

Ad hoc On Demand Distance Vector (AODV) Routing Protocol

CS: 647

Advanced Topics in Wireless Networks

Dr. Baruch Awerbuch & Dr. Amitabh Mishra

Department of Computer Science

Johns Hopkins

Reading

- Chapter 6, Sections 6.1-6.3, 6.5 - Ad Hoc Networking, Perkins, Addison Wesley, 2001

Outline

- ❑ Introduction
- ❑ AODV properties
- ❑ Route discovery
 - Flooding
- ❑ Route maintenance
- ❑ Local connectivity management - Hello Messages
- ❑ Summary

Introduction: DSDV vs. AODV

- DSDV broadcasts every change in the network to every node
- When two neighbors enter communication range of each other
 - This results in a network wide broadcast
- Similarly when two nodes drift apart from each other's range -> link breakage
 - Also results in a network wide broadcast
- Local movements have global effects
- In AODV such broadcasts are not necessary
- If a link breakage does not affect on going transmission -> no global broadcast occurs
- Only affected nodes are informed
- Local movements of nodes have local effects
- AODV reduces the network wide broadcasts to the extent possible
- Significant reduction in control overhead as compared to DSDV

Ad Hoc On Demand Distance-Vector (AODV) Routing (1)

- ❑ Reactive or on Demand
- ❑ Descendant of DSDV
- ❑ Uses bi-directional links
- ❑ Route discovery cycle used for route finding
- ❑ Maintenance of active routes
- ❑ Sequence numbers used for loop prevention and as route freshness criteria
- ❑ Provides unicast and multicast communication

Ad Hoc On Demand Distance-Vector (AODV) Routing (2)

- ❑ Whenever routes are not used -> get expired -> Discarded
 - Reduces stale routes
 - Reduces need for route maintenance
- ❑ Minimizes number of active routes between an active source and destination
- ❑ Can determine multiple routes between a source and a destination, but implements only a single route, because
 - Difficult to manage multiple routes between same source/destination pair
 - If one route breaks, its difficult to know whether other route is available
 - Lot of book-keeping involved

AODV Properties (1)

1. AODV discovers routes as and when necessary
 - Does not maintain routes from every node to every other
2. Routes are maintained just as long as necessary
3. Every node maintains its monotonically increasing sequence number -> increases every time the node notices change in the neighborhood topology

AODV Properties (2)

- AODV utilizes routing tables to store routing information
 1. A Routing table for unicast routes
 2. A Routing table for multicast routes

- The route table stores: *<destination addr, next-hop addr, destination sequence number, life_time>*

- For each destination, a node maintains a list of *precursor nodes*, to route through them
 - Precursor nodes help in route maintenance (more later)

- Life-time updated every time the route is used
 - If route not used within its life time -> it expires

AODV - Route Discovery (1)

- ❑ When a node wishes to send a packet to some destination -
 - It checks its routing table to determine if it has a current route to the destination
 - If Yes, forwards the packet to next hop node
 - If No, it initiates a **route discovery** process
- ❑ Route discovery process begins with the creation of a Route Request (RREQ) packet -> source node creates it
- ❑ The packet contains - source node's IP address, source node's current sequence number, destination IP address, destination sequence number

AODV - Route Discovery (2)

- Packet also contains broadcast ID number
 - Broadcast ID gets incremented each time a source node uses RREQ
 - Broadcast ID and source IP address form a unique identifier for the RREQ

