

# A Rigorous Analysis of AODV and its Variants

Peter Höfner, Rob van Glabbeek, Wee Lum Tan,  
Marius Portman, Annabelle McIver, Ansgar Fehnker

Paphos, Cyprus  
October 23, 2012



**Australian Government**  
**Department of Broadband, Communications  
and the Digital Economy**  
**Australian Research Council**

**NICTA Members**



Department of State and  
Regional Development




The University of Sydney



**NICTA Partners**


- Mobile Ad Hoc Networks (MANETs)  
Wireless Mesh Networks (WMNs)
  - key features: mobility, dynamic topology, wireless multihop backhaul
  - quick and low cost deployment
- Applications
  - public safety
  - emergency response, disaster recovery
  - transportation
  - smart grid
  - London buses
  - ...
- Limitations in reliability and performance




- **Goal**
  - model, analyse, verify and increase the performance of wireless mesh routing protocols
  - develop suitable formal methods techniques
- **Benefits**
  - more reliable protocols
  - finding and fixing bugs
  - better performance
  - proving correctness
  - reduce “time-to-market”

- Ad Hoc On-Demand Distance Vector Protocol
  - Routing protocol for WMNs and MANETs
  - Ad hoc (network is not static)
  - On-Demand (routes are established when needed)
  - Distance (metric is hop count)
  - Developed 1997-2001 by Perkins, Beldig-Royer and Das (University of Cincinnati)
  - One of the four protocols currently standardised by the IETF MANET working group (IEEE 802.11s)

- Main Mechanism
  - if route is needed  
BROADCAST RREQ
  - if node has information about a destination  
UNICAST RREP
  - if unicast fails or link break is detected  
GROUPCAST RERR




- Main Mechanism
  - if route is needed  
BROADCAST RREQ
  - if node has information about a destination  
UNICAST RREP
  - if unicast fails or link break is detected  
GROUPCAST RERR





- Main Mechanism
  - if route is needed  
BROADCAST RREQ
  - if node has information about a destination  
UNICAST RREP
  - if unicast fails or link break is detected  
GROUPCAST RERR



- **Standards (IETF RFCs) are not precise**
  - written in English
  - ambiguous (sometimes incomplete)
  - no formal specification
- **Compliant implementations**
  - have different behaviours
  - are not compatible
  - have serious flaws
- **Traditional evaluation techniques: simulation and test-bed**
  - expensive
  - limited to (a small number of) specific scenarios
  - errors found after years of evaluation
  - barely offer any guarantee for properties such as route discovery



# Why Formal Specification?



# Why Formal Specification?





# Complete and Accurate Formalisation of AODV



```
+ [ (oip, rreqid) ∉ rreqs ]      /* the RREQ is new to this node */
  [[rt := update(rt,(oip,osn,kno,val,hops + 1,sip,∅))]      /* update the route to oip in rt */
  [[rreqs := rreqs ∪ {(oip,rreqid)}]]      /* update rreqs by adding (oip, rreqid) */
  (
    [ dip = ip ]      /* this node is the destination node */
    [[sn := max(sn,dsn)]]      /* update the sqn of ip */
    /* unicast a RREP towards oip of the RREQ */
    unicast(nhop(rt,oip),rrep(0,dip,sn,oip,ip)) . AODV(ip,sn,rt,rreqs,store)
    ► /* If the transmission is unsuccessful, a RERR message is generated */
    [[dests := {(rip,inc(sqn(rt,rip))) | rip ∈ vD(rt) ∧ nhop(rt,rip) = nhop(rt,oip)}]]
    [[rt := invalidate(rt,dests)]]
    [[store := setRRF(store,dests)]]
    [[pre := ∪{precs(rt,rip) | (rip,*) ∈ dests}]]
    [[dests := {(rip,rsn) | (rip,rsn) ∈ dests ∧ precs(rt,rip) ≠ ∅}]]
    groupcast(pre,rerr(dests,ip)) . AODV(ip,sn,rt,rreqs,store)
  + [ dip ≠ ip ]      /* this node is not the destination node */
    (
      [dip ∈ vD(rt) ∧ dsn ≤ sqn(rt,dip) ∧ sqnf(rt,dip) = kno]      /* valid route to dip that is fresh enough */
      /* update rt by adding precursors */
      [[rt := addpreRT(rt,dip,{sip})]]
      [[rt := addpreRT(rt,oip,{nhop(rt,dip)})]]
      /* unicast a RREP towards the oip of the RREQ */
      unicast(nhop(rt,oip),rrep(dhops(rt,dip),dip,sqn(rt,dip),oip,ip)) .
```

- **Developed Process Algebra**
  - inspired by  $\pi$ -calculus and LOTOS; based on  $\omega$ -calculus
  - main process expressions

