

Routing in reti ad hoc

Sistemi Wireless, a.a 2011/2012

Un. of Rome "La Sapienza"

Chiara Petrioli[†]

[†] *Department of Computer Science – University of Rome "Sapienza" – Italy*



- A wireless multi-hop infrastructure-less network whose devices act as source/ destination of messages & as relay for packets generated by a node s and addressed to a node z (iff they are on a s - z route)
- Pros: No need for infrastructure \rightarrow low cost, enables communication where it is usually not needed or is not viable
- Must be: Self-organizing, self-configuring, self-maintaining



- Collaboration between users in office environments
- Disaster recovery applications
- Military networks
- Personal Area Networks
- Home Networking
- Wireless Sensor Networks (WSNs)
- Inter-vehicular communication



- Highly dynamic networks → device mobility, energy saving sleep/awake modes
- Need for low energy/resource-consuming, simple protocols
- Bandwidth and resource constrained environment
- Traffic:
 - All-pairs in general ad hoc networks, from sensors to sink(s) in sensor networks
 - In many case not high
- Scale: Application dependent
 - 10-100 nodes in traditional ad hoc networks
 - 1000-10000 in sensor networks



- Highly dynamic networks → due to device mobility (only in some specific applications), to the fact the active node set changes in time for sake of energy saving (always to be considered)
- Need to design low energy/resource-consuming, simple protocols → very critical, energy consumption a real bottleneck
- Traffic from sensors to sink(s)
- Scalability is a major issue
- Code must be simple (small storage capability, very simple, inexpensive, resource constrained devices)
- First solutions we will see for traditional ad hoc networks do not scale to high numbers and are not energy-saving



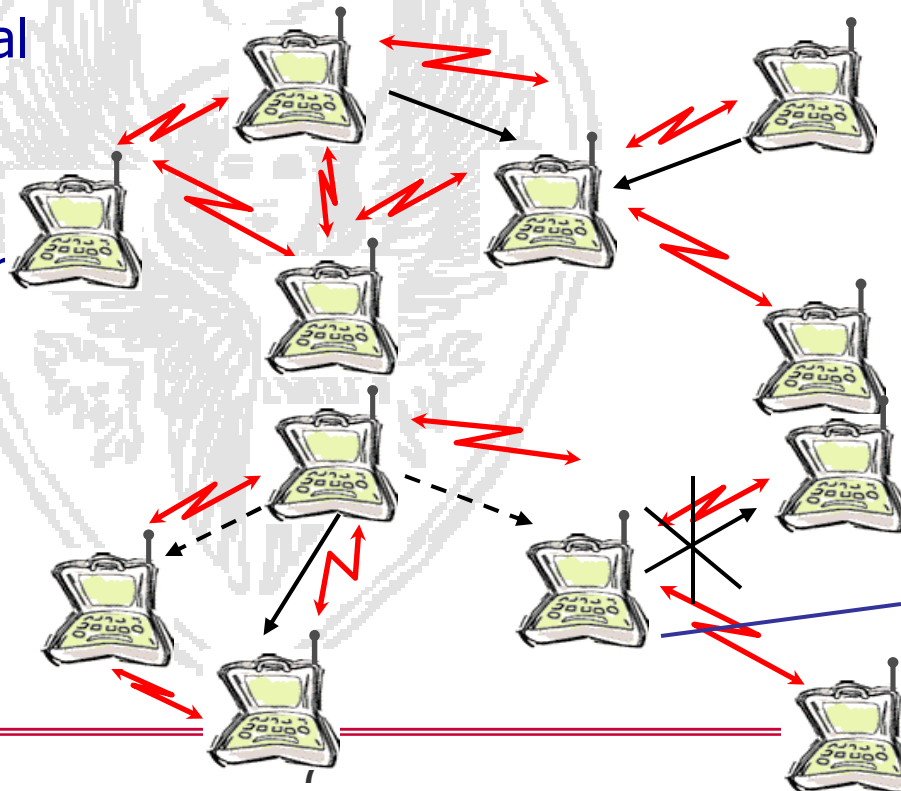
SAPIENZA
UNIVERSITÀ DI ROMA



***Medium Access Control nelle reti
ad hoc
Un esempio***



- Si usano approcci CSMA-like, e.g. CSMA/CA
 - Perche' non TDMA like?
- Perche' non CSMA/CD? nodi non ricevono/trasmettono contemporaneamente
 - Hidden terminal
 - Exposed terminal



*I due
nodi che
Tx non
si ascoltano*

*Il nodo
Potrebbe
trasmettere*




- Basato sul CSMA/CA
- Prima di trasmettere un frame una stazione fa il sensing del canale
- Se il canale e' libero per un intervallo superiore al Distributed InterFrame Space (DIFS) la stazione trasmette
- Altrimenti (canale gia' occupato) si aspetta la fine della trasmissione corrente + un intervallo casuale detto *backoff timer*.
 - Il backoff timer viene decrementato solo quando il canale e' idle e viene congelato quando invece il canale e' occupato (e' riattivato quando il canale e' libero per un DIFS) **DOMANDA:PERCHE'?**
 - La stazione trasmette quando il backoff timer raggiunge il valore zero.
 - Il valore del backoff timer e' scelto casualmente all'interno di una finestra di CW slots. Al primo tentativo CW e' settato al valore minimo previsto da standard settato a 16.
- Come fa la MS trasmittente a sapere se il frame e' stato ricevuto con successo (collisioni si possono verificare per trasmissioni simultanee o per effetto del terminale nascosto)? Viene inviato dal ricevente un ACK esplicito alla fine della corretta ricezione del frame, dopo aver atteso per un tempo pari allo Short InterFrame Space (SIFS), SIFS < DIFS
- Nel caso di collisione si aspetta un tempo random e si prova a ritrasmettere → backoff esponenziale (CW viene raddoppiata ad ogni ritrasmissione fino ad un massimo di 1024 slots)




- Per alleviare gli effetti sulle prestazioni del fenomeno del terminale nascosto il DCF puo' usare un meccanismo di *virtual carrier sensing*
- prima di trasmettere un frame una stazione fa il sensing del canale, aspetta un tempo DIFS e se possibile trasmette un pacchetto (corto) di controllo denominato RTS (Request to Send), annunciando che trasmettera' un frame di una certa lunghezza a quella destinazione
- Il pacchetto RTS contiene un campo NAV (Network Allocation Vector) con la stima esatta del tempo necessario dalla ricezione dell'RTS al completamento della trasmissione dell'ACK relativo alla conferma della corretta ricezione del frame → ricevendo l'RTS tutte le stazioni vicine al trasmittente sanno di non poter usare il canale per questo tempo
- Il destinatario risponde se riceve correttamente l'RTS ed il canale e' libero con un CTS (Clear to Send), che a sua volta contiene un campo NAV, dicendo 'OK, ricevero' il tuo messaggio' → tutti I vicini della destinazione sapranno che il canale e' occupato, per quanto tempo e si astengono dal tramettere in tale tempo
- Quale il vantaggio del meccanismo RTS/CTS? Che problemi risolve? Che problemi (se puo') crea?





- Per alleviare gli effetti sulle prestazioni del fenomeno del terminale nascosto il DCF può usare un meccanismo di *virtual carrier sensing*



- Il problema del terminale nascosto è causato dal fatto che il terminale nascosto non può sentire il canale occupato dal terminale visibile. Il problema del terminale nascosto può essere risolto utilizzando il meccanismo di *virtual carrier sensing*. In questo meccanismo, il terminale nascosto può sentire il canale occupato dal terminale visibile attraverso il terminale visibile. Il terminale nascosto può sentire il canale occupato dal terminale visibile attraverso il terminale visibile. Il terminale nascosto può sentire il canale occupato dal terminale visibile attraverso il terminale visibile.
- Il vantaggio del meccanismo RTS/CTS? Che problemi risolve? Che problemi (se può) crea?



- Intra-AS routing in the Internet
 - Link State Approaches
(info on the topology graph gathered at nodes which run shortest path algorithms-Dijkstra- to decide the routes to the different destinations –e.g. OSPF routing protocol)
 - Distance Vector approaches (e.g. RIP)



Given a graph $G=(N,A)$ and a node s find the shortest path from s to every node in N .

A shortest walk from s to i subject to the constraint that the walk contains at most h arcs and goes through node s only once, is denoted **shortest($\leq h$) walk and its length is D_i^h .**

Bellman-Ford rule:

Initiatilization $D_s^h=0$, for all h ; $c_{i,k} = \text{infinity}$ if (i,k) NOT in A ; $c_{k,k} = 0$;
 $D_i^0 = \text{infinity}$ for all $i \neq s$

Iteration:

$$D_i^{h+1} = \min_k [c_{i,k} + D_k^h]$$

Assumption: non negative cycles (this is the case in a network!!)

**The Bellman-Ford algorithm first finds the one-arc shortest walk lengths, then the two-arc shortest walk length, then the three-arc...etc.
→distributed version used for routing**



$$D_i^{h+1} = \min_k [c_{i,k} + D_k^h]$$

Can be computed locally.

What do I need?

For each neighbor k , I need to know

-the cost of the link to it (known info)

-The cost of the best route from the neighbor k to the destination
(←this is an info that each of my neighbor has to send to me via messages)

In the real world: I need to know the best routes among each pair of nodes → we apply distributed Bellman Ford to get the best route for each of the possible destinations



iterative:

- continues until no nodes exchange info.
- *self-terminating*: no "signal" to stop

asynchronous:

- nodes need *not* exchange info/iterate in lock step!

Distributed, based on local info:

- each node communicates *only* with directly-attached neighbors

Distance Table data structure

each node has its own

- row for each possible destination
- column for each directly-attached neighbor to node
- example: in node X, for dest. Y via neighbor Z:

Cost associated to the (X,Z) link




$$D^X(Y,Z) = \text{distance from X to Y, via Z as next hop}$$
$$= c(X,Z) + \min_w \{D^Z(Y,w)\}$$

Info maintained at Z. Min must be communicated



Link cost changes:

- good news travels fast
- ***bad news travels slow*** - "count to infinity" problem!





Which is the problem here?

the info exchanged by the protocol!! ‘the best route to X I have has the following cost...’ (no additional info on the route)

A Roman example...

