

Formal Modeling and Analysis of Bluetooth 4.0 Pairing Protocol

David Jia
Stanford University
djia@stanford.edu

Richard Hsu
Stanford University
richhsu@stanford.edu

ABSTRACT

Bluetooth is a wireless technology for exchanging data over short distances and is built into many of the devices used in daily life such as smartphones and laptops. Bluetooth allows communication between paired devices. The pairing is done through Bluetooth’s pairing protocol known as Simple Secure Pairing (SSP) which has been established since Bluetooth 2.1+. Bluetooth security is important because sensitive and confidential information such as phone conversations, text-based messages, and key strokes on a keyboard are often communicated through a pairing. So compromising Bluetooth could potentially compromise user privacy and confidential information. We employ a formal model to verify the security properties of the pairing protocol used in Bluetooth 4.0, the latest iteration of the protocol. Using our formal Murphi model, we demonstrate and confirm previously known attacks on Bluetooth that have not yet been formalized. We then discuss a new attack found by our model as well as its implications to Bluetooth 4.0 security. Finally, we discuss and recommend possible fixes that could be employed to avoid the attacks we have found. Through this formal modeling we expand on the security verification of Bluetooth pairing protocol.

1. INTRODUCTION

Bluetooth is a wireless technology that is ubiquitous in many modern applications. Users rely on it to transfer data between smartphones, laptops, on-board car systems, and headsets. Bluetooth technology also exists in more peripheral items such as printers, mice, keyboards, and speakers. Because Bluetooth serves as an underlying systems for so many communication systems that are used on a daily basis, sensitive and confidential information can often times be passed through the supposedly secure paired channel. For example, a security breach on Bluetooth could potentially lead to eavesdropping on phone conversations and keylogging password entries. Thus, we analyze Bluetooth in this paper through a formal modeling and analysis in hopes of shedding light on its security and providing recommendations to improve its security where possible.

In this paper, we focus our attentions on the Bluetooth 4.0 Protocol [12], which is the most recent version of the technology at the time of writing.

2. BLUETOOTH OVERVIEW

Bluetooth is a wireless technology standard used for exchanging data over a short distance. It utilizes short-wavelength

radio transmission as the underlying communication channel and is implemented in devices ranging from hands-free headsets to computer peripherals such as a mouse or keyboard. These devices are fairly ubiquitous today. Some examples of the data transferred between devices via Bluetooth include inputs from a wireless keyboard, contacts transferred between phones, files transferred between computers, and even patient health data from medical sensors to servers. Each Bluetooth responder device (such as a laptop) can be paired with up to seven Bluetooth initiator devices (such as a keyboard or mouse).

Because of the wide range in the application of Bluetooth, the security of the Bluetooth communication is vitally important. The major part of Bluetooth that becomes vulnerable to attackers is the the pairing protocol when two or more devices first initiate connections with one another. The connection process creates a personal area network and requires authentication and encryption to remain secure. Since Bluetooth 2.1, Secure Simple Pairing (SSP) [11] has been implemented and further improved in subsequent versions leading to the current version of Bluetooth 4.0 discussed in the current paper.

2.1 Secure Simple Pairing Protocol

Bluetooth 4.0 Secure Simple Pairing contains four main association models depending on the devices involved:

1. Just Works (JW) When at least one device does not have display nor input channel. Example: wireless headset.
2. Numeric Comparison (NC) When both devices have a display and at least one has a binary input channel for a yes or no response. Example: pairing between smartphones.
3. Passkey Entry (PE) When one device has an input channel but no display and the other device has a display but no input channel. Example: Wireless keyboard.
4. Out of Band (OoB) When both devices support a common additional wireless or wired communication technology for purpose of device discovery or cryptographic channels. Example: Near Field Communication (NFC).

The SSP protocol is completed in five phases. SSP follows the same steps in these five phases for all four association models with the exception of phase 2. The process of pairing through the five phases is described in detail below.

Note that in our analysis, we will only consider the JW and NC models since the PE model is just a special case of the NC model and OoB's security is dependent on the security of an out of band cryptographic channel.

2.1.1 Phase I: Public Key Exchange

In this phase, Elliptic Curve Diffie-Hellman (ECDH) key exchange [14] is completed by the initiating (Device A) and non-initiating, or responding device (Device B). Both devices generate a public-private key pair using ECDH. The public key of each device is sent to the other device. Note that according to the official Bluetooth Specifications, this key pair does not need to be generated each time a pairing occurs. Each device may discard its key pair and generate a new one at any time, but this is not a requirement.