$X(exp_1, \dots, exp_n)$	process calls
$P + Q$	nondeterministic choice
$[\varphi]P$	if-construct
$\llbracket \text{var} := \text{exp} \rrbracket P$	assignment followed by $P$
$\text{broadcast}(ms).P$	broadcast message followed by $P$
$\text{unicast}(dest, ms).P \blacktriangleright Q$	unicast $ms$ to $dest$ ; if successful proceed with $P$ ; otherwise with $Q$
$\text{receive}(msg).P$	receive message

- "Formal languages are useful tools for specifying parts of protocols. However, as of today, there exists no well-known language that is able to capture the full syntax and semantics of reasonably rich IETF protocols."  
[IETF]
- IETF Requirements (for formal methods)
  - relatively easy to extract code
  - complete specification
  - implementation independent
- Easy to use
  - only a few (well-known) programming constructs

- **Achievements**

- full concise specification of AODV (RFC 3561)
  - 6 processes (~140 lines; instead of 40 pages English prose)
  - without time
- verified/disproved properties
  - route discovery
  - packet delivery
  - loop freedom
    - first (correct) proof
    - disproved loop freedom for variants of AODV  
(implemented in at least 3 open source implementations, e.g. AODV-ns2)
- found several ambiguities, contradictions, shortcomings
- found solutions for some limitations



- **Proofs**
  - independent of topology
    - no need for specific scenarios (e.g., Random Way Point Model)
    - without having to simulate each and every network topology
  - modularity / reusability
    - line-by-line analysis
    - mainly based on invariants
    - loop-freedom proof based on sequence of smaller results
  - analysed 432 “reasonable” interpretations (112 are loop free and “correct”)
  - even some interpretations we never thought about
- simulation and test-bed experiment would have to be repeated for each interpretation
  - only a few topologies can be checked

- One of the Ambiguities (RFC3561)

sequence number field is set to false. The route is only updated if the new sequence number is either

- (i) higher than the destination sequence number in the route table, or
- (ii) the sequence numbers are equal, but the hop count (of the new information) plus one, is smaller than the existing hop count in the routing table, or
- (iii) the sequence number is unknown.

# Analysing Variants of AODV




```
+ [ (oip, rreqid) ∉ rreqs ]      /* the RREQ is new to this node */
  [[rt := update(rt,(oip,osn,kno,val,hops + 1,sip,0))]      /* update the route to oip in rt */
  [[rreqs := rreqs ∪ {(oip,rreqid)}]]      /* update rreqs by adding (oip, rreqid) */
  (
    [ dip = ip ]      /* this node is the destination node */
    [[sn := max(sn,dsn)]]      /* update the sqn of ip */
    /* unicast a RREP towards oip of the RREQ */
    unicast(nhop(rt,oip),rrep(0,dip,sn,oip,ip)) . AODV(ip,sn,rt,rreqs,store)
    ▶ /* If the transmission is unsuccessful, a RERR message is generated */
    [[dests := {(rip,inc(sqn(rt,rip))) | rip ∈ vD(rt) ∧ nhop(rt,rip) = nhop(rt,oip)}]]
    [[rt := invalidate(rt,dests)]]
    [[store := setRRF(store,dests)]]
    [[pre := ∪ {precs(rt,rip) | (rip,*) ∈ dests}]]
    [[dests := {(rip,rsn) | (rip,rsn) ∈ dests ∧ precs(rt,rip) ≠ ∅}]]
    groupcast(pre,rerr(dests,ip)) . AODV(ip,sn,rt,rreqs,store)
  + [ dip ≠ ip ]      /* this node is not the destination node */
  (
    [ dip ∈ vD(rt) ∧ dsn ≤ sqn(rt,dip) ∧ sqnf(rt,dip) = kno ]      /* valid route to dip that is fresh enough */
    /* update rt by adding precursors */
    [[rt := addpreRT(rt,dip,{sip})]]
    [[rt := addpreRT(rt,oip,{nhop(rt,dip)})]]
    /* unicast a RREP towards the oip of the RREQ */
    unicast(nhop(rt,oip),rrep(dhops(rt,dip),dip,sqn(rt,dip),oip,ip)) .
```

