publishable identity of work
2. Look for new ideas and see if they work
3. Try to get people to start doing something the new way
4. Try to get all the answers yourself so you can establish priority

NOKIA