

Modeling AODV

(Ad hoc On-Demand Distance Vector Routing Protocol)

Università di Pisa, Dipartimento di Informatica boerger@di.unipi.it
Universität Ulm, Abteilung Informatik alexander.raschke@uni-ulm.de
See Ch. 6.1 of ModelingCompanion

Why is a rigorous AODV model needed?

- C.E.Perkins/E.M.Belding-Royer/S.R.Das: *AODV* RFC 3561 de facto standard docu of 2003 is **partly ambiguous, incomplete, contradictory**
 - various implementations show different behavior on protocol relevant features (e.g. loops, route discovery, packet delivery, optimal routes)
 - a huge number of different interpretations of RFC 3561 are possible
- widely believed, performance relevant loop freedom claim (by Perkins and Royer in 1999) established only in 2013 for a ‘correct’ interpretation of AODV
 - in: A. Fehnker, R. van Glabbeek, P. Höfner, A. McIver, M. Portmann, W.L. Tan: A process algebra for wireless mesh networks used for *modelling, verifying and analysing AODV* (TR 5513, NICTA, 2013)
 - after some wrong/partial proofs in the literature
- TR 5513 identifies also some performance relevant shortcomings of AODV and five key implementations

Goal of this lecture

- **explain** the functional core behavior of AODV for users and programmers, **'from scratch' and reliably**, by stepwise developing an Abstract State Machine (ASM) model
 - reflecting a correct and complete understanding of the (core) requirements in the de facto standard document RFC 3561
 - but informed by the professional analysis in the NICTA TR 5513
- as a prerequisite for a rigorous high-level analysis, long **before coding**
- performed in the NICTA TR 5513, in process algebraic terms, for different interpretations and implementations of the RFC 3561 wrt
 - loop freedom
 - route discovery
 - packet deliveryand related correctness and performance relevant issues

Mobile Ad hoc Network (MANET) routing protocols

In MANETs *every network agent*

- can move independently to change its position
- can (try to) send messages to every 'directly connected' network node it knows ('neighbor')
- can (try to) broadcast messages to all its 'neighbors'
- for (wireless) communication with any other network agent must ask a routing protocol to indicate a communication path to that destination

The *routing protocol*

- receives and elaborates *route request* messages, by forwarding them and generating reply messages once a route has been found
- receives and elaborates *route reply* messages by forwarding them back to the original requestor
- creates and propagates *route error* messages if some broken direct link is detected

Background and agent/protocol interaction structure

Background structure:

- network: graph $(Agent, Link)$ with dynamic sets of nodes and edges
- determines for each $a \in Agent$ a dynamic set $neighb(a)$

Agent/protocol interaction: when a *WantsToCommunicateWith* d

- in case a *KnowsActiveRouteTo* (d) , it can right away `STARTCOMMUNICATIONWITH` (d) , without entering *WaitingForRouteTo* (d) (which is initialized by *false*)
- otherwise it must `GENERATEROUTEREQ` (d) and becomes *WaitingForRouteTo* (d) until via the protocol run it eventually *KnowsActiveRouteTo* (d)
- to `GENERATEROUTEREQ` (d) only once per required communication, it is called only when a is not already *WaitingForRouteTo* (d)
 - easily refinable to permit repeated route requests to a same d

Agent/protocol interaction component

PREPARECOMM =

if *WantsToCommunicateWith(destination)* **then**

if *KnowsActiveRouteTo(destination)*

then

STARTCOMMUNICATIONWITH(*destination*)

WantsToCommunicateWith(destination) := *false*

WaitingForRouteTo(destination) := *false*

else

if not *WaitingForRouteTo(destination)* **then**

GENERATEROUTEREQ(*destination*)

WaitingForRouteTo(destination) := *true*

NB. *WantsToCommunicateWith* is assumed to be set to *true* only by the application program and to *false* only by PREPARECOMM.

Router components and AODV program structure

The main AODV program each node is equipped with:

AODV**SPEC** = **one of** `PREPARECOMM` `ROUTER` -- main program

`PREPARECOMM`

`ROUTER`

The `ROUTER` consists of components to process request, reply and error msgs:

`ROUTER` = **one of**

`PROCESSROUTEREQ`

`PROCESSROUTEREP`

`PROCESSROUTEERR`

`PROCESSROUTEERR` = **one of**

`GENERATEROUTEERR`

`PROPAGATEROUTEERR`

AODV data structure: route table $RT(a)$ information

Each route table *entry* keeps agent a 's knowledge that

- the destination $dest(entry)$ of the *entry* may be reachable
- in the direction as indicated by a neighbor node $nextHop(entry)$
- on a path of length $hopCount(entry)$ (distance to the destination)
- which can be considered as *Active* (without *LinkBreak*)