2.1.2 Phase II: Authentication Stage 1

In this stage, each of the four associated models has a different, but relatively similar authentication process. In each of the associated models, both devices generate a random nonce, N_a and N_b , respectively for Device A and Device B. Device B, the non-initiating device then computes a commitment value $C_b = f_1(PK_b, PK_a, N_b, 0)$, where PK_a is the public key for Device A and PK_b is the public key for Device B, and f_1 is SHA256 dependent one-way hash-function that generates a 128-bit value with the given input.

After C_b is generated, it along with N_b is sent to Device A. Device A also sends N_a to Device B. Next, Device A computes the same commitment value $C_a = f_1(PK_b, PK_a, N_b, 0)$ and compares it with C_b . If they do not equal, then the pairing process is aborted. If they are equal, then the process proceeds, and diverges for each of the four associated models as follows until the end of Phase 2.

1. Just Works (JW) In this associated model, since one of the devices can neither display nor input values, Phase 2 ends here and moves on to Phase 3.
2. Numeric Comparison (NC) In this associated model, Phase 2 goes on to generate two verification codes to be displayed by the user. Each device x generates a six-digit verification code $V_x = g(PK_a, PK_b, N_a, N_b)$, and displays it to the user. The user can then determine if $V_a = V_b$ and confirm the pairing and moves on to Phase 3.
3. Passkey Entry (PE) In PE, the two devices decide upon a secret value, created by the users of the devices. If we are pairing two devices that both do not have screens, then the values are agreed upon by the user. If one device has a display, then the device generates a random value, which is shown on the device's display to be inputted into the other (input) device. This shared secret is turned into an n -bit number (n depends on the length of the shared secret). Let the i -th bit for device x be denoted as r_{x_i} . Since the secret is shared, the assumption is that all $r_{a_i} = r_{b_i}$ for all i . Then for each i , a commitment value, $C_{x_i} = f_1(PK_b, PK_a, N_{x_i}, r_{x_i})$ is generated for device x . The twenty commitment values are sent to the other device and compared. If any commitment values do not equal, then we abort the process. Otherwise, we move on to Phase 3.

4. Out of Band (OoB) In OoB, after the commitment stage described completes, a shared secret comparison is done similar to the PE mode, except the external communication is complete through an external channel such as NFC.

2.1.3 Phase III: Authentication Stage 2

Phase 3 confirms that both devices have successfully completed the exchange during pairing. Device A generates exchange codes $E_a = f_3(DHKey, N_a, N_b, r_b, IOcapA, A, B)$ and Device B generates a corresponding exchange code $E_b = f_3(DHKey, N_b, N_a, r_a, IOcapB, B, A)$, where $DHKey$ is the Diffie-Hellman key created in Phase 1. Both are sent to the other device. Each device then computes the exchange code for the other device and verifies that it is the same.

2.1.4 Phase IV: Link Key Calculation

A Link Key is calculated by both devices by computing $LK = f_2(DHkey, N_A, N_B, "btlk", BD_{ADDR_A}, BD_{ADDR_B})$. Since both devices have been authenticated at this stage, the link key generated by both devices should be equal. This link key will be used in phase 5.

2.1.5 Phase V: LMP Authentication and Encryption

Phase 5 is the actual communication between the paired devices. The Link Manager Protocol (the controller that handles authentication and encryption) utilizes the link key developed in Phase 4 to perform authentication and encryption during the communication between the paired devices. For example if a device wants to verify a paired device it can perform a challenge-response by sending a nonce to the claimant device in which it generates a value and sends back to the verifier who can then verify the identity of the claimant device.

3. RELATED WORKS

Because of Bluetooth's prevalence and the emergence of mobile technology, several related works have been done in an effort to analyze the security of Bluetooth. Most recently, in 2010, Phan and Mingard [15] analyzed SSP in Bluetooth 4.0 in "Analyzing the Secure Simple Pairing in Bluetooth v4.0" which hand modeled and analyzed Bluetooth 4.0 and described three Man in the Middle (MitM) attacks for the NC and PE modes. In their 2007 paper, "Formal Analysis of Authentication in Bluetooth Device Pairing", Chang and Shmatikov [2] modeled Bluetooth 2.0 using ProVerif cryptographic protocol verifier and verified a known guessing attack and discovered a potential vulnerability which produces concurrent execution of authentication. Similarly, Haataja and Toivanen [6] discussed in "Practical Man-in-the-Middle Attacks Against Bluetooth Secure Simple Pairing" two more MitM attacks on Bluetooth 2.0 handsets and hand-free devices as well as modeled a vulnerability on the OoB mode given that the attacker has visual contact to the victim devices. Other papers are more general and informative, presenting a high-level view of the Bluetooth protocol, or suggesting potential vulnerabilities and risks associated with Bluetooth without providing formal analysis or proof.