-assumption: there is only one route going from Colosseo to Altare della Patria: Via dei Fori Imperiali. Let us now consider a network, whose nodes are Colosseo., Altare della Patria, Piazza del Popolo





The Colosseo. and Alt. Patria nodes exchange the following info

- Colosseo says ‘the shortest route from me to P. Popolo is 2 Km’
- Alt. Patria says ‘the shortest path from me to P. Popolo is 1Km’

Based on this exchange from Colosseo you go to Al. Patria, and from there to Piazza del Popolo OK Now due to the big dig they close Via del Corso (Al. Patria—P.Popolo)

• Al. Patria thinks ‘I have to find another route from me to P.Popolo. Look there is a route from Colosseo to P.Popolo that takes 2Km, I can be at Colosseo in 1Km → I have found a 3Km route from me to P.Popolo!!’ Communicates the new cost to Colosseo that updates ‘OK I can go to P.Popolo via Al. Patria in 4Km’

VERY WRONG!! Why is it so? I didn’t know that the route from Colosseo to P.Popolo was going through Via del Corso from Al.Patria to P.Popolo (which is closed)!!



- Numero massimo di hop count
 - L'impatto di loop può essere limitato tramite un campo del pacchetto che tenga conto degli hop attraversati. I pacchetti che sono stati in rete per un numero di hop superiore ad una soglia prefissata vengono scartati.
 - se la dimensione della rete è limitata si può limitare (come si fa in RIP) il numero di hop `num_max_hops` al diametro della rete (pari ad es. al valore 15). Se, come in RIP, il costo dei link è pari a 1 ponendo infinito=16 il protocollo opera correttamente e converge rapidamente nel caso in cui si verificano problemi di tipo count to infinity.
- Split horizon con poison reverse
 - Si limita l'informazione trasmessa. Se A usa informazioni ricevute da B per scegliere la rotta verso la destinazione D (A passa tramite B per raggiungere D), A non comunicherà un costo per raggiungere D a B oppure (variante Poison Reverse) comunicherà un costo infinito.
 - ✓ La tecnica assume un mezzo non broadcast ed un vicinato noto
 - ✓ Non risolve tutti i loop
- Trigger Updates (velocizzazione della convergenza)
 - Anziché trasmettere periodicamente aggiornamenti tali aggiornamenti possono essere immediatamente inviati nel caso in cui si individuino cambiamenti.



- Multi-hop path routing capability
- Dynamic topology maintenance
- “No loops”
- Minimal control overhead
- Low processing overhead
- Self-starting



- Proactive

- Based on traditional distance-vector and link-state protocols
- Each node maintains route to each other network node
- Periodic and/or event triggered routing update exchange
- Higher overhead in most scenarios
- Longer route convergence time
- Examples: DSDV, OLSR



- Proactive, distance vector approach (uses distributed asynchronous Bellman Ford). Updates on route costs transmitted periodically or when significant new information is available.
- Difference wrt Bellman Ford: in ad hoc networks there are frequent changes in the topology, solutions must try to avoid loops (approaches such as Poison reverse non effective in broadcast channels, we seek solutions which are simple and fully distributed)
- Metrics: fresh routes better than stale routes, number of hops used to select among the fresh routes
- How to identify fresh routes? By means of sequence numbers identifying the freshness of the communicated information. When changes occur, the sequence number increase.



- Periodically destination nodes transmit updates with a new sequence number (and such updates are propagated by the other nodes).
- Updates periodically sent by nodes contain information on the costs to achieve the different destinations and the freshness of the route
- Data broadcast include multiple entries each with:
 - Destination address
 - Number of hops required to reach the destination
 - Sequence number of the information received regarding that destination as originally stamped by the destination
- In the header the data broadcast also include:
 - Address (HW address/Net address) of the sender of the message
 - Sequence number created by the transmitter
- Two types of updates (full dump or incremental-only changes- to decrease bandwidth consumption.



- How can the costs be modified? Cost=number of hops, target: using fresh routes as short as possible → a link cost changes from 1 to inf and from inf to 1
- How do we detect that a link is 'broken'? At layer 2 (no hello messages received for some time, or attempts to retransmit a frame exceeds the MAC protocol threshold) or at layer 3 (do not receive periodic updates by a neighbor)
- Link cost increase (1 → inf):
 - The nodes incident to that link (A,B) discover it (see above)
 - Routes going through that link get assigned an inf cost in nodes A and B routing tables
 - A new sequence number is generated by the mobile node. Mobile nodes different from the destination use odd SN, the destination even SN.
 - Updates with routes with infinite cost are immediately transmitted by nodes
- Link cost decrease (inf → 1):
 - Immediately transmits updates



- When a node receives updates it sees if costs to reach the different destinations can be improved:
 - routes with more recent sequence numbers to a given destination are used
 - if more routes available with the same SN the shortest is used
- Newly recorded routes are scheduled for immediate advertisement (inf → finite value)
- Routes with improved metric are scheduled for advertisement at a time which depends on the estimated average settling time for routes to that particular destination (based on previous history) → delayed advertisements to decrease the overall overhead
- As soon as a route cost changes the node may delay informing its neighbors but immediately starts using the new information for its forwarding



- Assuming routing tables are stable and a change occurs
 - let $G(x)$ denotes the routes graph from the sources to x BEFORE the change (assume no loop)
 - change occurs at i when 1) the link from i to its parent $p(i)$ in $G(x)$ breaks $\rightarrow i$ sets to inf that route (no loop can occur) 2) node i receives from one of its neighbors k a route to x with sequence number SN_k^x and metric m which is selected to replace the current metric i has to reach x (this occurs only if SN_k^x greater than the previous SN I had stored SN_i^x or if the two SN are equal but the new route has a lower hop cost \rightarrow in the first case if selecting k leads to a loop then $SN_k^x \leq SN_i^x$ which is a contradiction, in the second case the claim comes from the observation reduction in the costs do not bring to loops).



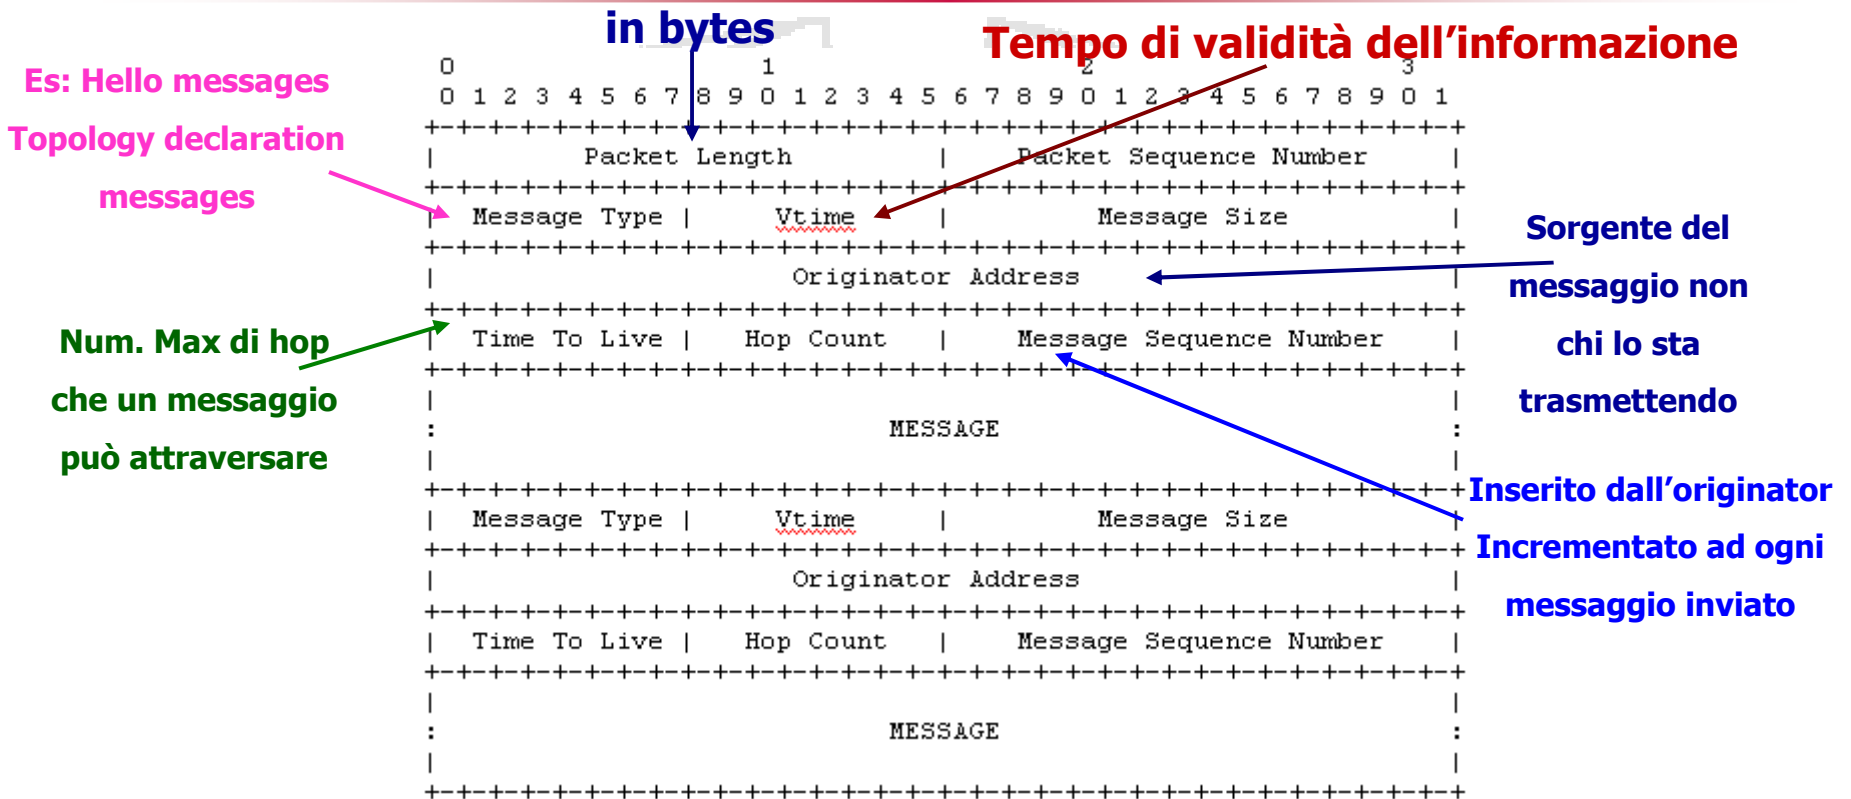
- Optimized Link State Routing (OLSR) is a link state protocol for MANETs
 - suited for large and dense ad hoc networks
- The key concept is to decrease the overhead of flooding by means of identifying a subset of nodes (multipoint relays) in charge of forwarding the information during the flooding process
 - **Multipoint relay Y**: a node selected by at least one of its 1-hop neighbors (say node X) to relay all valid broadcast information it receives from X (the broadcast information is valid if it is not expired and not duplicate).
 - ✓ $MPR(X)$ =set of multipoint relays of node X
 - ✓ neighbors of node X which are not in $MPR(X)$ receive and store the broadcast messages transmitted by X but DO NOT retransmit them
 - A node X which has selected a neighbor Y as multi-point relay is called a **multipoint relay selector** of node Y
- Requires only partial link state to be flooded
 - links from MPR to their selectors must be declared
 - ✓ enough to ensure routes to each destination can be found
 - additional link state information MAYBE advertised



