

Provable security for ad hoc network routing protocols

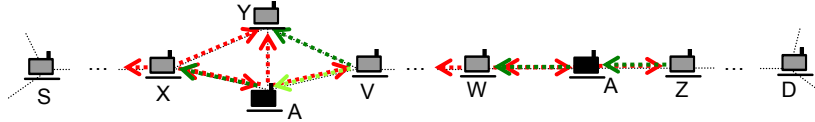
motivation for a rigorous analysis framework;
simulation-based approach:

- real world model;
- ideal world model;
- definition of security;
- proof technique;
- proof of endairA;

Provable security for ad hoc network routing protocols

- several “secure” routing protocols have been proposed for wireless ad hoc networks
 - SRP, Ariadne, S-AODV, ARAN, SEAD, ...
- their security have been analyzed mainly by informal means
- informal reasoning about security protocols is prone to errors
 - lessons learnt in the field of key exchange protocols
 - some attacks have been found against SRP, Ariadne, and S-AODV
- we need more assurances
 - mathematical models
 - precise definitions
 - sound proof techniques

An attack on Ariadne



$X \rightarrow * : [\text{RREQ}, S, D, \text{id}, h_X, (\dots, X), (\dots, \text{mac}_{XD})]$
 $A \rightarrow * : [\text{RREQ}, S, D, \text{id}, *, (\dots, X, A), (\dots, \text{mac}_{XD}, h_X)]$
 \dots
 $W \rightarrow * : [\text{RREQ}, S, D, \text{id}, *, (\dots, X, A, V, \dots, W), (\dots, \text{mac}_{XD}, h_X, \dots, \text{mac}_{WD})]$
 $A : h_A = H(A \parallel h_X)$
 $A \rightarrow * : [\text{RREQ}, S, D, \text{id}, h_A, (\dots, X, A), (\dots, \text{mac}_{XD}, \text{mac}_{AD})]$
 \dots
 $Z \rightarrow A : [\text{RREP}, D, S, (\dots, X, A, Z, \dots), \text{mac}_{DS}]$
 $A \rightarrow W : [\text{RREP}, D, S, (\dots, X, Y, V, \dots, W, A, \dots), \text{mac}_{DS}]$
 \dots
 $V \rightarrow Y : [\text{RREP}, D, S, (\dots, X, Y, V, \dots, W, A, \dots), \text{mac}_{DS}]$
 $A \rightarrow X : [\text{RREP}, D, S, (\dots, X, A, Z, \dots), \text{mac}_{DS}]$
 \dots
 $? \rightarrow S : [\text{RREP}, D, S, (\dots, X, A, Z, \dots), \text{mac}_{DS}]$ **(a non-existent route!)**

Mathematical framework

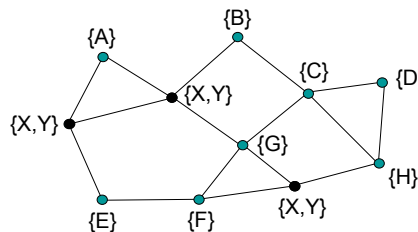
- based on the simulation paradigm
 - real-world model
 - describes the real operation of the protocol
 - ideal-world model
 - captures what the protocol wants to achieve in terms of security
 - definition of security in terms of indistinguishability of the two models from the point of view of honest participants

Mathematical framework (cont'd)

- communication model
 - multi-hop communication and the broadcast nature of radio channels are explicitly modeled
- adversary model
 - power of the adversary is limited
 - it has communication capabilities similar to regular nodes
 - it cannot fully control when some nodes send and receive messages
- model of computation
 - computation is not scheduled by the adversary
 - computation is performed in rounds (synchronous model), but ...
 - knowledge of the current round number is never exploited
- ideal-world model and ideal-world adversary
 - they are essentially the same as the real-world model and adversary
 - the ideal world is ideal in the following sense:
 - route reply messages that contain incorrect routes are marked and filtered out
 - incorrect routes are never returned in the ideal world