- Broadcasting is done via Flooding

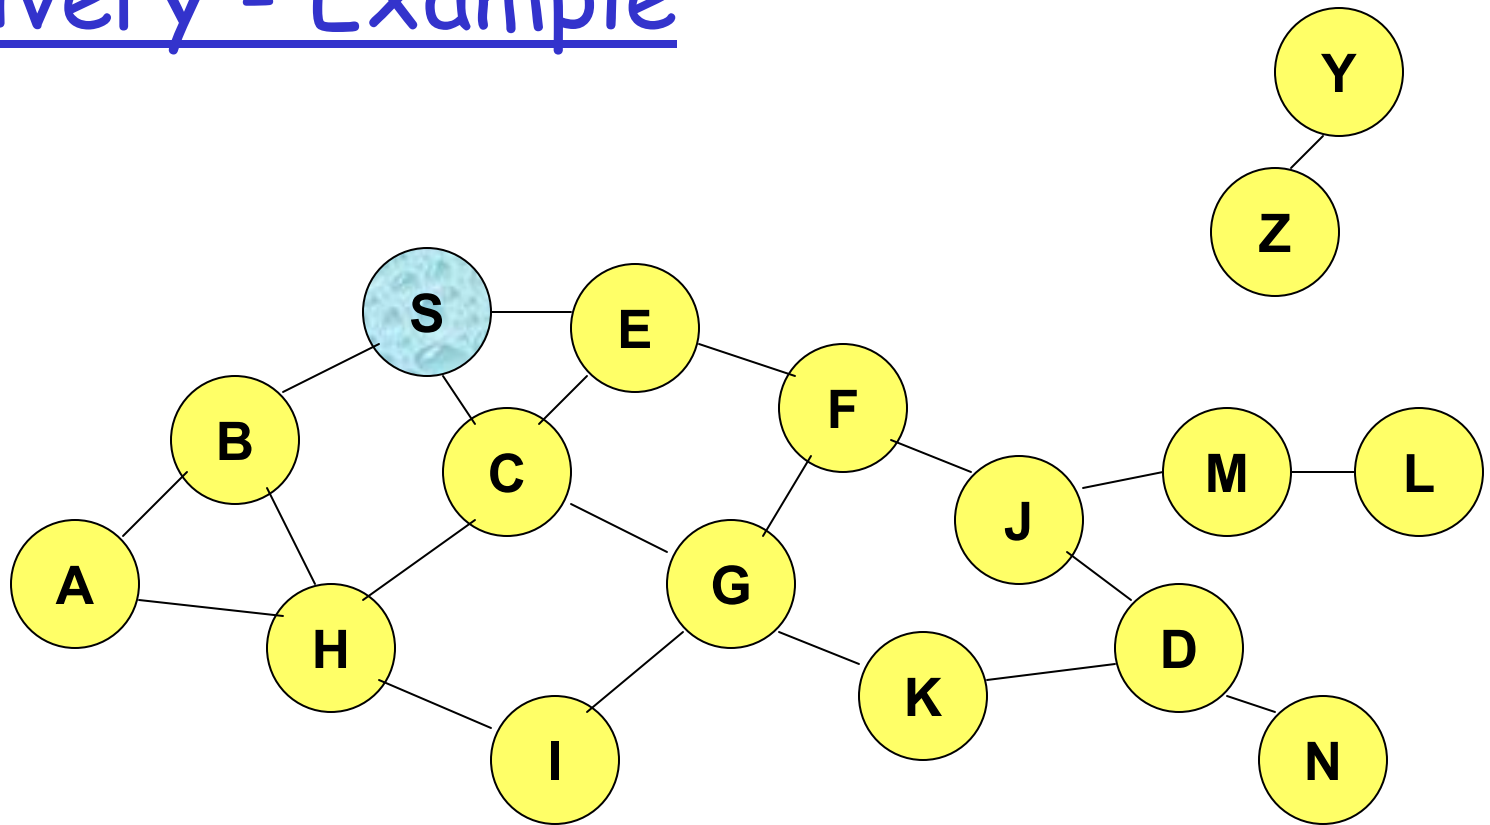
Outline

- ❑ Introduction
- ❑ AODV properties
- ❑ Route discovery
 - Flooding
- ❑ Route maintenance
- ❑ Local connectivity management - Hello Messages
- ❑ Summary


Flooding for Control Packet Delivery

- ❑ Sender S broadcasts a control packet P to all its neighbors
- ❑ Each node receiving P forwards P to its neighbors
- ❑ Sequence numbers help to avoid the possibility of forwarding the same packet more than once
- ❑ Packet P reaches destination D provided that D is reachable from sender S
- ❑ Node D does not forward the packet