# Analysing Variants of AODV



```
+ [ (oip, rreqid) ∉ rreqs ]      /* the RREQ is new to this node */
  [[rt := update(rt,(oip,osn,kno,val,hops + 1,sip,0) )]]      /* update the route to oip in rt */
  [[rreqs := rreqs ∪ {(oip,rreqid)}]]      /* update rreqs by adding (oip, rreqid) */
  (
    [ dip = ip ]      /* this node is the destination node */
    [[sn := max(sn,dsn)]]      /* update the sqn of ip */
    /* unicast a RREP towards oip of the RREQ */
    unicast(nhop(rt,oip),rrep(0,dip,sn,oip,ip)) . AODV(ip,sn,rt,rreqs,store)
    ► /* If the transmission is unsuccessful, a RERR message is generated */
    [[dests := {(rip,inc(sqn(rt,rip))) | rip ∈ vD(rt) ∧ nhop(rt,rip) = nhop(rt,oip)}]]
    [[rt := invalidate(rt,dests)]]
    [[store := setRRF(store,dests)]]
    [[pre := ∪ {precs(rt,rip) | (rip,*) ∈ dests}]]
    [[dests := {(rip,rsn) | (rip,rsn) ∈ dests ∧ precs(rt,rip) ≠ ∅}]]
    groupcast(pre,rerr(dests,ip)) . AODV(ip,sn,rt,rreqs,store)
  + [ dip ≠ ip ]      /* this node is not the destination node */
    (
      [ dip ∈ vD(rt) ∧ dsn ≤ sqn(rt,dip) ∧ sqnf(rt,dip) = kno ]      /* valid route to dip that is fresh enough */
      /* update rt by adding precursors */
      [[rt := addpreRT(rt,dip,{sip})]]
      [[rt := addpreRT(rt,oip,{nhop(rt,dip)})]]
      /* unicast a RREP towards the oip of the RREQ */
      unicast(nhop(rt,oip),rrep(dhops(rt,dip),dip,sqn(rt,dip),oip,ip)) .
        unicast(nhop(rt,dip),rrep(hops + 1,oip,osn,dip,ip)).
```

# Non-Optimal Route Selection




- during route discovery *only* nodes lying on route from source to destination find optimal routes

[MiskovicKnightly10]

- modification: forward route request
  - only a few more messages on average
  - specification changes in **5 lines** only
  - invariants to be checked only for these lines

# Failure of Route Discovery Process




- route replies are dropped if they do not carry new information; this might yield route discovery failure

[IETF Mailing List]

- modification: forward route request
  - only a few more messages in average
  - specification changes in **1 line** only (2 lines can be dropped)
  - invariants to be checked only for this line (it is easy to adapt the proof)



# Failure of Route Discovery Process




- route replies are dropped if they do not carry new information; this might yield route discovery failure

[IETF Mailing List]

- modification: forward route request
  - only a few more messages in average
  - specification changes in **1 line** only (2 lines can be dropped)
  - invariants to be checked only for this line (it is easy to adapt the proof)

# Failure of Route Discovery Process



- route replies are dropped if they do not carry new information; this might yield route discovery failure  
[IETF Mailing List]
- modification: forward route request
  - only a few more messages in average
  - specification changes in **1 line** only (2 lines can be dropped)
  - invariants to be checked only for this line  
(it is easy to adapt the proof)

- So far concentrated on AODV
  - well known
  - IETF standard
- Full concise specification of AODV (RFC 3561)
  - verified/disproved properties (independent of topologies)
  - modular proofs
  - found several ambiguities, contradictions, shortcomings
  - adapted proofs to verify interpretations and variants of AODV

- Extend formal methods to other protocols
  - OSLR, B.A.T.M.A.N., ...
- Add further necessary concepts
  - time
  - probability (links, (quantitative) measurements)
  - quality