- The protocol is fully distributed
- Proactive approach: routes are always available when needed
- Other features:
 - Time between updates can be tuned to increase reactivity to topological changes
 - does not require reliable transmission
 - ✓ some losses are tolerated ← needed info are periodically transmitted
 - OLSR control packets are embedded in UDP datagrams
 - ✓ sequenced delivery of the messages is not needed ← proper reconstruction of the sequence is possible due to use of sequence numbers
 - support to other MANET related issues
 - ✓ Sleep mode operation
 - ✓ Multicasting



Packet format



Ogni nodo mantiene triple <originator address, sequence number, se il messaggio e' stato già inviato> relative a messaggi ricevuti recentemente. In questo modo i duplicati vengono scartati

Sono scartati anche pacchetti con TTL a zero o non coerenti con le specifiche



- Gli hello messages servono a
 - verificare se i link sono ancora su (link sensing)
 - ✓ se non si riceve un hello dal vicino entro un timeout si assume che il link non sia più attivo
 - ✓ scambiare con i vicini il proprio elenco di vicini
 - consente ai nodi di aggiornare le informazioni relative al loro vicinato a due hop
 - » Serve per poter individuare i multipoint relay



- Each node X selects its MPRs among its one hop neighbors
- The set is selected to cover node X 2-hop neighborhood
 - MPR(X) is an arbitrary subset of node X one hop neighbors such that each node z in node X's two hop neighborhood have a neighbor in MPR(X)
 - ✓ can be selected with a greedy protocol
 - $MPR(X)=null$, $C(X)=$ vicinato a due hop da X
 - per ciascun vicino Y del nodo X si calcola il suo degree D (Y) escludendo il vicinato a 1-hop di X e X stesso
 - si inserisce nell'insieme MPR(X) un nodo Y se è l'unico vicino di X in grado di coprire un nodo a due hop da X
 - » $C(X)=C(X) \setminus \{\text{nodi coperti da Y}\}$
 - fino a quando $C(X)=null$
 - » inserisci in MPR(X) il vicino di X che consente di coprire più nodi rimasti scoperti in C(X) (a parità si seleziona in base al degree D)
 - » $C(X)=C(X) \setminus \{\text{nodi coperti dal vicino selezionato}\}$
 - The smaller MPR(X) the less control overhead exchanged



Upon receiving a message m at node Y

- If the received message is not a duplicate, is valid and has a non zero TTL
 - if it is received by an MPR selector of Y
 - ✓ retransmit m
 - *reduce by one the message TTL*
 - *increase by one the message hop count*
 - *broadcast on all node Y interfaces*
 - ✓ update or create the entry for the message in the duplicate set (→upon receiving the same message the node can identify it was already received and retransmitted → can be discarded)



- Information on topology are disseminated to all the network nodes
 - in an efficient way
 - ✓ exploiting the backbone of multipoint relays
 - ✓ limiting as much as possible topology information
- Each node then locally runs a shortest path algorithm to determine paths to the different destination
 - fills a routing table
 - upon reception of a data packet forwarding is performed according to the routing table



SAPIENZA
UNIVERSITÀ DI ROMA





- Proactive protocols are costly in terms of overhead (the bandwidth and energy are critical resources)
- The cost of maintaining routes updated may not make sense in an environment in which
 - Medium-high mobility
 - Medium-high dynamicity (awake/asleep states)Motivate frequent changes in the the optimum route (requiring updates) while
 - Traffic is generally low (so the cost of maintaining always updated routes is not balanced by their use)If this is the scenario what can we do?



- Reactive (on-demand)
 - Source build routes on-demand by “flooding”
 - Maintain only active routes
 - Route discovery cycle
 - Typically, less control overhead, better scaling properties
 - Drawback: route acquisition latency
 - Example: AODV, DSR

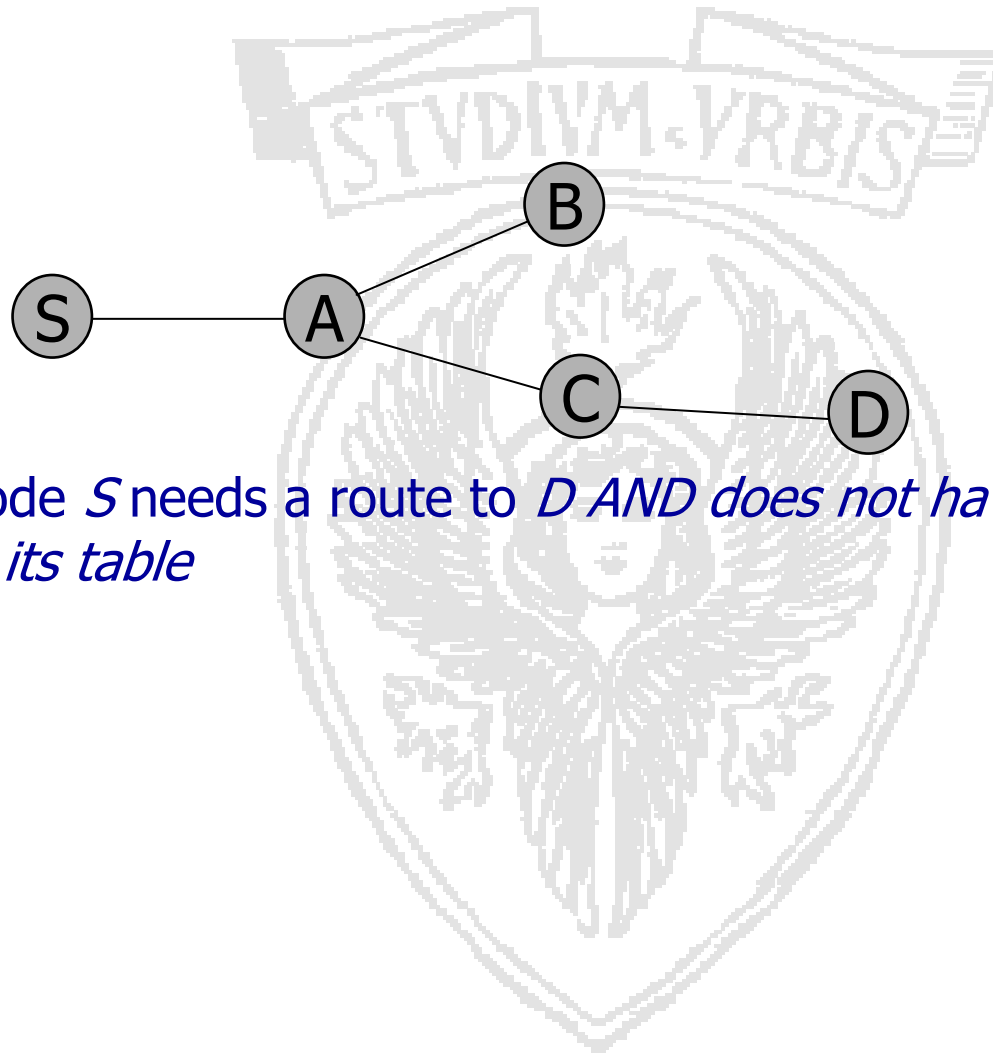


Ad hoc On-Demand Distance Vector Routing

- AODV: Reactive (nodes that do not lie on active paths neither maintain any routing information nor participate in any periodic routing table exchange; a node does not have to discover/maintain a route to a destination till it is on a path to it or has to send messages to it)
- *Route discovery cycle* used for route finding
- Maintenance of *active* routes
- Sequence numbers used for loop prevention and as route freshness criteria
- Descendant of DSDV (standard distance vector approach mapped to ad hoc networks), in AODV no periodic updates but pure on-demand operation.
- Provides unicast and multicast communication




AODV: Route Discovery



1. Node *S* needs a route to *D* AND does not have routing info for it in its table




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), broadcast ID




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), broadcast ID
3. Node *S* broadcasts RREQ to neighbors




AODV: Route Discovery



4. Node *A* receives RREQ
 - Makes reverse route entry for *S*
dest=*S*, nexthop=*S*, hopcnt=1, expiration time for reverse path
Source node, Source node SN, D, broadcastID also maintained
 - It has no route to *D*, so it rebroadcasts RREQ (hopcount increased)
 - If it has already received that request (same source and broadcast ID) it discards the RREQ
 - if it knows a valid path to *D* it will send back a reply to the source



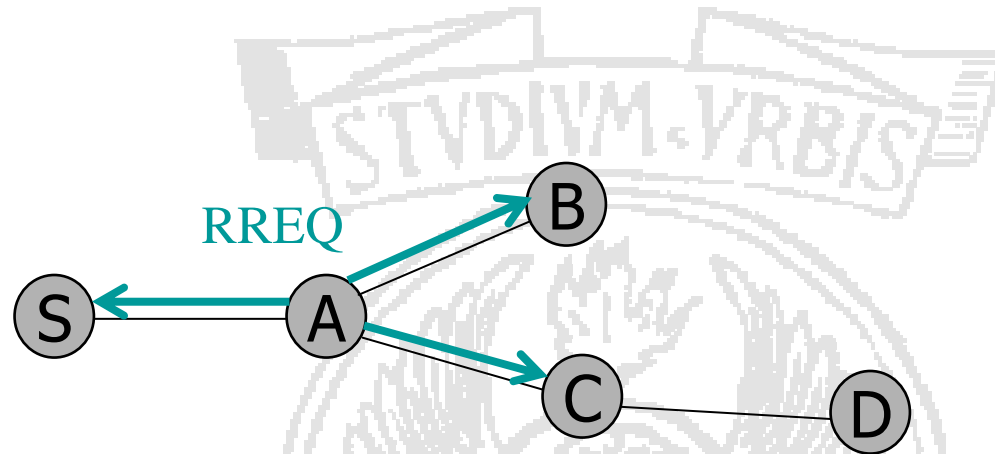
AODV: Route Discovery



4. Node *A* receives RREQ
 - Makes reverse route entry for *S*
dest=*S*, nexthop=*S*, hopcnt=1
 - It has no route to *D*, so it rebroadcasts RREQ