Flooding for Control Packet Delivery - Example




Represents a node that has received packet P



Represents that connected nodes are within each other's transmission range

Flooding for Control Packet Delivery

Broadcast transmission




Represents a node that receives packet P for the first time




Represents transmission of packet P

Flooding for Control Packet Delivery




- Node H receives packet P from two neighbors:
potential for collision

Flooding for Control Packet Delivery




- Node C receives packet P from G and H, but does not forward it again, because node C has **already forwarded packet P** once

Flooding for Control packet Delivery




- ❑ Nodes J and K both broadcast packet P to node D
- ❑ Since nodes J and K are **hidden** from each other, their transmissions may collide
=> Packet P may not be delivered to node D at all, despite the use of flooding

Flooding for Control Packet Delivery



- ❑ Node D **does not forward** packet P, because node D is the **intended destination** of packet P

Flooding for Control Packet Delivery




❑ Flooding completed

❑ Nodes **unreachable** from S do not receive packet P (e.g., node Z)

❑ Nodes for which paths go through the destination D also do not receive packet P (example: node N)

Flooding for Control Packet Delivery



- ❑ Flooding may deliver packets to too many nodes (in the **worst case**, all nodes reachable from sender may receive the packet)

Flooding of Control Packets - Summary

- ❑ Many protocols perform flooding of **control** packets e.g. AODV & DSR, instead of **data** packets
- ❑ The control packets are used to discover routes
- ❑ Discovered routes are subsequently used to send data packet(s)
- ❑ Overhead of control packet flooding is **amortized** over data packets transmitted between consecutive control packet floods

Flooding for Data Delivery -Advantages

- Simplicity
- May be more efficient than other protocols when rate of information transmission is low enough that the overhead of explicit route discovery and maintenance incurred by other protocols is relatively higher
 - this scenario may occur, for instance, when nodes transmit **small data packets** relatively infrequently, and many topology **changes occur** between consecutive packet transmissions
- Potentially higher reliability of data delivery
 - Because packets may be delivered to the destination on multiple paths

Flooding for Data Delivery - Disadvantages

- Potentially, very high overhead
 - Data packets may be delivered to too many nodes who do not need to receive them
- Potentially lower reliability of data delivery
 - Flooding uses broadcasting -- hard to implement reliable broadcast delivery without significantly increasing overhead
 - Broadcasting in IEEE 802.11 MAC is unreliable
 - In this example, nodes J and K may transmit to node D simultaneously, resulting in loss of the packet
 - in this case, destination would not receive the packet at all


Route Discovery (Contd.)

- Once an intermediate node receives a RREQ, the node sets up a reverse route entry for the source node in its route table
 - Reverse route entry consists of *<Source IP address, Source seq. number, number of hops to source node, IP address of node from which RREQ was received>*
 - Using the reverse route a node can send a RREP (Route Reply packet) to the source
 - Reverse route entry also contains - life time field
- RREQ reaches destination -> In order to respond to RREQ a node should have in its route table:
 1. Unexpired entry for the destination
 2. Seq. number of destination at least as great as in RREQ (for loop prevention)

Route Discovery (Contd.)


- RREQ reaches destination (contd.)
 - If both conditions are met & the IP address of the destination matches with that in RREQ → the node responds to RREQ by sending a RREP back using **unicasting and not flooding** to the source using reverse path
 - If conditions are not satisfied, then node increments the hop count in RREQ and broadcasts to its neighbors
- Ultimately the RREQ will make to the destination

AODV - Route Discovery - Example




1. Node S needs a route to D

AODV - Route Discovery




1. Node S needs a route to D
2. Creates a **Route Request (RREQ)**
 - Enters D's IP addr, seq#, S's IP addr, seq#, hopcount (=0)

AODV - Route Discovery



1. Node S needs a route to D
2. Creates a **Route Request (RREQ)**
 - Enters D's IP addr, seq#, S's IP addr, seq#, hopcount (=0)
3. Node S broadcasts RREQ to neighbors