From imagination to **impact**





From imagination to **impact**




AODV




DEMO

# AODV – An Example




s is looking for a route to d


# AODV – An Example




# AODV – An Example



# AODV – An Example




# AODV – An Example






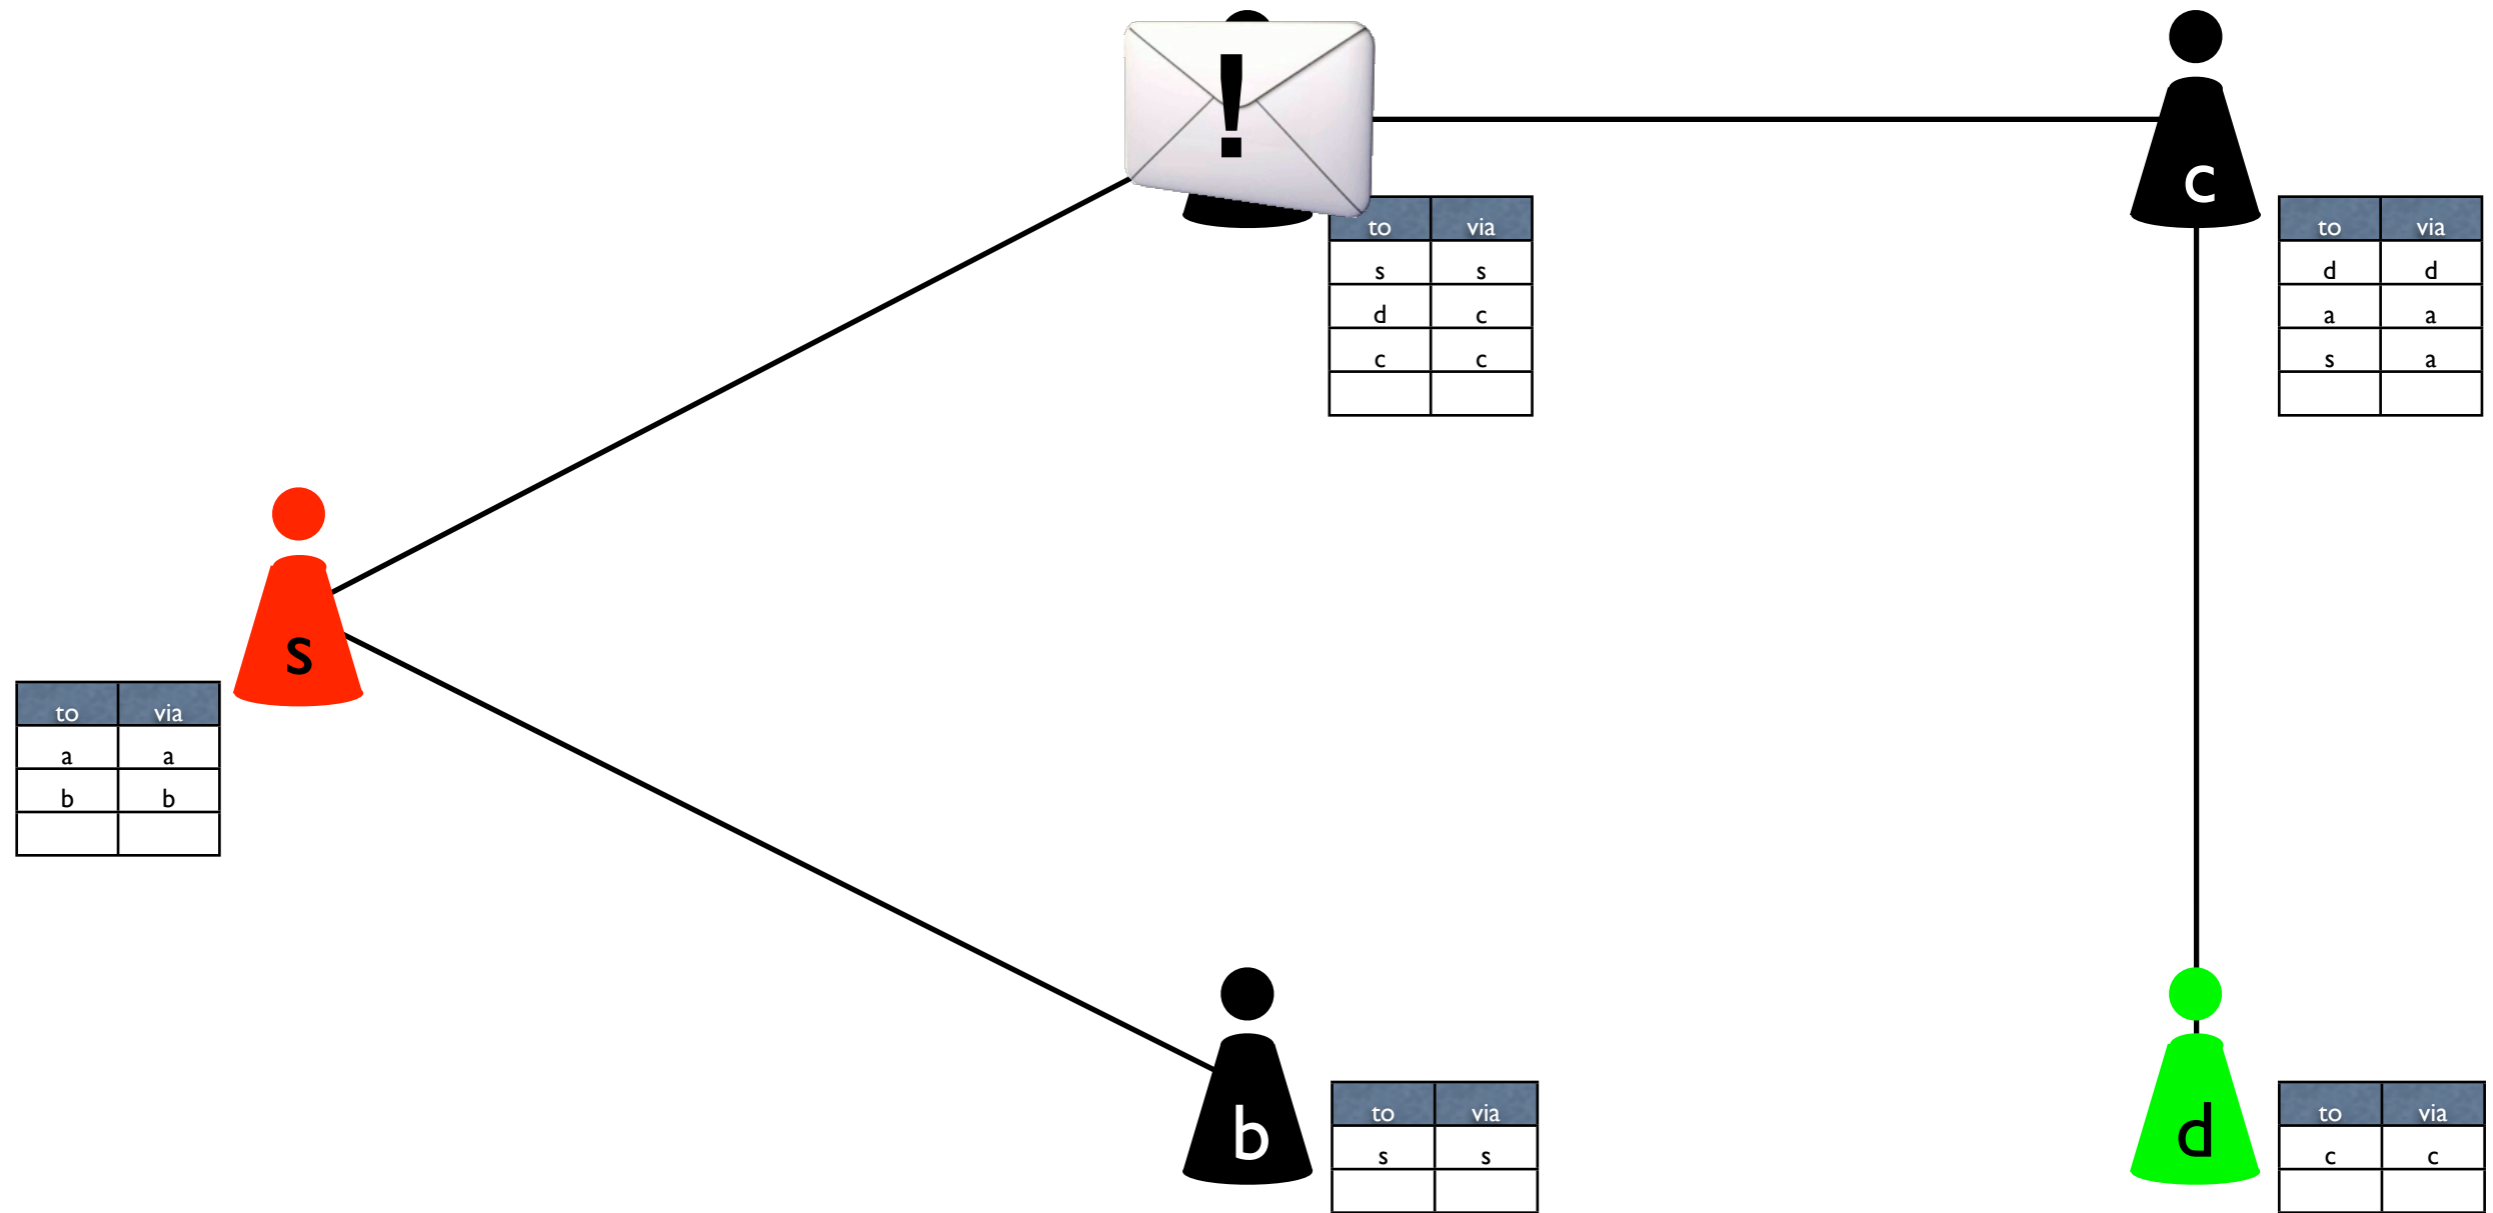
# AODV – An Example




# AODV – An Example



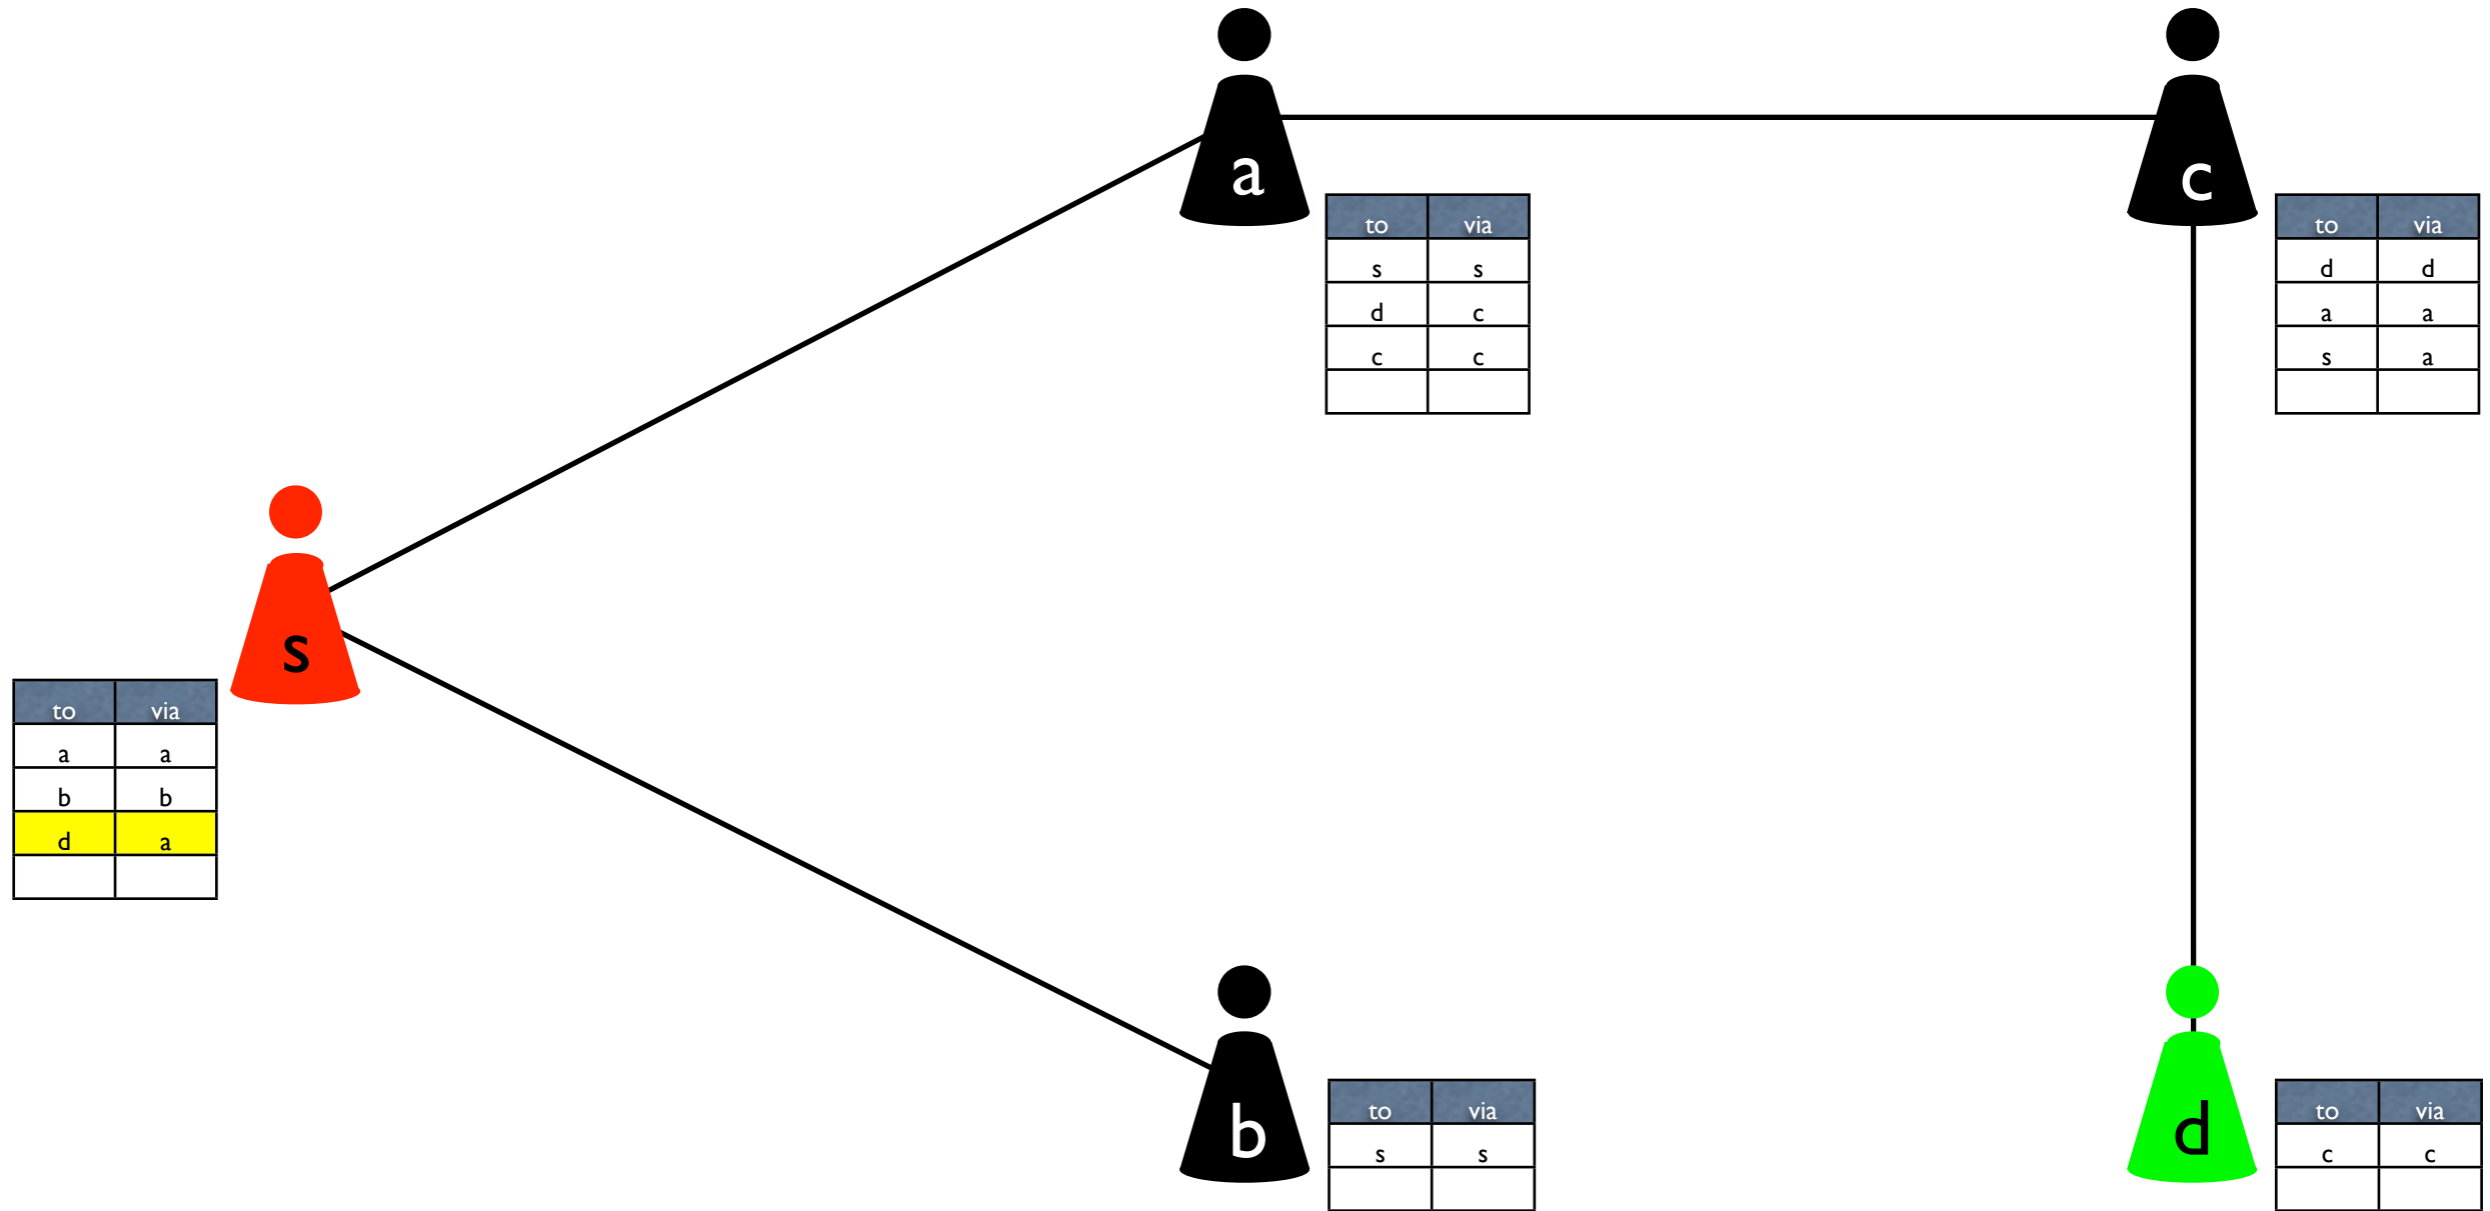
# AODV – An Example



# AODV – An Example




# AODV – An Example






# AODV – An Example



# AODV – An Example



s has found a route to d

# Ambiguities and Loop Freedom

<b>1. Updating the Unknown Sequence Number in Response to a Route Reply</b>		
1a.	the destination sequence number (DSN) is copied from the RREP message (Sect 6.7)	decrement of sequence numbers and loops
1b.	the routing table is not updated when the information inside is “fresher” (Sect. 6.1)	loop free