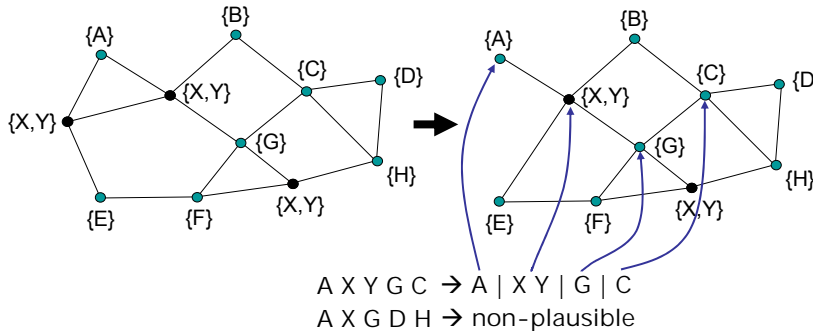
Configuration

- an ad hoc network is represented by a graph $G(V, E)$
 - V : vertices are network nodes (honest and adversarial)
 - E : edges represent communication links (radio or wormhole)
- $V^* \subset V$ is a set of distinguished nodes (under the adversary's control)
- \mathcal{L} is a labeling function (assigns IDs to nodes) with the following restrictions:
 - each honest node has a unique, uncompromised ID
 - each adversarial node is labeled with *all* the compromised IDs
 - we assume that ID's are authenticated during neighbor discovery (Sybil attack is excluded)
- a *configuration* is a triplet: (G, V^*, \mathcal{L})



Plausible routes

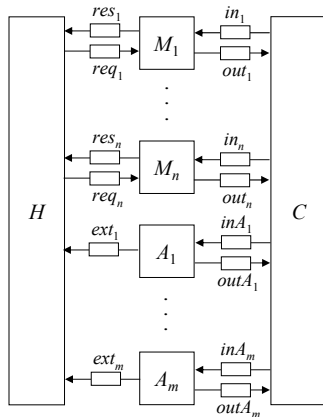
- reduced configuration: $(\underline{G}(\underline{V}, \underline{E}), \underline{V}^*, \underline{\mathcal{L}})$
 - neighboring adversarial nodes are joined
- a route is *plausible* in a given configuration, if it doesn't contain repeating IDs and it can be partitioned in a way that each partition P can be associated with a node v in \underline{G} such that
 - $P \subseteq \underline{\mathcal{L}}(v)$, and
 - neighboring partitions are associated with neighboring nodes in \underline{G}



The rationale behind plausible routes

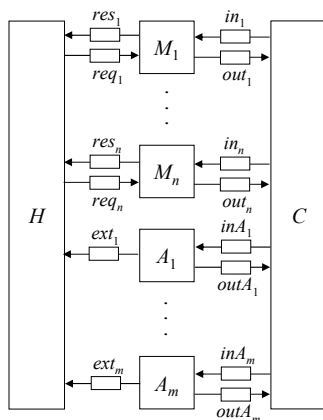
- adversarial nodes can emulate the execution of the routing protocol (locally) using any subset of the compromised IDs in any order
- they can also pass information to each other in a proprietary way
- these are *tolerable imperfections*, which are embedded in the notion of plausible routes

Real-world model (1)



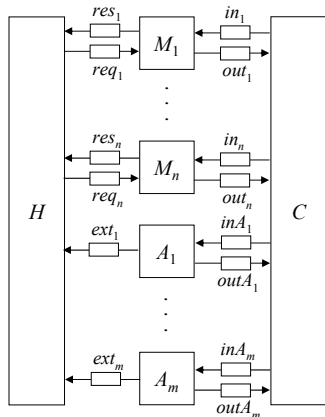
- $H, M_1, \dots, M_n, A_1, \dots, A_m, C$ are interacting, probabilistic Turing machines
 - M_1, \dots, M_n represent honest nodes in \underline{G}
 - A_1, \dots, A_m represent adversarial nodes in \underline{G}
 - C models the communication links (edges of \underline{G})
- each machine is initialized with some input data (e.g., crypto keys) and some random input
- each machine operates in a reactive manner (must be activated)
 - reads input tape
 - performs state transition and writes output tape
 - goes back to sleep
- machines are activated by a hypothetical scheduler in rounds in a fix order in each round: H, \dots, C
- the computation ends when H reaches a final state

Real-world model (2)



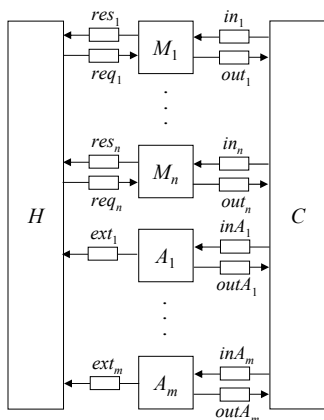
- C models the communication links
 - when activated, it moves the content of the output tape of each protocol machine (M_i and A_j) onto the input tape of all neighboring machines in \underline{G} (in a random order)
- H models higher layer protocols (and ultimately the end-users) of non-corrupted nodes
 - it can initiate a route discovery process at any machine M_i by placing a request on req_i
 - a response may be returned to the request via res_i
 - the response contains a set of routes (maybe empty set)
 - it can receive out-of-band requests from the adversarial machines via ext_j