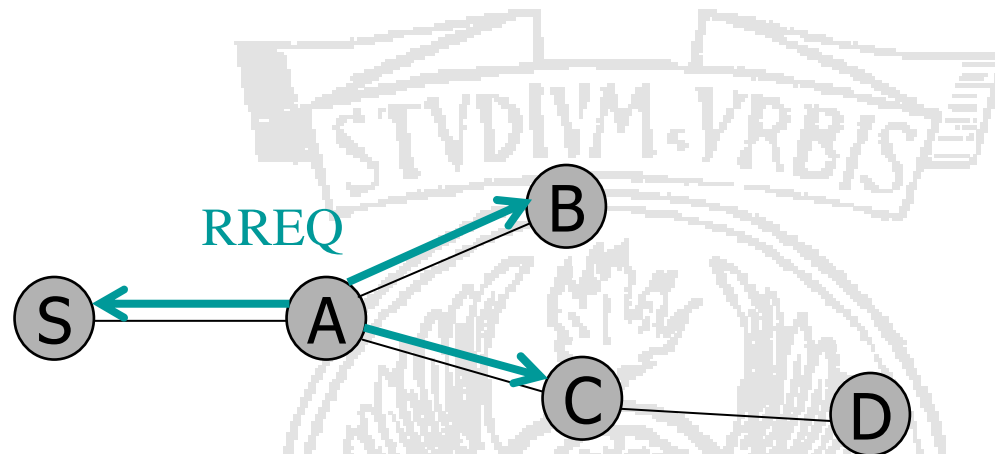
AODV: Route Discovery



5. Node *C* receives RREQ
 - Makes reverse route entry for *S*
dest=*S*, nexthop=*A*, hopcnt=2
 - It has a route to *D*, and the seq# for route to *D* is $\geq D$'s seq# in RREQ



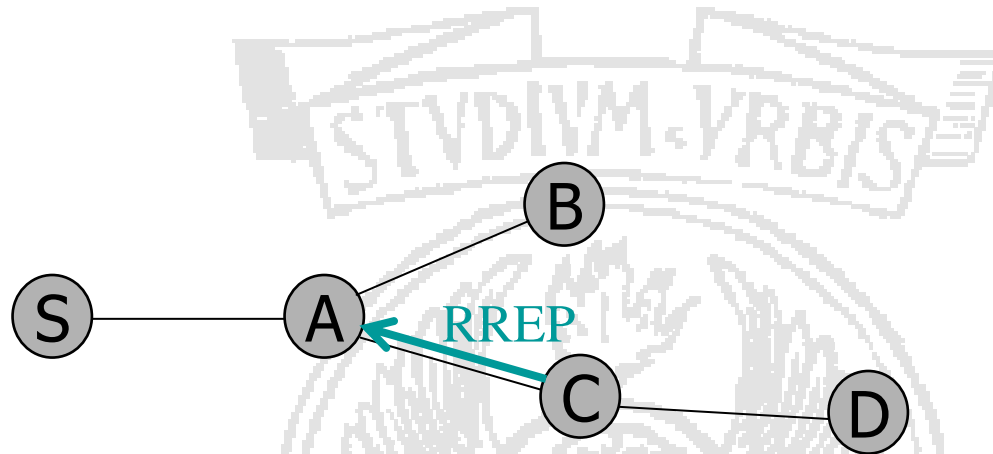
AODV: Route Discovery



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), lifetime of the forward route
 - Unicasts RREP to *A*



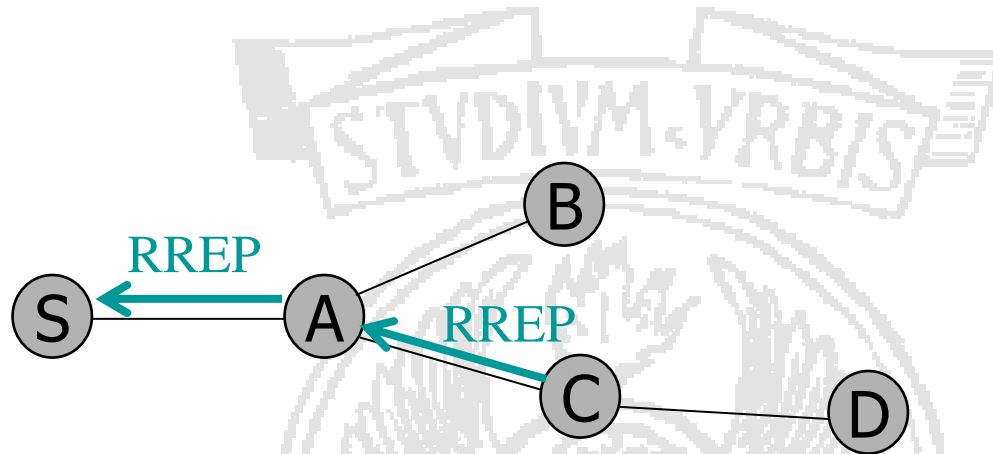
AODV: Route Discovery



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: Route Discovery



6. Node *A* receives RREP
 - Makes forward route entry to *D*
dest = *D*, nexthop = *C*, hopcount = 2
 - Unicasts RREP to *S*



AODV: Route Discovery




7. Node *S* receives RREP
 - Makes forward route entry to *D*
dest = *D*, nexthop = *A*, hopcount = 3

Also the latest SN of the destination is updated when receiving the RREP

Nodes not along the path determined by the RREP will timeout after ACTIVE_ROUTE_TIMEOUT (3000ms) and will delete the reverse pointer



AODV: Route Discovery

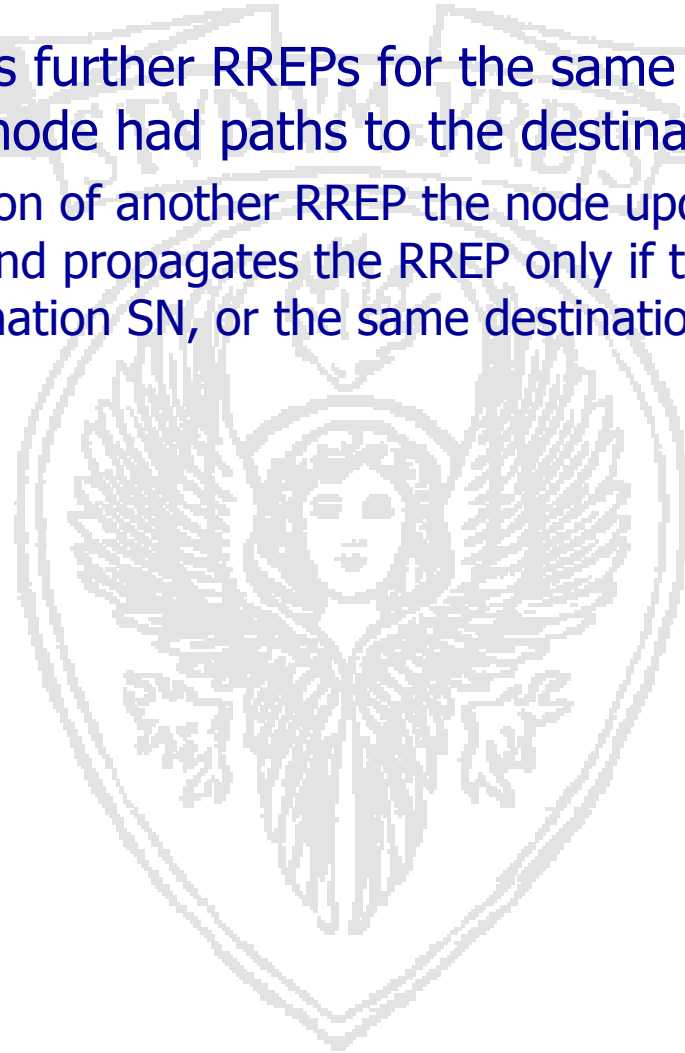


7. Node *S* receives RREP
 - Makes forward route entry to *D*
dest = *D*, nexthop = *A*, hopcount = 3
 - Sends data packets on route to *D*



What if....

- A node receives further RREPs for the same request? (e.g. more neighbors of a node had paths to the destination in cache)?
 - upon reception of another RREP the node updates its routing information and propagates the RREP only if the RREP contains either a greater destination SN, or the same destination SN with a smaller hopcount



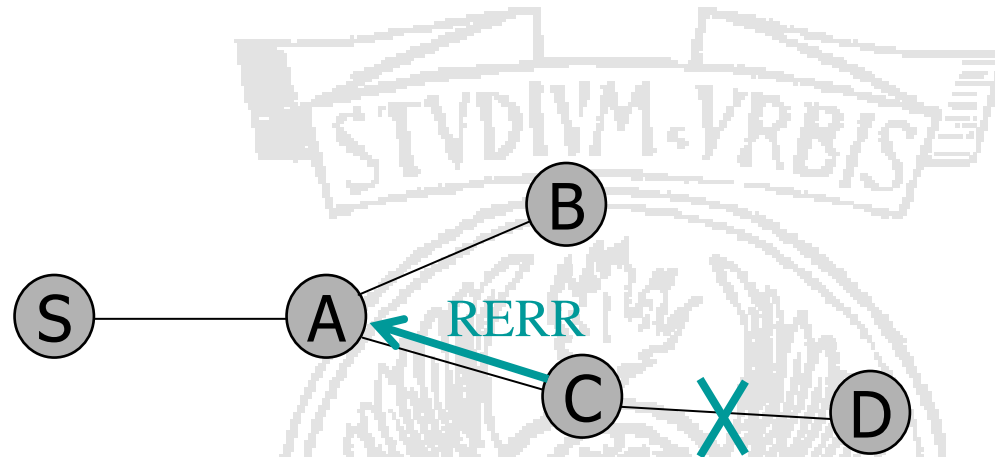


Other info maintained

- Each node maintains the list of active neighbors, neighbors sending to a given destination through it
 - useful for route maintenance
- Routing table entries: dest,next hop, hopcount, dest SN, active neighbors for this route, expiration time for route table entry (updates each time the route is used for transmitting data → routes entries are maintained if the route is active)




AODV: Route Maintenance



1. Link between *C* and *D* breaks
2. Node *C* invalidates route to *D* in route table
3. Node *C* creates Route Error (RERR) message
 - Lists all destinations which are now unreachable
 - Sends to upstream neighbors
 - Increases of one the SN of the destination



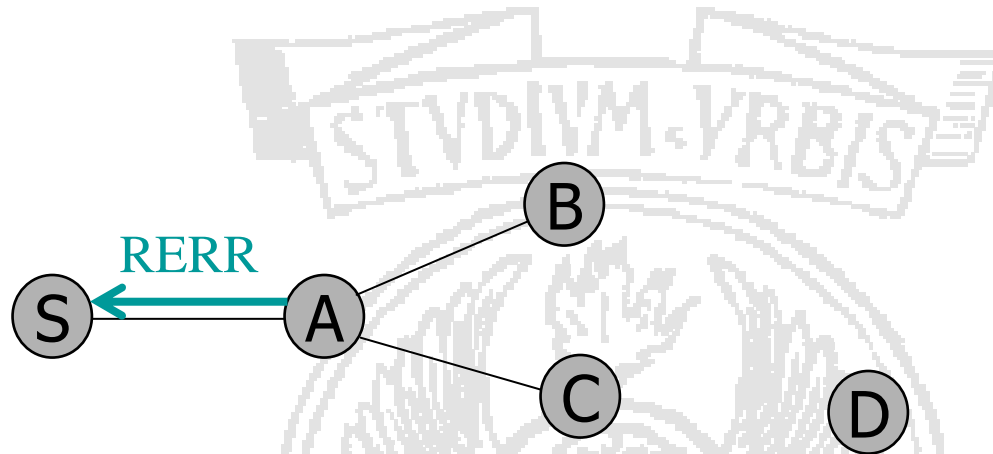
AODV: Route Maintenance



4. Node *A* receives RERR
 - Checks whether *C* is its next hop on route to *D*
 - Deletes route to *D*
 - Forwards RERR to *S*



AODV: Route Maintenance



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 (in that case it sends a RREQ with a SN which is equal to the last known destination Sequence Number +1)



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
 - Advantage: Less overhead when successful
 - Disadvantage: Longer delay if route not found immediately



AODV: Optimizations (cont.)