AODV - Route Discovery



4. Node A receives RREQ

- Makes a reverse route entry for S
dest=S, nexthop=S, hopcount=1
- It has no routes to D, so it rebroadcasts RREQ


AODV - Route Discovery



4. Node A receives RREQ

- Makes a **reverse route** entry for S
dest=S, nexthop=S, hopcount=1
- It has no routes to D, so it rebroadcasts RREQ

AODV - Route Discovery



5. Node C receives RREQ

- Makes a **reverse route** entry for S
dest=S, nexthop=A, hopcount=2
- It has a route to D, and the seq# for route to D is \geq D's seq# in RREQ


Route Reply in AODV - Comments

- An **intermediate node** (not the destination) may also send a **Route Reply (RREP)** provided that it knows a **more recent path** than the one previously known to sender S

- To determine whether the path known to an intermediate node is more recent, *destination sequence numbers* are used

- The likelihood that an intermediate node will send a Route Reply when using AODV not as high as DSR (Later)
 - A new Route Request by node S for a destination is assigned a higher destination sequence number. An intermediate node which knows a route, but with a smaller sequence number, **cannot send** Route Reply

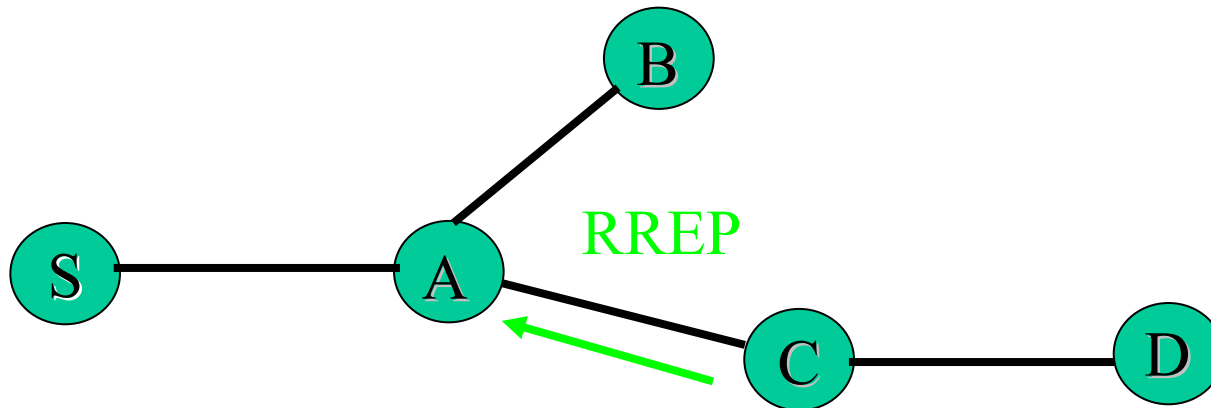
AODV - Route Discovery



5. Node C receives RREQ (Cont.)

- C creates a **Route Reply** (RREP)
Enters D's IP addr, seq#, S's IP addr, hopcount to D (=1)
- Unicasts RREP to A


AODV - Route Discovery



5. Node C receives RREQ (Cont.)

- C creates a **Route Reply** (RREP)
 - Enters D's IP addr, seq#, S's IP addr, hopcount to D (=1)
- Unicasts RREP to A

AODV - Forward Path Setup



5. Node A receives RREP

- Makes a **forward route** entry to D
dest=D, nexthop=C, hopcount=2
- Unicasts RREP to S


Forward Path Setup (1)

- When a node determines that it has a current route to respond to RREQ i.e. has a path to destination - It creates RREP (Route Reply)
- RREP contains <IP address of source and destination>
 - If RREP is being sent by destination
 - The RREP will also contain the <current sqn # of destination, hop-count=0, life-time>
 - If RREP is sent by an intermediate node
 - RREP will contain its record of the <destination sequence number, hop-count=its distance to destination, its value of the life-time>

Forward Path Setup (2)

- When an intermediate node receives the RREP, it sets up a **forward path** entry to the destination in its route table
 - Forward path entry contains
 - <IP Address of destination, IP address of node from which the entry arrived, hop-count to destination, life-time>
 - To obtain its distance to destination i.e. hop-count, a node increments its distance by 1
 - If route is not used within the life time, its deleted
- After processing the RREP, the node forwards it towards the source