To avoid loops in communication paths, each agent keeps a local request/-level counter:

- $localReqCount(a)$ and $curSeqNum(a)$, both initialized by 0
- $curSeqNum(a)$ possibly incremented when a receives a new request to reach a , after a *LinkBreak* which made an *Active* path to a *InActive*

Each route table *entry* for $d = dest(entry)$ also records

- as $destSeqNum(entry)$ the last known value of $curSeqNum(d)$
- by $known(entry) \in \{true, false\}$ whether this number is valid

NB. For $precursor(entry)$ info for route error handling see below

AODV data structure: route table entry attributes

$entryFor(d, RT) =$

$$\begin{cases} entry & \text{if forsome } entry \in RT \text{ } dest(entry) = d \\ \text{undef else} \end{cases}$$

$KnowsActiveRouteTo(destination) \text{ iff}$

$Active(entryFor(destination, RT))$

$lastKnownDestSeqNum(d, RT) =$

$$\begin{cases} destSeqNum(entry) & \text{if forsome } entry \in RT \text{ } dest(entry) = d \\ unknown & \text{else} \end{cases}$$

$ValidDestSeqNum(entry) \text{ iff } known(entry) = true$

AODV data structure: *RouteRequest*

Each route request msg $rreq \in RouteRequest$, when generated, records information about:

- the request destination $dest(rreq) \in Agent$
- the last known value $destSeqNum(rreq) \in NAT \cup \{unknown\}$ the request originator knows about the $curSeqNum(dest(rreq))$
- $known(rreq)$ indicating whether $destSeqNum(rreq)$ is reliable
- the request originator $origin(rreq) \in Agent$
- the originator's $originSeqNum(rreq) \in NAT$
- the length $hopCount(rreq) \in NAT$ of the path the $rreq$ traveled from its $origin(rreq)$ to its current $rreq$ -sender
- the value $localId(rreq) \in NAT$ of $localReqCount + 1$ at the $origin(rreq)$ when the $rreq$ is generated

NB. Global identification of $rreq$ via the following equation:

$$globalId(rreq) = (localId(rreq), origin(rreq))$$

AODV data structure: *RouteReply*

Each route reply msg $rrep \in RouteReply$, when generated, records information about:

- the destination $d = dest(rrep) \in Agent$ of the detected route
- the value $destSeqNum(rrep) \in NAT$ of $curSeqNum(d)$
 - as known at an intermediate node (not d), if $rrep$ is generated there
 - or as updated when the destination node d generates $rrep$
- the request originator $origin(rrep) \in Agent$ to whom the reply is addressed
- the length $hopCount(rrep) \in NAT$ of the current route from the $rrep$ -sender to $dest(rrep)$

NB. We abstract from concerns about $rrep$ lifetime, network traffic and performance properties.

AODV data structure: *RouteError*

Each route error msg $rerr \in RouteError$, sent by an agent a ,

- indicates a set of destinations, together with their increased *destSeqNum* value, which became unreachable via a , i. e. cannot be reached at present using a as *nextHop* of a route entry
- inactivates every route table *entry* which uses the *rerr* sender a as *nextHop* to any relevant unreachable destination communicated by *rerr*
- is forwarded along the *precursor* chain

GENERATEROUTEREQ(*destination*)

let $r = \text{new}$ (*RouteRequest*) **in**

$\text{dest}(r) := \text{destination}$

$\text{destSeqNum}(r) :=$

$\text{lastKnownDestSeqNum}(\text{destination}, RT)$

if $\text{entryFor}(\text{destination}, RT) \neq \text{undef}$

then $\text{known}(r) := \text{known}(\text{entryFor}(\text{destination}, RT))$

else $\text{known}(r) := \text{false}$

$\text{origin}(r) := \text{self}$ $\text{originSeqNum}(r) := \text{curSeqNum} + 1$

$\text{hopCount}(r) := 0$ $\text{localId}(r) := \text{localReqCount} + 1$

BROADCAST(r) -- i.e. **forall** $n \in \text{neighb}$ **do** SEND(r , **to** n)

INCREMENT(curSeqNum) INCREMENT(localReqCount)

BUFFER(r) -- i.e. INSERT($\text{globalId}(r)$, *ReceivedReq*)

NB. BUFFER(r) helps to recognize whether r has been ‘seen’ already

PROCESSROUTEREQ(*rreq*)

```
if Received(rreq) and rreq ∈ RouteRequest then  
  if not AlreadyReceivedBefore(rreq) then      -- rreq processed once  
    BUFFER(rreq)  
    if HasNewReverseRouteInfo(rreq) then  
      BUILDREVERSEROUTE(rreq)  
  seq  
    if FoundValidPathFor(rreq)  
      then GENERATEROUTEREPLY(rreq)  
      else FORWARDREFRESHEDREQ(rreq)  
  CONSUME(rreq)
```

NB. GENERATEROUTEREPLY sends *rreply* to *nextHop* in—possibly by BUILDREVERSEROUTE updated—*entryFor*(*origin*(*rreq*), *RT*), which could be different from *sender*(*rreq*).