<b>2. Updating with the Unknown Sequence Number (Sect. 6.5)</b>		
2a.	no update occurs	loop free, but opportunity to improve routes is missed.
2b.	overwrite any routing table entry by an update with an unknown DSN	decrement of sequence numbers and loops
2c.	use the new entry with the old DSN	loop free
<b>3. More Inconclusive Evidence on Dealing with the Unknown Sequence Number (Sect. 6.2)</b>		
3a.	update when <i>incoming</i> sequence number is unknown	supports Interpretations 2b or 2c above
3b.	update when <i>existing</i> sequence number is unknown	decrement of sequence numbers and loops
3c.	update when no <i>existing</i> sequence number is known	supports Interpretation 2a above
<b>4. (Dis)Allowing Self-Entries</b>		
4a.	allow self-entries	loop free if used with appropriate <b>invalidate</b>
4b.	disallow self-entries; if self-entries would occur, ignore mess.	loop free
4c.	disallow self-entries; if self-entries would occur, forward	loop free
<b>5. Storing the Own Sequence Number</b>		
5a.	store sequence number as separate value	loop free
5b.	store sequence number inside routing table	excludes non-trivial self-entries (4b–c)
<b>6. Invalidating Routing Table Entries in Response to a RERR message</b>		
6a.	copy DSN from RERR message (Sect. 6.11)	decrement of sequence numbers and loops (when allowing self-entries (Interpretation 4a))
6b.	no action if the DSN in the routing table is larger than the one in the RERR mess. (Sect. 6.1 & 6.11)	loops (when allowing self-entries)
6c.	take the maximum of the DSN of the routing table and the one from the RERR message	loops (when allowing self-entries)
6d.	take the maximum of the increased DSN of the routing table and the one from the RERR mess.	loop free