- Local Repair
 - 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
 - Disadvantage: longer delay and greater packet loss when unsuccessful



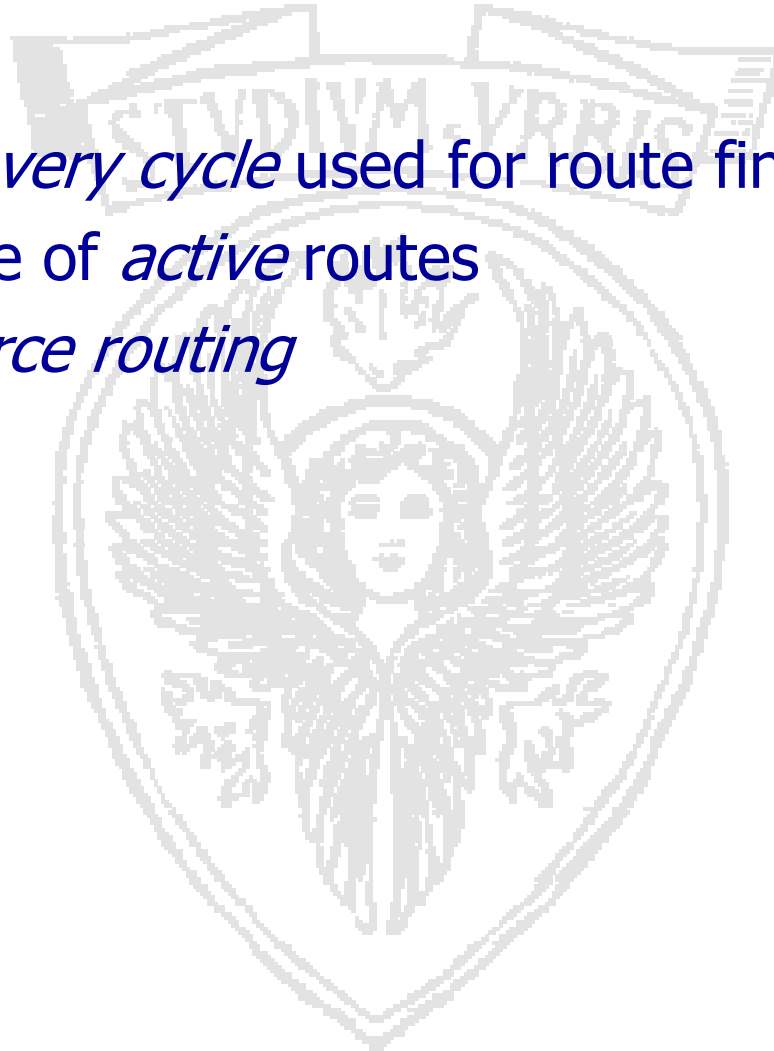
AODV: Summary

- Reactive/on-demand
- Sequence numbers used for route freshness and loop prevention
- Route discovery cycle
- Maintain only active routes
- Optimizations can be used to reduce overhead and increase scalability



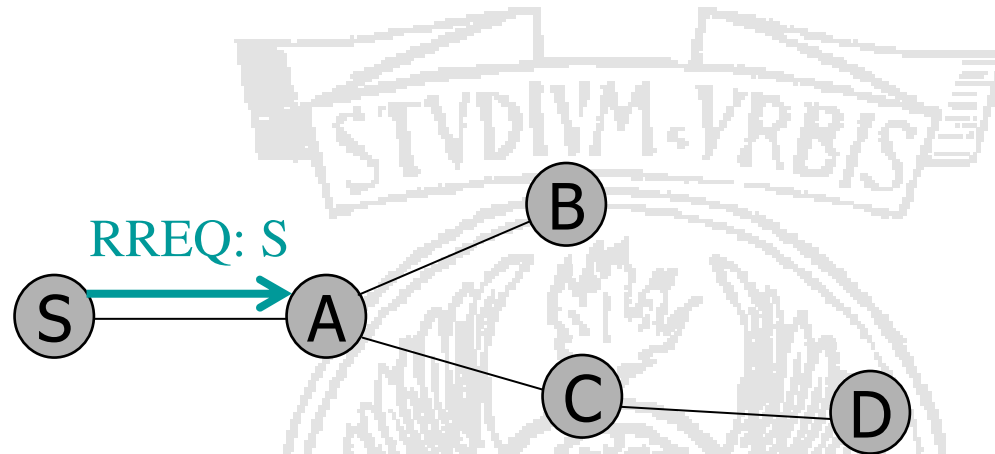
Dynamic Source Routing (DSR)

- Reactive
- *Route discovery cycle* used for route finding
- Maintenance of *active* routes
- Utilizes *source routing*





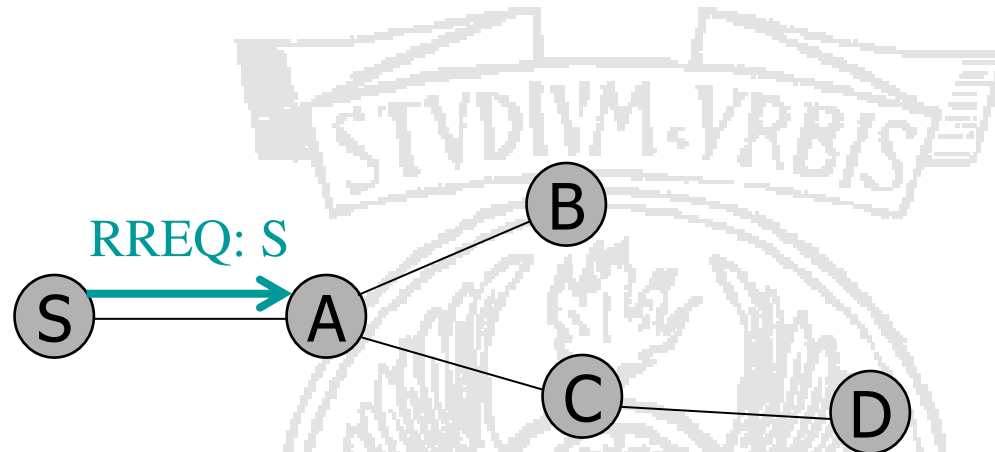
DSR: Route Discovery



1. Node *S* needs a route to *D*
2. Broadcasts RREQ packet
 1. RREQ identifies the target of the route discovery, contains a route record in which the traversed route is accumulated, contains a pair <initiator, request id> uniquely identifying the request



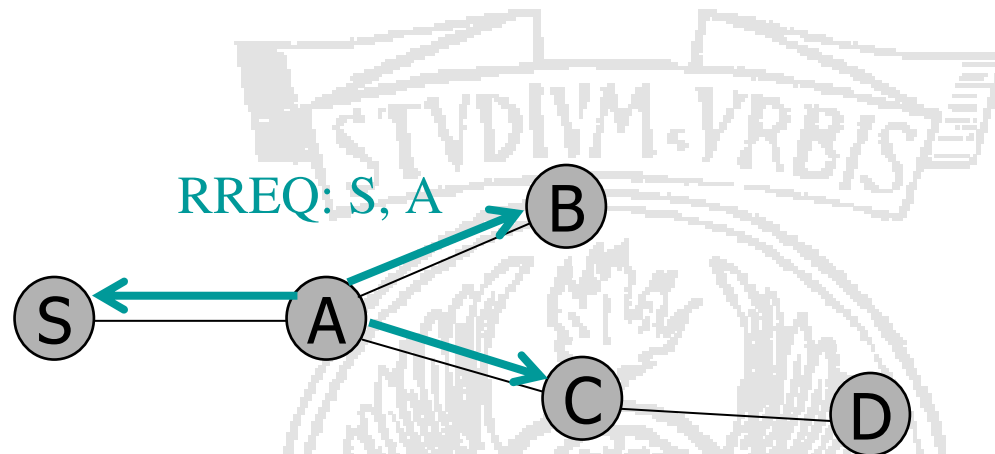
DSR: Route Discovery



1. Node *S* needs a route to *D*
2. Broadcasts RREQ packet
3. Node *A* receives packet, has no route to *D* *AND* is *NOT* *D*
 - Rebroadcasts packet after adding its address to source route



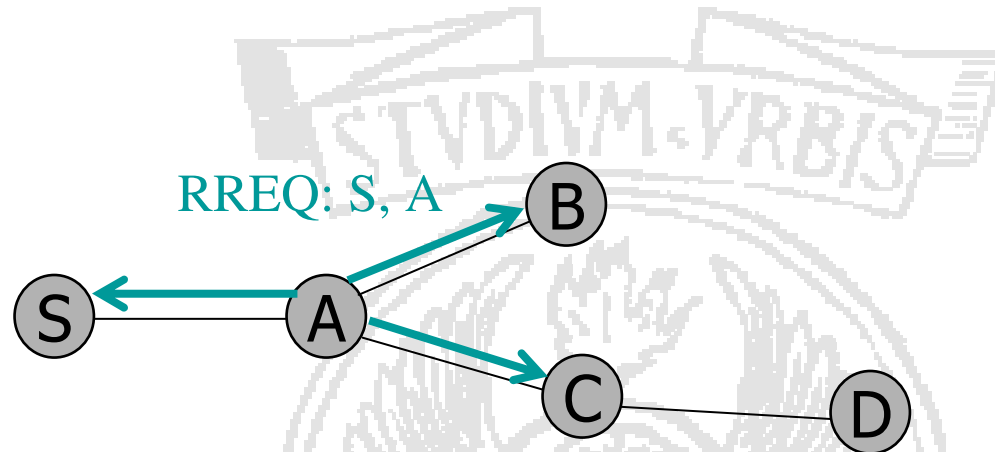
DSR: Route Discovery



1. Node *S* needs a route to *D*
2. Broadcasts RREQ packet
3. Node *A* receives packet, has no route to *D*
 - Rebroadcasts packet after adding its address to source route