Real-world model (3)



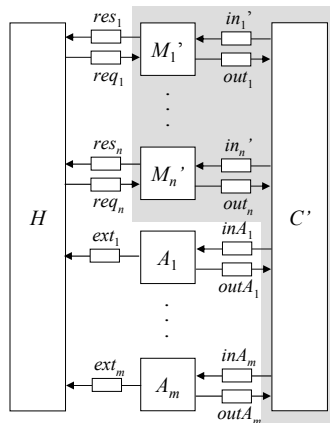
- M_i models the operation of the routing algorithm in the i -th non-corrupted node
 - it receives requests from H via req_i and may return a response via res_i
 - it sends and receives routing messages to and from its neighbors via out_i and in_i
 - initialized with its own ID and those of its neighbors, some cryptographic material, and random input
- A_j models the j -th adversarial node
 - it uses out_{A_j} and in_{A_j} to communicate with its neighbors
 - it can use ext_j to “force” H to start a route discovery between any two *honest nodes*
 - it is *non-adaptive*: it places its requests on ext_j at the beginning of the computation, and doesn’t use ext_j anymore
 - its behavior is not restricted apart from being polynomial-time in the security parameter

Real-world model (4)



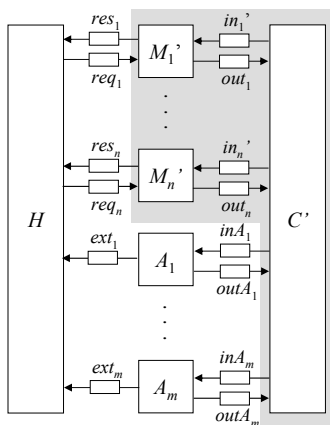
- output of the real-world model
 - sets of routes returned to H
 - denoted by $real_out_{conf, \mathcal{A}}(r)$, where $r = (r_p, r_M, r_A, r_C)$
 - r_I – random input of cryptographic initialization (key generation)
 - r_M – random input of M_1, \dots, M_n
 - r_A – random input of A_1, \dots, A_m
 - r_C – random input of C
 - $real_out_{conf, \mathcal{A}}$ denotes the random variable describing the output when r is chosen uniformly at random

Ideal-world model (1)



- difference between C and C' :
 - C' marks every route reply message that contains a non-plausible route as corrupted before placing it on the input tape in_i' of a non-corrupted protocol machine M_i
 - otherwise C' works in the same way as C
- difference between M_i and M_i' :
 - when M_i' receives a route reply message that belongs to a route discovery process initiated by itself, it processes the message as follows:
 - it performs all the verifications required by the routing protocol
 - if the message passes all verifications, then it also checks the corruption flag attached to the message
 - if the message is corrupted (contains a non-plausible route), then M_i' drops the message
 - otherwise M_i' behaves as M_i

Ideal-world model (2)



- output of the ideal-world model
 - sets of routes returned to H
 - denoted by $ideal_out_{conf, \mathcal{A}}(r')$, where $r' = (r'_p, r'_{M_i}, r'_A, r'_C)$
 - $ideal_out_{conf, \mathcal{A}}$ denotes the random variable describing the output when r' is chosen uniformly at random

Definition of (statistical) security

A routing protocol is said to be (statistically) **secure** if, for any configuration $conf$ and any real-world adversary \mathcal{A} , there exists an ideal-world adversary \mathcal{A}' , such that

$$real_out_{conf,\mathcal{A}} =^s ideal_out_{conf,\mathcal{A}'}$$

where $=^s$ means **statistically indistinguishable**.

notes:

- two random variables are statistically indistinguishable if the L_1 distance of their distributions are negligibly small
- if this definition is satisfied by a protocol, then a non-plausible route can be returned in the real system only with negligible probability (for every configuration and arbitrary adversary)

Proof technique

- let $\mathcal{A}' = \mathcal{A}$
- if, for a given r , no message is dropped due to its corruption flag in the ideal-world model, then the ideal-world model perfectly simulates the real-world model:

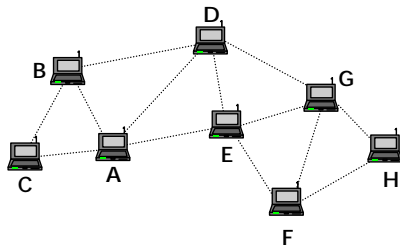
$$real_out_{conf,\mathcal{A}}(r) = ideal_out_{conf,\mathcal{A}}(r)$$