A req where $HasNewReverseRouteInfo(req)$ is false

$AlreadyReceivedBefore(req)$ **iff** $globalId(req) \in ReceivedReq$

// i.e. req has been BUFFERED when received for the first time

Such route req msgs are simply discarded. Nothing else happens.

Otherwise, if the req brings no new reverse route information to the RT of the receiving agent, the existing reverse route entry is kept unchanged and the protocol proceeds to either $GENERATEROUTEREPLY(rreq)$ or $FORWARDREFRESHEDREQ(rreq)$.

$HasNewReverseRouteInfo(req)$ **iff** $req \in RouteRequest$ **and**

$ThereIsNoRouteInfoFor(origin(req), RT)$ **or**

$(ThereIsRouteInfoFor(origin(req), RT)$ **and**

$HasNewOriginInfo(req, RT))$

When $HasNewOriginInfo(req)$ to update the reverse route

$HasNewOriginInfo(req, RT)$ **iff**

let $entry = entryFor(origin(req), RT)$

$originSeqNum(req) > destSeqNum(entry)$

or $originSeqNum(req) = destSeqNum(entry)$ **and**

$(hopCount(req) < hopCount(entry)$ **or**

not $Active(entry)$)

NB. For comparison with $destSeqNumbers$, every natural number n is stipulated to be better than $unknown$, formally $unknown < n$.

BUILDREVERSEROUTE(*rreq*) component

if *ThereIsRouteInfoFor*(*origin*(*rreq*), *RT*) **then**

 UPDATEREVERSEROUTE(*entryFor*(*origin*(*rreq*), *RT*), *rreq*)

else EXTENDREVERSEROUTE(*RT*, *rreq*)

where

UPDATEREVERSEROUTE(*e*, *req*) =

destSeqNum(*e*) := *originSeqNum*(*req*) -- freshest *destSeqNum*

known(*e*) := *true* *Active*(*e*) := *true*

nextHop(*e*) := *sender*(*req*)

hopCount(*e*) := *hopCount*(*req*) + 1 -- maybe shorter path

EXTENDREVERSEROUTE(*RT*, *req*) =


let *e* = **new** (*RT*)

dest(*e*) := *origin*(*req*) *precursor*(*e*) := \emptyset


 UPDATEREVERSEROUTE(*e*, *req*)

Redirecting Reverse Route example

- Reverse routes for segments of $rreq$ -path from a_1 to a_{n+1} are created.



- a_i receives a new req from a_1 to another destination and redirects the $rreq$ reverse route at a_i .¹



¹ Figures © 2018 Springer-Verlag Berlin Heidelberg, reprinted with permission

The meaning of $FoundValidPathFor(req)$

- either req is received by the *destination* node
- or an 'intermediate' node which $KnowsFreshEnoughRouteFor$ the $dest(req)$ receives the req :

$FoundValidPathFor(req)$ **iff**

$dest(req) = \mathbf{self}$ **or** $KnowsFreshEnoughRouteFor(req, RT)$

A 'fresh enough' route entry besides being *Active* must have a $ValidDestSeqNumber$ that is not smaller than the $destSeqNumber$ of the received route request :

$KnowsFreshEnoughRouteFor(req, RT)$ **iff** **forsome** $entry \in RT$

$dest(entry) = dest(req)$ **and** $ValidDestSeqNum(entry)$

and $destSeqNum(entry) \geq destSeqNum(req)$

and $Active(entry)$

req forwarding increases hopCount and possibly destSeqNum

FORWARDREFRESHEDREQ(r) =

let $r' = \text{new}$ (*RouteRequest*)

COPY($dest, origin, originSeqNum, localId, known,$
from r **to** r')