**Table 2: Analysis of Different Interpretations of the RFC 3561 (AODV)**

# Update Function



$$\text{update}(rt, r) := \begin{cases} rt \cup \{r\} & \text{if } \pi_1(r) \notin \mathbf{kD}(rt) & //r \text{ is new} \\ nrt \cup \{nr\} & \text{if } \mathbf{sqn}(rt, \pi_1(r)) < \pi_2(r) & //fresher \\ nrt \cup \{nr\} & \text{if } \mathbf{sqn}(rt, \pi_1(r)) = \pi_2(r) \\ & \wedge \mathbf{dhops}(rt, \pi_1(r)) > \pi_4(r) & //shorter \\ nrt \cup \{nr\} & \text{if } \mathbf{sqn}(rt, \pi_1(r)) = \pi_2(r) & //replaces \\ & \wedge \mathbf{flag}(rt, \pi_1(r)) = \mathbf{inv} & \text{invalid} \\ nrt \cup \{nr'\} & \text{if } \pi_2(r) = 0 & //unk. sqn \\ nrt \cup \{ns\} & \text{otherwise ,} \end{cases}$$

**Pro. 1, Lines 10, 14, 18:** The entry  $(\mathbf{sip}, 0, \mathbf{val}, 1, \mathbf{sip}, \emptyset)$  is used for the update; its destination is  $dip := \mathbf{sip}$ . We assume this entry is actually inserted in the routing table of  $ip$ . Since  $dip = \mathbf{sip} = \mathbf{nhop}_N^{ip}(dip) = \mathbf{nhip}$ , the antecedent of the invariant to be proven is not satisfied.

**Pro. 2, Line 5:** The entry  $(\mathbf{oip}, \mathbf{osn}, \mathbf{val}, \mathbf{hops}+1, \mathbf{sip}, *)$  is used for the update; again we assume it is inserted into the routing table of node  $ip$ . So  $dip := \mathbf{oip}$ ,  $\mathbf{nhip} := \mathbf{sip}$ ,  $\mathbf{nsqn}_N^{ip}(dip) := \mathbf{osn}$  and  $\mathbf{dhops}_N^{ip}(dip) := \mathbf{hops}+1$ . This information is distilled from a received route request message (cf. Lines 1 and 8 of Pro. 1). By Proposition 7.1 of [8], this message was sent before, say in state  $N'$ ; by Proposition 7.8 of [8], the sender of this message has identified itself correctly, and is  $\mathbf{sip}$ .