- if, for some r , there exist messages that are dropped due to their corruption flag in the ideal-world model, then there may be a *simulation failure*:

$$real_out_{conf,\mathcal{A}}(r) \neq ideal_out_{conf,\mathcal{A}}(r)$$

- in proofs, we want to show that simulation failures occur with negligible probability
- if this is not the case, then
 - in theory, we haven't proven anything (there may be another $\mathcal{A}' \neq \mathcal{A}$, for which we have statistical indistinguishability)
 - in practice, there's a problem with the protocol

Reminder on endairA



target verifies:

- there's no repeating ID in the node list
- last node in the node list is a neighbor

each intermediate node verifies:

- its own ID is in the node list
- there's no repeating ID in the node list
- next and previous nodes in the node list are neighbors
- all signatures are valid

source verifies:

- there's no repeating ID in the node list
- first node in the node list is a neighbor
- all signatures are valid

$A \rightarrow * : [RREQ, A, H, id, ()]$

$E \rightarrow * : [RREQ, A, H, id, (E)]$

$F \rightarrow * : [RREQ, A, H, id, (E, F)]$

$H \rightarrow F : [RREP, A, H, id, (E, F), (sig_H)]$

$F \rightarrow E : [RREP, A, H, id, (E, F), (sig_H, sig_F)]$

$E \rightarrow A : [RREP, A, H, id, (E, F), (sig_H, sig_F, sig_E)]$

Analysis of endairA (1)

Theorem:

endairA is statistically secure if the signature scheme is secure against chosen message attacks (and remote adversarial nodes cannot collaborate).

sketch of the proof:

- it is enough to prove that, for any configuration $conf$ and attacker \mathcal{A} , a route reply message in the ideal-world system is dropped due to its corruption flag set to true with negligible probability
- let us suppose that the following message is dropped due to its corruption flag:

$$[RREP, S, D, (N_1, N_2, \dots, N_p), (sig_D, sig_{N_p}, \dots, sig_{N_1})]$$

- we know that

- there are no repeating IDs in $(S, N_1, N_2, \dots, N_p, D)$
- N_1 is a neighbor of S
- all signatures are valid
- S and D are honest
- $(S, N_1, N_2, \dots, N_p, D)$ is a non-plausible route in \underline{G}

- we prove that \mathcal{A} must have forged a signature to achieve this

Analysis of endairA (2)

sketch of the proof (cont'd):

- in the reduced configuration adversarial nodes are non-adjacent
- thus each sequence of non-repeating IDs has a unique partitioning
 - IDs of honest nodes form distinct partitions
 - consecutive adversarial IDs form a partition
- if the route is non-plausible, then (at least) one of the following must hold:
 - $P_j = \{N_i\}$ and $P_{j+1} = \{N_{i+1}\}$ are non-adversarial partitions and the nodes v and v' that belong to N_i and N_{i+1} are not adjacent in \underline{G}
 - $P_j = \{N_i\}$, $P_{j+1} = \{N_{i+1}, \dots, N_{i+k}\}$, $P_{j+2} = \{N_{i+k+1}\}$ are two non-adversarial (P_j , P_{j+2}) and an adversarial partition (P_{j+1}) and the nodes that belong to N_i and N_{i+k+1} have no common neighbor that belongs to V^*
- in the first case, the node that uses N_i would detect that the next ID in the list doesn't belong to a neighbor and wouldn't sign the message

Analysis of endairA (3)

▪ sketch of the proof (cont'd):

– in the second case:

- assume the adversary didn't forge any signatures
- the node using N_i must have received
[RREP, S , D , (N_1, N_2, \dots, N_p) , $(sig_D, sig_{N_p}, \dots, sig_{N_{i+1}})$]
from an adversarial node, say A (why?)
- A must have received
[RREP, S , D , (N_1, N_2, \dots, N_p) , $(sig_D, sig_{N_p}, \dots, sig_{N_{i+k+1}})$]
from N_{i+k+1} (why?)
- A must be a common neighbor of N_i and N_{i+k+1} , which is a contradiction

→ the adversary must have forged some signatures

Summary

- attacks against secured ad hoc network routing protocols exist
- flaws are subtle and difficult to discover by informal analysis
- the simulation-based analysis approach used in cryptography can be adopted for reasoning about the security of ad hoc network routing protocols
 - we showed this for on-demand source routing protocols, but the same ideas work for other types of protocols too
- unfortunately, hand-written proofs are tedious and prone to errors
- open question: How to automate the case analysis in proofs?