$hopCount(r') := hopCount(r) + 1$

$destSeqNum(r') := \max\{destSeqNum(r),$
 $lastKnownDestSeqNum(dest(r), RT)\}$

BROADCAST(r')

where

COPY($f_1, \dots, f_n, \text{from } arg \text{ to } arg'$) =

forall $1 \leq i \leq n$ **do** $f_i(arg') := f_i(arg)$

GENERATEROUTEREPLY(*rreq*) at *dest* or intermediate node

```
let r = new (RouteReply)
let revEntry = entryFor(origin(rreq), RT)
  dest(r) := dest(rreq)
  origin(r) := origin(rreq)
if dest(rreq) = self then                                -- reply at destination node
  hopCount(r) := 0
  destSeqNum(r) := max{curSeqNum, destSeqNum(rreq)}
  curSeqNum := max{curSeqNum, destSeqNum(rreq)}
else let fwdEntry = entryFor(dest(rreq), RT)           -- at intermediate
  hopCount(r) := hopCount(fwdEntry)                       -- node
  destSeqNum(r) := destSeqNum(fwdEntry)
  PRECURSORINSERTION(nextHop(revEntry), fwdEntry)
SEND(r, to nextHop(revEntry))                            -- maybe to sender(rreq)
```

The role of precursor nodes

Consider the case that a node a has an $entryFor(d, RT)$ and by a $rerror$ msg, received by a , node d is reported as unreachable.

Then each $neighbor(a)$ which has a route entry to d that uses a as $nextHop$ is called a *precursor* of a .

The *precursor* set is recorded in $entryFor(d, RT)$ so that such a $rerror$ msg can be propagated to its elements.

$$\text{PRECURSORINSERTION}(node, entry) = \\ \text{INSERT}(node, precursor(entry))$$

NB. We leave 'gratuitous' replies as an exercise. Through gratuitous replies, a destination node obtains a reverse route to the request originator without having requested a route, namely in case an intermediate node answered the request. For this case one needs also a $\text{PRECURSORINSERTION}$ of $nextHop(fwdEntry)$ into $precursor(revEntry)$.

PROCESSROUTE_{REP}(*rrep*)

if *Received*(*rrep*) **and** *rrep* ∈ *RouteReply* **then**

if *HasNewForwardRouteInfo*(*rrep*) **then** -- else just discard *rrep*

BUILDFORWARDROUTE(*rrep*)

if *MustForward*(*rrep*) **then** **FORWARDREFRESHEDREP**(*rrep*)

CONSUME(*rrep*)

where -- note symmetry to PROCESSROUTE_{REQ}

HasNewForwardRouteInfo(*rep*) **iff** *rep* ∈ *RouteReply* **and**

ThereIsNoRouteInfoFor(*dest*(*rep*), *RT*) **or**

(*ThereIsRouteInfoFor*(*dest*(*rep*), *RT*) **and**

HasNewDestInfo(*rep*, *RT*))

MustForward(*rep*) **iff**

origin(*rep*) ≠ **self** **and** *Active*(*entryFor*(*origin*(*rep*), *RT*))

Meaning of *HasNewDestInfo*

HasNewDestInfo(*rep*, *RT*) **iff** -- *rep* has either

let *entry* = *entryFor*(*dest*(*rrep*), *RT*)

destSeqNum(*rep*) > *destSeqNum*(*entry*) -- better *destSeqNum*

or (*destSeqNum*(*rep*) = *destSeqNum*(*entry*)

and *hopCount*(*rep*) + 1 < *hopCount*(*entry*)) -- or shorter path

or (*destSeqNum*(*rep*) = *destSeqNum*(*entry*)

and *Active*(*entry*) = *false*) -- or **not** *Active*(*entry*)

NB. Remember *unknown* < *n* for each *n* = 0, 1, ...

BUILDFORWARDROUTE($rrep$) with fresh info

if $ThereIsRouteInfoFor(dest(rrep), RT)$
then UPDATEFORWARDROUTE($entryFor(dest(rrep), RT), rrep$)
else EXTENDFORWARDROUTE($RT, rrep$) -- create new entry

where

UPDATEFORWARDROUTE(e, rep) = -- copying fresh info

$destSeqNum(e) := destSeqNum(rep)$ $known(e) := true$

$nextHop(e) := sender(rep)$ $hopCount(e) := hopCount(rep) + 1$

$Active(e) := true$ SETPRECURSOR(rep, e)

EXTENDFORWARDROUTE(RT, rep) =

let $e = \text{new}(RT)$

$dest(e) := dest(rep)$ UPDATEFORWARDROUTE(e, rep)

SETPRECURSOR(rep, e) = **if** $MustForward(rep)$ **then**

INSERT($nextHop(entryFor(origin(rep), RT)), precursor(e)$)

FORWARDREFRESHEDREP(*rep*)

Upon forwarding *rep*, only the *hopCount* is updated:

FORWARDREFRESHEDREP(*rep*) =

let *rep'* = **new** (*RouteReply*)