AODV - Forward Path Setup




5. Node S receives RREP

- Makes a **forward route** entry to D
dest=D, nexthop =A, hopcount = 3

Receipt of Multiple RREP

- A node may receive multiple RREP for a given destination from more than one neighbor
 - The node only forwards the first RREP it receives
 - May forward another RREP if that has greater destination sequence number or a smaller hop count
 - Rest are discarded -> reduces the number of RREP propagating towards the source
- Source can begin data transmission upon receiving the first RREP


AODV - Data Delivery



5. Node S receives RREP

- Makes a **forward route** entry to D
dest=D, nexthop =A, hopcount = 3
- Sends data packet on route to D


Example 2: Route Requests in AODV



Represents a node that has received RREQ for D from S


Route Requests in AODV

Broadcast transmission




.....→ Represents transmission of RREQ

Route Requests in AODV




← Represents links on Reverse Path

Reverse Path Setup in AODV




- Node C receives RREQ from G and H, but does not forward it again, because node C has **already forwarded RREQ** once

Reverse Path Setup in AODV




Reverse Path Setup in AODV




- Node D **does not forward** RREQ, because node D is the **intended target** of the RREQ

Route Reply in AODV



 Represents links on path taken by RREP

Forward Path Setup in AODV




Forward links are setup when RREP travels along the reverse path



Represents a link on the forward path

Data Delivery in AODV



Routing table entries used to forward data packet.

Route is *not* included in packet header.

Timeouts

- A routing table entry maintaining a **reverse path** is purged after a timeout interval
 - timeout should be long enough to allow RREP to come back

- A routing table entry maintaining a **forward path** is purged if *not used* for a *active_route_timeout* interval
 - if no is data being sent using a particular routing table entry, that entry will be deleted from the routing table (even if the route may actually still be valid)

Outline

- Introduction
- Flooding
- AODV properties
- Route discovery
- Route maintenance - link failures
- Local connectivity management

Link Failure Reporting

- A neighbor of node X is considered **active** for a routing table entry if the neighbor sent a packet within **active_route_timeout** interval and was forwarded using that entry

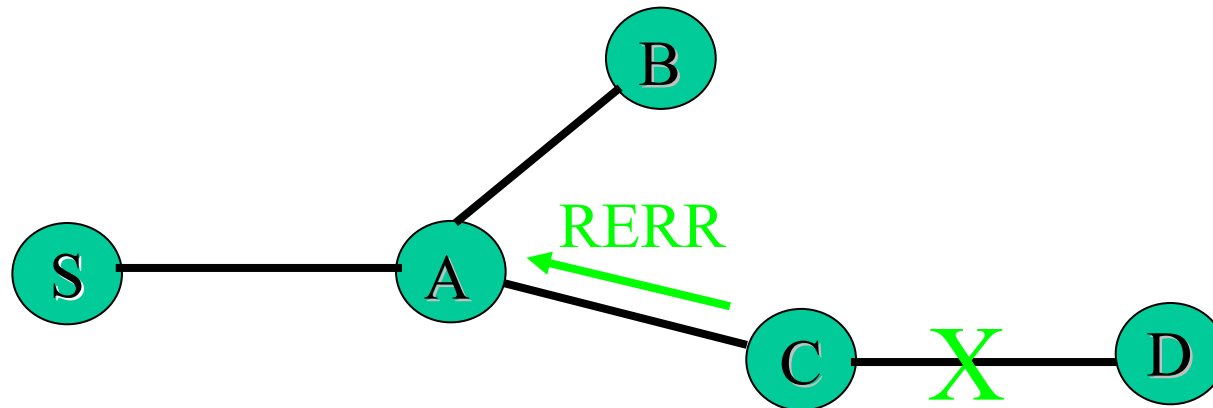
- If a source node moves, a new route discovery process is initiated