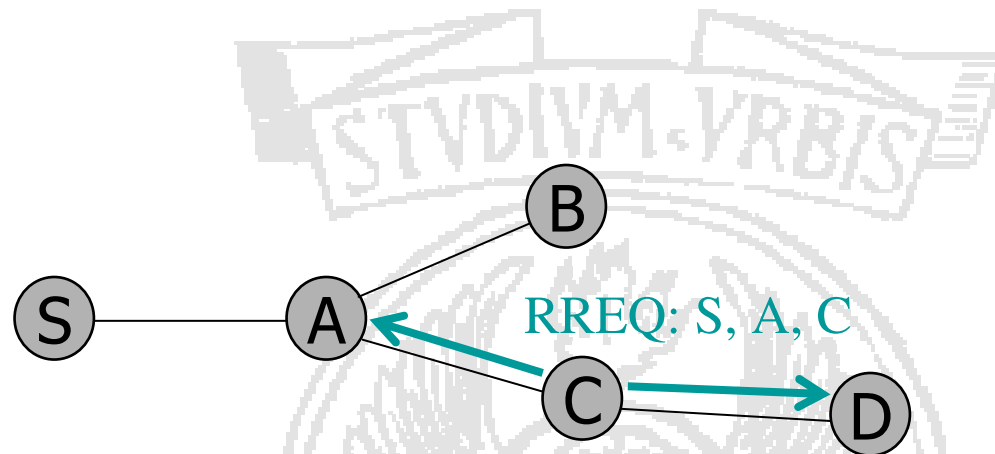
DSR: Route Discovery



4. Node *C* receives RREQ, has no route to *D*
 - Rebroadcasts packet after adding its address to source route




DSR: Route Discovery



4. Node *C* receives RREQ, has no route to *D*
 - Rebroadcasts packet after adding its address to source route



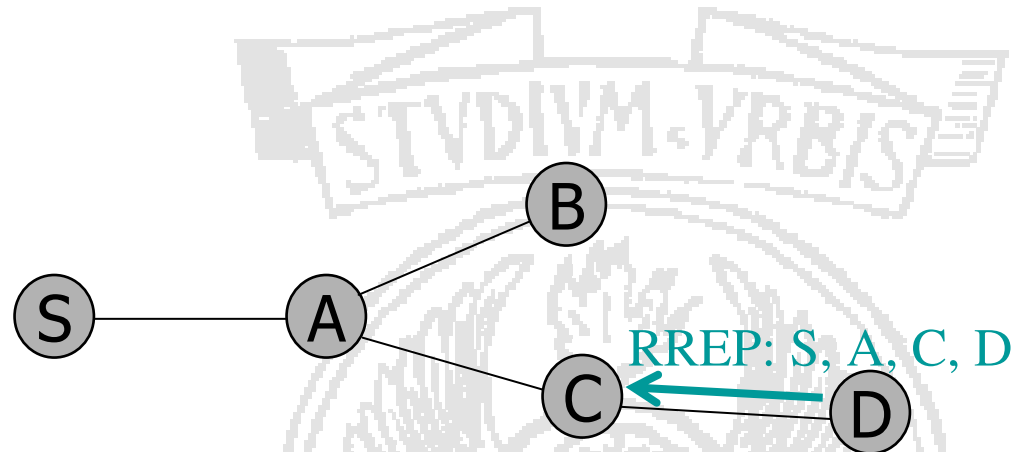
DSR: Route Discovery



4. Node *C* receives RREQ, has no route to *D*
 - Rebroadcasts packet after adding its address to source route
5. Node *D* receives RREQ, unicasts RREP to *C*
 - Puts source route accumulated in RREQ into RREP



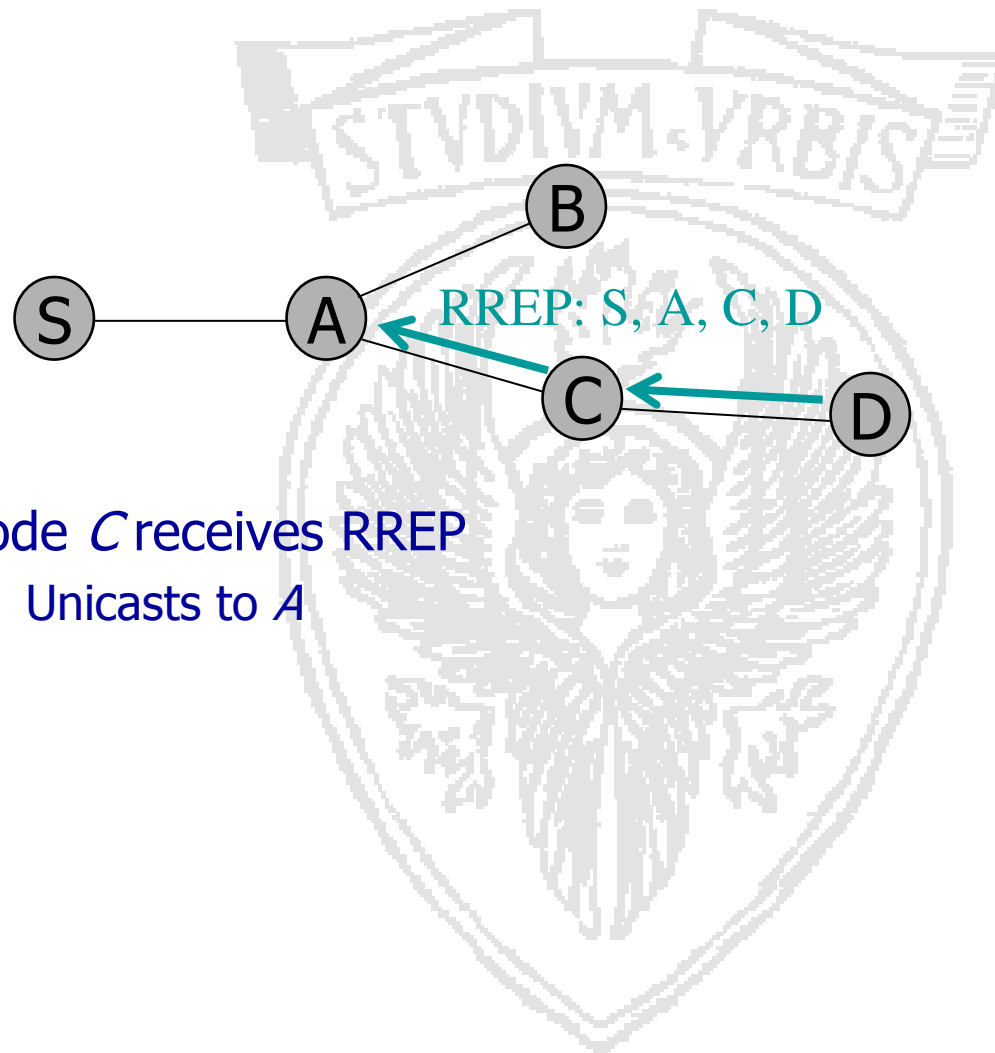
DSR: Route Discovery



4. Node *C* receives RREQ, has no route to *D*
 - Rebroadcasts packet after adding its address to source route
5. Node *D* receives RREQ, unicasts RREP to *C*
 - Puts source route accumulated in RREQ into RREP