COPY(*dest*, *destSeqNum*, *origin*, **from** *rep* **to** *rep'*)

hopCount(*rep'*) := *hopCount*(*rep*) + 1


SEND(*rep'*, **to** *nextHop*(*entryFor*(*origin*(*rep*), *RT*)))

Redirecting Forward Route example (1)

- Destination a_{n+1} answers $rreq$ before req by $rrep$ with $destSeqNum$ 0




- a_{n+1} broadcasts a new rrq for destination z , $curSeqNum := 1$




Redirecting Forward Route example (2)

- a_{n+1} answers req by rep with $destSeqNum(rep) := 1$
- a_i receives rep before $rrep$ and establishes $entryFor(a_{n+1}, RT(a_i))$ with $destSeqNum$ 1
- $rrep_{a_1, a_{n+1}}(0)$ is discarded at a_i



Redirecting Forward Route example (3)

- A new route request may be answered by a_i as intermediate node establishing the dotted communication path (\dots) from a_1 to a_{n+1}
- The communication path from a_{n+1} to a_1 still goes along the reverse route established by $rreq_{a_1, a_{n+1}}[1]^2$



² Figures (1)-(3) © 2018 Springer-Verlag Berlin Heidelberg, reprinted with permission

Generation of error messages

GENERATEROUTEERR =

```
let BrokenEntry =           -- compute Active entries with broken link
    { entry ∈ RT | LinkBreak(nextHop(entry)) and Active(entry) }
forall entry ∈ BrokenEntry           -- if there are any
    Active(entry) := false           -- Invalidate entry
    INCREMENT(destSeqNum(entry))    -- and destSeqNum
    let rerr = { (dest(e), destSeqNum(e) + 1) |
                e ∈ BrokenEntry and precursor(e) ≠ ∅ }
        forall a ∈ precursor(entry) SEND(rerr, to a)
```

NB. Stipulation $unknown + 1 = unknown$

PROPAGATEROUTEERR(*rerr*)

if *Received*(*rerr*) **and** *rerr* \in *RouteError* **then**

let *UnreachDest* = $\{(d, s) \in rerr \mid$ **forsome** *entry* \in *RT*

$d = dest(entry)$ **and** $nextHop(entry) = sender(rerr)$

and $Active(entry)$ **and** $destSeqNum(entry) < s\}$

forall $(d, s) \in UnreachDest$ **let** *entry* = *entryFor*(*d*, *RT*)

$Active(entry) := false$

$destSeqNum(entry) := s$

forall $a \in precursor(entry)$ SEND(*rerr'*, **to** *a*)

if *WaitingForRouteTo*(*d*) **then** REGENERATEROUTEREQ(*d*)

CONSUME(*rerr*)

where *err'* =

$\{(d', s') \in UnreachDest \mid precursor(entryFor(d', RT)) \neq \emptyset\}$

REGENERATEROUTEREQ(*d*) = ($WaitingForRouteTo(d) := false$)

Conclusion: What else to do with a rigorous model?

- **system debugging** by rigorous model analysis
 - (dis)prove system properties of interest (e.g. using PVS, KIV,...)
 - for AODV: loop freedom and correctness proved, route discovery and packet delivery disproved for process algebra model in TR 5513 (for loop freedom using Isabelle)
 - identify shortcomings (here e.g. non-optimal routes)
 - testing/modelchecking of executable model refinements (e.g. in CoreASM, Asmeta, ...)
- **system evaluation** by comparison of implementations with the model
 - in TR 5513 done for five key AODV implementations, three of them shown to possibly produce routing loops
- **model reuse** for experimenting with system extensions/variations, prior to coding
- **documentation** for maintenance needs

References

- C.E. Perkins, E.M. Belding-Royer, S. Das: *Ad hoc On-Demand Distance Vector (AODV) Routing*
 - RFC 3561, Networking Group. <http://www.ietf.org/rfc/rfc3561.txt>
- A. Fehnker, R. van Glabbeek, P. Höfner, A. McIver, M. Portmann, W.L. Tan: *A process algebra for wireless mesh networks used for modelling, verifying and analysing AODV*
 - TR 5513, NICTA, 2013. <http://www.nicta.com.au/pub?id=5513>
- E. Börger, A. Raschke: *Modeling Companion for Software Practitioners*
 - Springer 2018. (Ch.6.1)
<http://modelingbook.informatik.uni-ulm.de>

Copyright Notice

It is permitted to (re-) use these slides under the CC-BY-NC-SA licence

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

i.e. in particular under the condition that

- the original authors are mentioned
- modified slides are made available under the same licence
- the (re-) use is not commercial