- If intermediate nodes or the destination move ->
 - The next hop links break resulting in link failures
 - Routing tables are updated for the link failures
 - All **active** neighbors are informed by **RERR** message

Route Maintenance - RERR


- RERR is initiated by the node upstream (closer to the source) of the break
 - Its propagated to all the affected destinations
 - RERR lists all the nodes affected by the link failure -> Nodes that were using the link to route messages (precursor nodes)
 - When a node receives an RERR, it marks its route to the destination as invalid -> Setting distance to the destination as infinity in the route table
- When a source node receives an RRER, it can reinitiate the route discovery

AODV - Route Maintenance- Example



1. Link between C and D breaks
2. Node C invalidates route to D in route table
3. Node C creates **Route Error** message
 - Lists all destinations that are now unreachable
 - Sends to upstream neighbors


AODV - Route Maintenance - Example



4. Node A receives RERR

- ❑ Checks whether C is its next hop on route to D
- ❑ Deletes route to D (makes distance \rightarrow infinity)
- ❑ Forwards RERR to S

AODV - Route Maintenance- Example



5. Node S receives RERR

- ❑ Checks whether A is its next hop on route to D
- ❑ Deletes route to D
- ❑ Rediscovered route if still needed

Route Error

- ❑ When node X is unable to forward packet P (from node S to node D) on link (X,Y) , it generates a RERR message
- ❑ Node X increments the destination sequence number for D cached at node X
- ❑ The incremented sequence number N is included in the RERR
- ❑ When node S receives the RERR, it initiates a new route discovery for D using destination sequence number at least as large as N

Destination Sequence Number

- Continuing from the previous slide ...
- When node D receives the route request with destination sequence number N , node D will set its sequence number to N , unless it is already larger than N

Outline


- Introduction
- Flooding
- AODV properties
- Route discovery
- Route maintenance - link failures
- Local connectivity management - Hello Messages

Link Failure Detection

- *Hello* messages: Neighboring nodes periodically exchange hello message
- Absence of hello message is used as an indication of link failure
- Alternatively, failure to receive several MAC-level acknowledgements may be used as an indication of link failure


Why Sequence Numbers in AODV

- ❑ To avoid using old/broken routes
 - To determine which route is newer
- ❑ To prevent formation of loops
 - A had a route to D initially
 - Assume that A does not know about failure of link C-D because RERR sent by C is lost



- Now C performs a route discovery for D. Node A receives the RREQ (say, via path C-E-A)
- Node A will reply since A knows a route to D via node B
- Results in a loop (for instance, C-E-A-B-C)

Why Sequence Numbers in AODV



- Loop C-E-A-B-C
- But because of usage of sequence number, A will not use the route A-B-C, because the sequence numbers will be lower than what A receives from A

Optimizations

- ❑ Route Requests are initially sent with small Time-to-Live (TTL) field, to limit their propagation
 - DSR also includes a similar optimization
- ❑ If no Route Reply is received, then larger TTL tried

AODV: Optimizations

- Expanding Ring Search
 - Prevents flooding of network during route discovery
 - Control Time to Live (TTL) of RREQ to search incrementally larger areas of network
 - Advantages: Less overhead when successful
 - Disadvantages: Longer delay if route not found immediately

AODV: Optimizations (Cont.)

□ Local Repairs

- Repair breaks in active routes locally instead of notifying source
- Use small TTL because destination probably hasn't moved far
- If first repair attempt is unsuccessful, send RERR to source
- Advantage: repair links with less overhead, delay and packet loss
- Disadvantages: longer delay and greater packet loss when unsuccessful

AODV: Summary (1)

- ❑ Routes need not be included in packet headers (DSR does it)
- ❑ Nodes maintain routing tables containing entries only for routes that are in active use
- ❑ At most one next-hop per destination maintained at each node
 - DSR may maintain several routes for a single destination
- ❑ Unused routes expire even if topology does not change

AODV: Summary (2)

- ❑ Reactive/On - demand
- ❑ Sequence numbers used for route freshness and loop prevention
- ❑ Route discovery cycle
- ❑ Optimizations can be used to reduce overhead and increase scalability