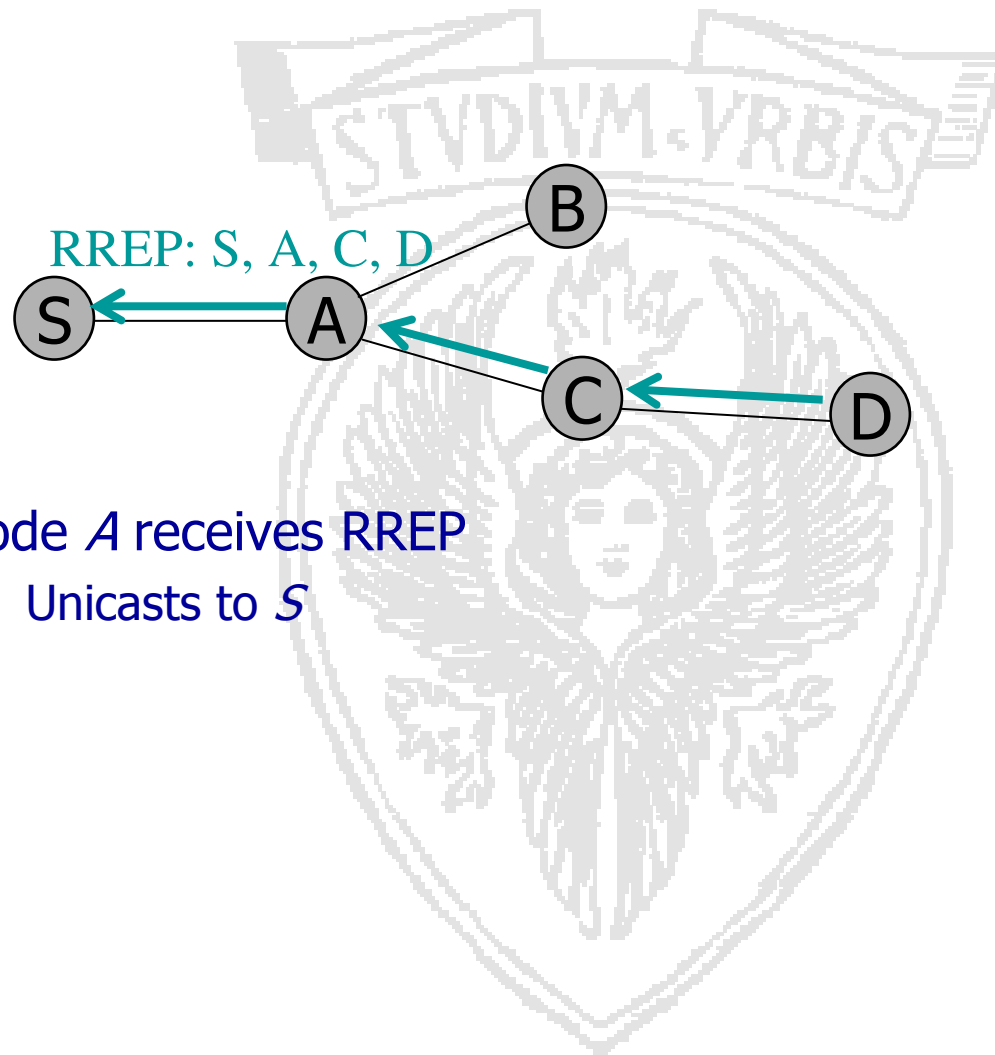
DSR: Route Discovery



6. Node *C* receives RREP
 - Unicasts to *A*



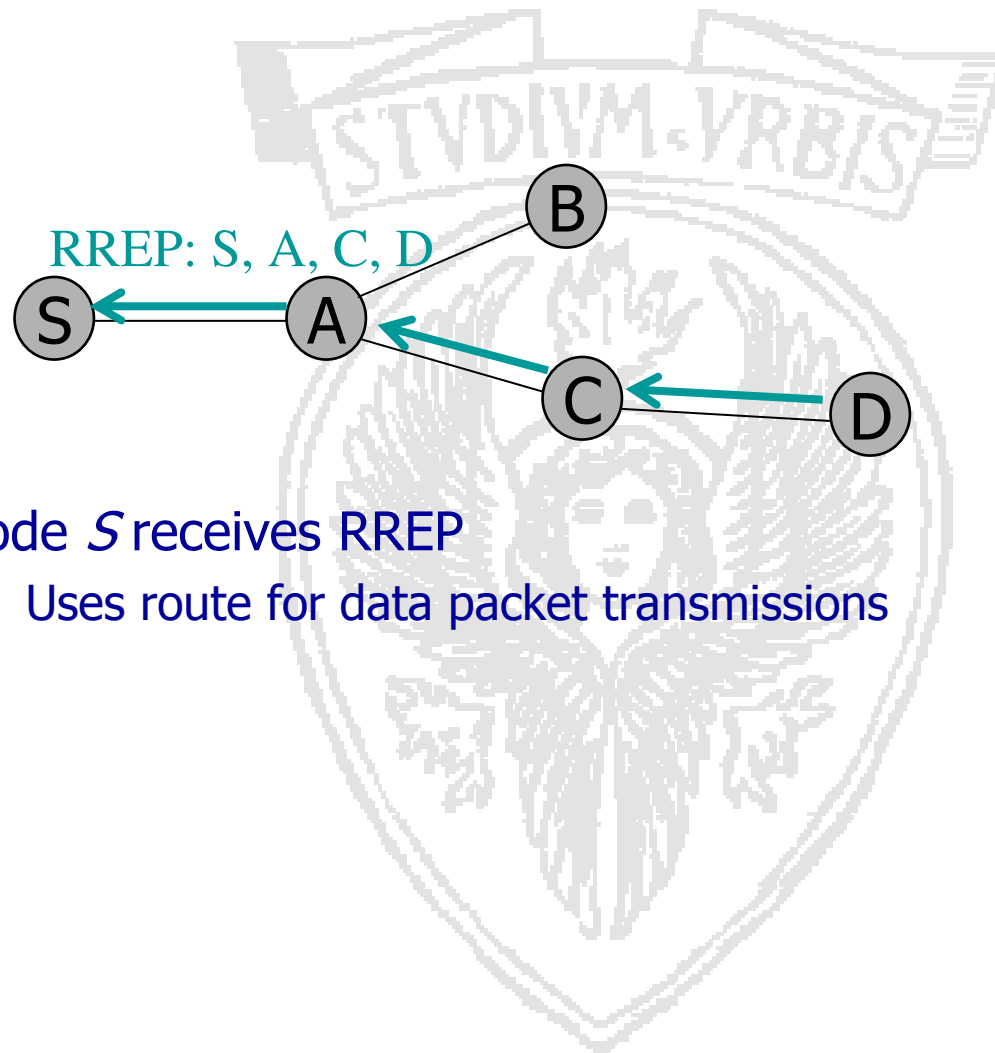
DSR: Route Discovery



6. Node A receives RREP
 - Unicasts to S



DSR: Route Discovery



8. Node *S* receives RREP
 - Uses route for data packet transmissions



General node operation upon receiving RREQ

- If the pair <initiator address, request ID> has recently been seen, DISCARD
- If the node ID is already listed in the source route DISCARD → avoids loops
- If I'm the destination, send a RREP
- Otherwise, append my ID in the source route and rebroadcast (orange cases already seen in the previous slides)



- The two endpoints of a failed link are transmitted to the source in a route error packet
- Upon a receiving a RERR packet a node invalidates all the routes going through that link
- If the route is invalidated and it is needed, a new route must be discovered




- Extensive use of caching (caching source routes means that I already know all the route to intermediate destinations, discovery a better route to an intermediate destination also brings me to improve the route to the final destination). Transmitting packets or sending back replies make me learn routes.
- A node that knows a route to a given destination (has a source route in cache) can immediately answer a RREQ
 - Broadcast storm? Each nodes waits for a time which is $C*(h-1+r)$, r random in $(0,1)$, h length of the route I'm advertising. Only if I haven't received other routes –listen to other routes tx in the meanwhile-I transmit mine.



- Operation in promiscuous mode (I keep discovering new routes by transmission of routes by my neighbours)
- RREQ overhead minimization: first set a TTL=1, if I do not get answer I set it to infinity
- Path shortening: if Y receives a packet by node X but in the source route we have X, B,...,C,Y, Y signals the path can be shortened (unsolicited RREP)
- What if the network is disconnected? Exponential back-off to decrease the quantity of RREQ sent



- DSR uses source routing; **AODV uses next hop entry**
- DSR uses route cache; **AODV uses route table**
- DSR route cache entries do not have lifetimes;
AODV route table entries do have lifetimes





Proactive Solutions: Drawbacks

- Updates overhead, especially in presence of high mobility
- Overhead for enforcing loop freedom
- Large routing tables
- Low *scalability*
- Is it really necessary to maintain a consistent view of the network topology?



Reactive Protocols: Drawbacks

- The discovery phase introduces long delays
- Route discovery and maintenance is very sensitive to node mobility
- Route caching is memory greedy
- The size of the header of a data packet can become cumbersome in approaches such as DSR (no scalability)
- Operating in promiscuous mode is energy-consuming.
- Relying on flooding based route discovery is resource consuming.
- Is the dependency on the network topology avoidable?



Geographically-Enabled Routing

- Outline
 - Problems with proactive approaches
 - Problems with reactive approaches
 - A new way of naming\locating the destination node: geographic routing
 - Two seminal protocols
 - ✓ DREAM & LAR
 - Geo-enable routing costs: I need to know where I am, where the destination is.



Location-Enabled Ad Hoc Routing

- Nodes are equipped with positioning system devices (e.g., Global Positioning System receivers) that make them aware of their position
- This enables “directional” routing
- Possible solutions differ on how the location information of the destination nodes is achieved



Strengths

- No need to update big routing tables, no need to piggyback routes to the packet
 - Destination position must be known at the source.
- No need to know the nodes on the way to the destination: they can be moving while the packet travels



Drawbacks

- Needs extra hardware
- Depends on the extra hardware limitation (and resource requirements)
- Scalability is an issue (indeed the problem here translates to how to maintain correct estimates of the nodes positions)



Location-Aided Routing (LAR)

- Exploits location information to limit scope of RREQ flood
- **Expected Zone**: region that is expected to hold the current location of the destination
 - Expected region determined based on potentially old location information, and knowledge of the destination's speed
- RREQs limited to a **Request Zone** that contains the Expected Zone and location of the sender node




Figure 2. Examples of expected zone.




Figure 3. Request zone. An edge between two nodes means that they are neighbors.

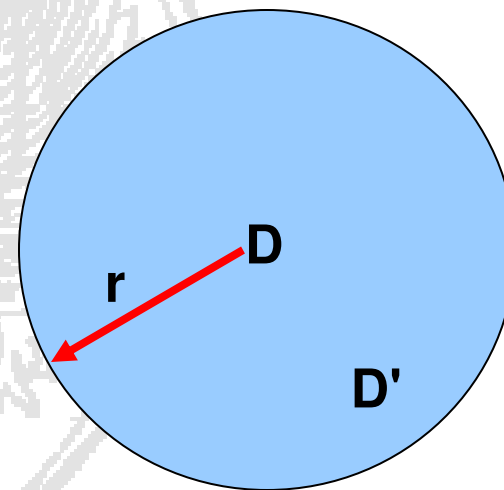


LAR: Expected Zone

**D = last known location of node
D, at time t_0**

**D' = location of node D at current
time t_1 , unknown to node S**

$r = (t_1 - t_0) * \text{estimate of D's speed}$




Expected Zone



LAR-1

- The **request zone** is the smallest rectangle that includes the current location of the source and the expected zone
- Only nodes **within the request zone** forward route requests
 - Node A does not forward RREQ, but node B does
- Request zone explicitly specified in the RREQ
- Each node must know its physical location to determine whether it is within the request zone





LAR, Possible Failures

- If route discovery using the smaller request zone fails to find a route, the sender initiates another route discovery (after a timeout) using a larger request zone
 - The larger request zone may be the entire network
- Rest of route discovery protocol similar to DSR
 - The directional flooding approach proposed in the paper is used to reach the destination with a limited overhead wrt traditional flooding
 - Destination then answers with a route reply packets
 - This allows the source to compute based on traditional reactive protocols approaches the route towards the destination



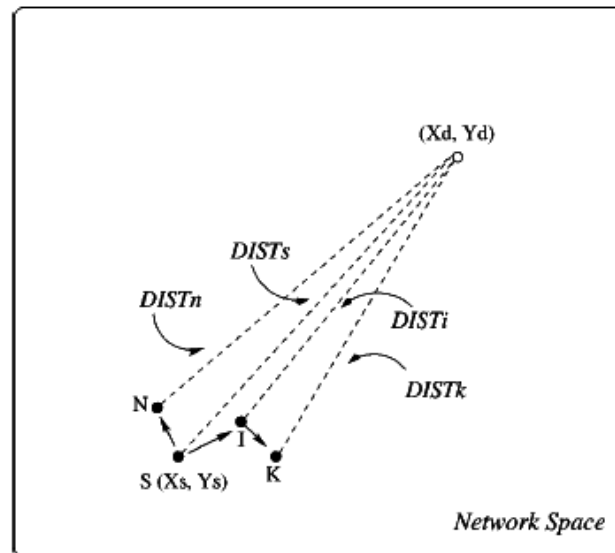
LAR, the Routing

- The basic proposal assumes that, *initially*, location information for node X becomes known to Y only during a route discovery
- This location information is used for a future route discovery
 - according to the paradigm explained before
 - Updates on the node position are piggybacked in the route reply message
 - This allows to reduce overhead associated to route discovery process provided that the time between two route discovery is not too much
- Destination may also proactively distribute its location information
 - But in this case the control traffic for geographic information updates could be high
 - ✓ How to solve this issue? Later (DREAM)



LAR - 2

- Each protocol relies RREQ if it is closer to the destination than the source (greedy forwarding)
- Assume that node S knows the location (X_d, Y_d) of node D at some time t_0 – the time at which route discovery is initiated by node S is t_1 , where $t_1 \geq t_0$. Node S calculates its distance from location (X_d, Y_d) , denoted as $DIST_s$, and includes this distance with the route request message.
- The coordinates (X_d, Y_d) are also included with the route request.
- For some parameters α and β , if $\alpha(DIST_s) + \beta \geq DIST_i$, then node I forwards the request to its neighbors. When node I forwards the route request, it now includes $DIST_i$ and (X_d, Y_d) in the route request (i.e., it replaces the $DIST_s$ value received in the route request by $DIST_i$, before forwarding the route request).
- Else $\alpha(DIST_s) + \beta < DIST_i$. In this case, node I discards the route request.





DREAM

- Distance routing effect algorithm for mobility [Basagni+, 1998]
- A proactive, effective way to spread location information
- Directional routing




DREAM: Directional Routing

- Source S determines the location of destination D at time t_0 based on its location table
- Based on the current time t_1 and t_0 S determines the area in which D can be found (hence, D's direction)
- S transmits the data packet to all its neighbors in D's direction
- Each neighbor does the same till D is reached



DREAM: Routing a Data Packet






- Need to periodically update the location of a moving node.
 - Efficient broadcast of location information
 - Determining how far each location packet should travel
 - Determining how often a location packet should be sent



- Distance effect
- Rate of updates is bound to the mobility of the node



The Distance Effect



- “Closer nodes look like they are moving faster”
- Need to receive more location updates from closer node
- Each location update packet is associated with an age that determines how far that packet must travel



DREAM: Rate of updates

- Triggered by the mobility of the nodes
- The faster the node the more updates it sends
- A plus: slow moving nodes impose little overhead



DREAM, Strengths

- First of its kind: after that, the deluge!
- Robustness: multiple routes to the destination
 - directional flooding

DREAM, Weaknesses

- It is flooding, although only directional
- It is not that scalable, geographic info updates have to be periodically transmitted (even if mechanisms to limit such overhead are enforced)



- For solutions which scale
- Which are energy saving
 - Which are well integrated with awake/asleep schedules
- Which do not require to maintain routing tables
- Which are simple
- Solutions such as AODV and DSR have been proven to work well iff they exploit intensively caching and promiscuous mode operation (energy inefficient ← work by L. Feeney et al, 2001) and have been shown not to scale to the volumes of nodes expected in sensor networks (work by E. Belding Royer and S.J. Lee)
- What can we use?
 - communication sensors – sink
 - Info such as localization and some level of synchronization often needed by the application (if I sense an event I have to say WHERE and WHEN it occurred, otherwise the information is not very interesting)



An example: GeRaF


- Integrates
 - geographic routing
 - awake/asleep schedule
 - MAC
- How do nodes alternate between awake and asleep states? According to a duty cycle (time ON/time ON+OFF)

ON

ON

OFF
94

OFF



Geographic routing:
each node needs to
know its location, the
destination (sink)
location, and the
location of whom is
transmitting
(communicated in the
packet)
Greedy approach:
tries to select relays
so to advance as
much as possible
toward the destination
at each hop





GeRaF: operations

- Main problem to be solved: how to make a contention-based scheme/routing work in the presence of sleep modes
 - Flat solution
 - Integrated MAC/routing/awake-asleep but awake-asleep schedule and routing decoupled → each node does not know its neighbor and their schedules → low overhead
- Tightly integrated with the routing layer (no clear separation really)
 - Without requiring routing tables/routing table updates
 - Based on the knowledge of the nodes location and on the knowledge of the sink location



- RTS invites all awake neighbors to become a relay
- Nodes in best position should win
 - Nodes within tx range are divided in areas depending on how close they are to the final destination (the closest the better as relay)



•Need of location awareness



- Node i sends RTS with the identity of the area it is polling now (starting from the closer to the sink, among the slices in which its tx range has been divided)
- Each node, upon receiving the RTS, decides whether it belongs to the polled area or not (based on location info)
- Only nodes in the polled area answer with a CTS
 - No node answers \rightarrow node i polls next area (no node available for forwarding in the area-there are no nodes or they are sleeping)
 - One answer, CTS correctly received, send DATA
 - Multiple answer COLLISION, sender sends a collision packet, MAC needed to solve collision (next slide)



- 1) A node receiving a collision packet tosses a coin and with probability p transmits a CTS iff it was participating to the previous phase (it had previously sent a CTS resulting in collision)
 - if only one node answers node i sends data
 - If no node answers node i asks these nodes to toss a coin again..
 - if more nodes answer COLLISION. Collision packet is sent. GO TO 1) (only the nodes which have lead to collision survive to the next phase)

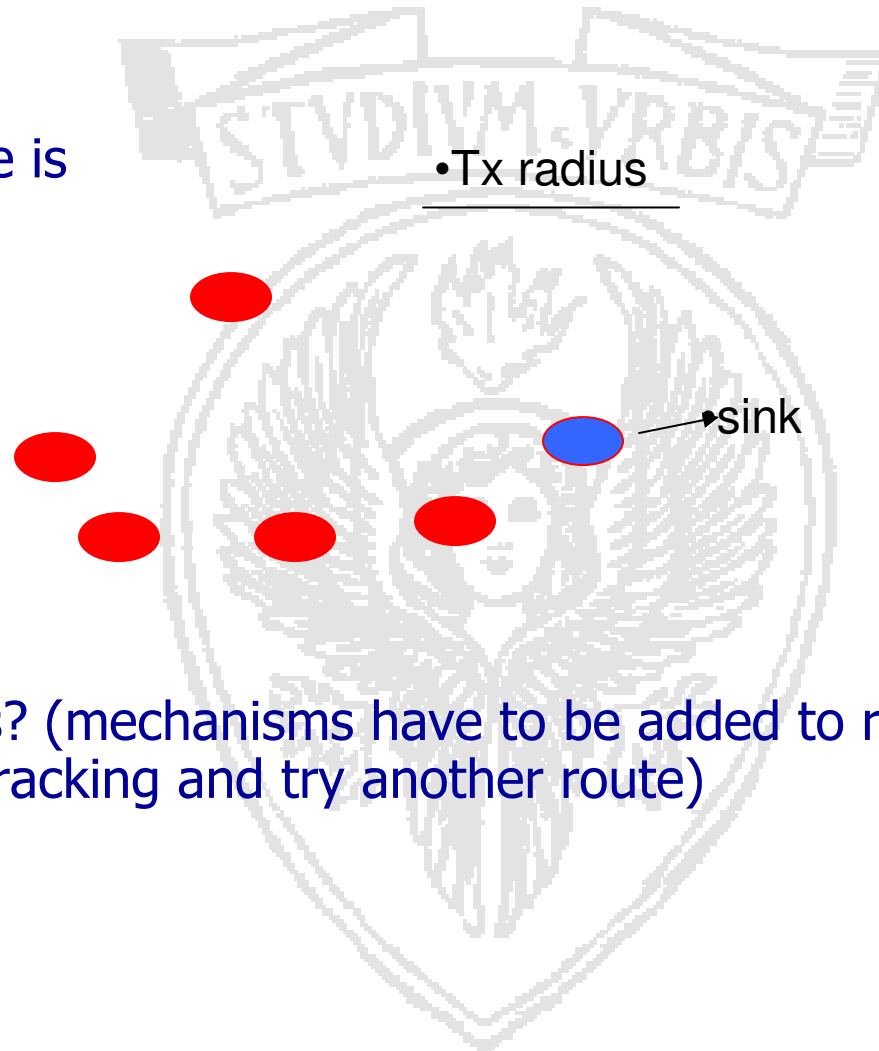


- All areas are polled unsuccessfully?
 - Try again after some time (exponential backoff)
- Can I always reach the destination in this way?





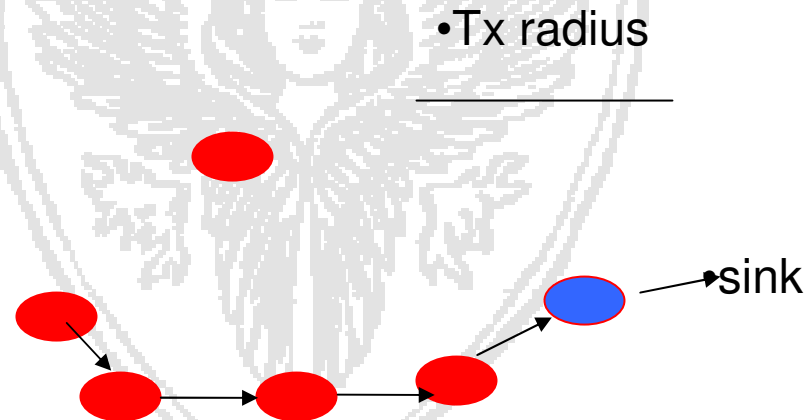
- No. Here is
An example



- Solutions? (mechanisms have to be added to recognize the problem, do backtracking and try another route)



- A problem only at low density
- We can set a maximum number of attempts to find a relay. When a node fails to find a relay it starts decreasing its duty cycle/or the probability to propose itself as relay ...over time nodes along paths to dead-ends are less and less selected as next hop relays and other paths able to bring to the destination are instead found
- Still..we may have problems...







Casari, Marcucci, Nati, Petrioli, Zorzi IEEE MILCOM 2005

- square area 320m x 320m
- Transmission range=40m
- 100-1000 randomly deployed nodes (avg degree 5-50)
- Duty cycle =0.01,0.1,0.5
- Comparable costs for tx/rx/idle
- Poisson packet arrival
- Channel data rate 38Kbps





Casari, Marcucci, Nati, Petrioli, Zorzi IEEE MILCOM 2005





Casari, Marcucci, Nati, Petrioli, Zorzi IEEE MILCOM 2005



Con i meccanismi per evitare dead end si sale a % di successfully Delivered packets nel caso di 200 nodi pari a 93-97% (evitando nel tempo cammini che portino a dead ends). Non si risolve il problema nel caso $n=100$. Soluzione completa in uno schema che abbiamo proposto: ALBA.