While these papers present a good starting point for modeling Bluetooth 4.0, they are insufficient. First, many of the analysis of Bluetooth are either not formal or performed on

outdated versions, such as v2.x [17, 5, 10, 9]. The most up-to-date analysis of Bluetooth 4.0 by Phan and Mingard [15] encapsulated a non-formal analysis of the NC and PE modes without formal modeling with a cryptographical tool such as Murphi.

In our paper, we perform a formal analysis and modeling (with Murphi) of Bluetooth 4.0, which has several improved and modified security features from its predecessors. We first confirm the findings of other papers, such as that of Phan and Mingard, by modeling MitM attacks on Bluetooth in JW and NC modes using Murphi. We then describe a novel attack on the NC association model found by our formal Murphi modeling.

4. MURPHI MODEL CHECKER

We use the Murphi Model Checker [1] in order to model the states of the Bluetooth 4.0 pairing protocol. The Murphi Model Checker allows us to model the different states possible in the protocol and be able to model the packets being sent between device parties while modeling an adversary that has certain access rights to these packets depending on the protocol mode and encryption methods of the packets. We begin by modeling two devices being paired and then from there model multiple devices connecting to a single parent device. Bluetooth only allows at most seven devices paired to a single responder device, therefore scaling can be done in a natural way and we can model larger ecosystems of Bluetooth devices given the right computational resources.

5. BLUETOOTH MODELING

In order to formally model the Bluetooth 4.0 Pairing Protocols we utilize the Murphi modeling tool. In Murphi, the variables represent the values being passed throughout the communication of the pairing as well as the hash values being created by the devices. The states are modeled by the state variables representing the actual states during the pairing protocol as described by the Bluetooth 4.0 core specification files. We model transitions between steps of each phase as transitions in our state variables. We model the intruder as being a general active attacker adhering to the Dolev-Yao model [4]. In conjunction with the possible state executions, the Murphi model checker will explore all possible states and verify our security invariants for plausible executions. Our overall model consists of roughly 1000 lines of Murphi code. Please note that for the sake of conciseness, we refrain from showing all of our code in this paper and instead include the entire codebase in separate repository. In this paper, for purposes of understanding our formalization, we will present a high level overview of our model.

5.1 Design Choices

We model our pairing protocol by defining separate models for the different device roles in SSP. The major models are the initiator—the device that begins the pairing protocol, and the responder—the device that with which the initiator is being paired. The message model is utilized to describe the transitions throughout the various phases of the pairing protocol. In general only phases prior to and up until phase IV are modeled using Murphi rulesets because phase V simply involves encryption and the start of message transmission after the Link Key is calculated in phase IV. After phase III,

the link key generation can be done without a transition, given all prior information from phases I to III. Phase IV is represented in our model by denoting two paired devices as “having” a link key by the linkKey boolean variable of initiators and responders (as we will discuss below). Thus, there is no loss of generality of our model because phase IV is the last phase in which any network communication occurs and the final execution stage that verifies the pairing is complete.

5.2 Modeling Protocol Entities

5.2.1 Initiators

In the initiator model, there are seven distinct states defined by the enumeration InitiatorStates. Three states, I_SLEEP, I_SENT_KEY, and I_PHASEONE_DONE describe the initiator’s interactions in phase I of the Bluetooth SSP protocol. The states I_WAIT_ONCE, I_NC_VERIF_SET, and I_PHASETWO_DONE model the initiator’s interactions in phase II. I_WAIT_EVALUE and I_PAIREED model the initiator’s interactions in phase III and IV.

Each initiator record holds the current state the initiator is in, the responder with which the initiator is trying pair, and a set of saved variables passed from the responder to the initiator. Finally, the linkKey boolean variable denotes whether or not the link key was generated.

```
Initiator: record
  state: InitiatorStates;
  responder: AgentId;
  -- recieved public key in phase 1
  responder_pk: AgentId;
  -- recieved rb value if any
  responder_r: AgentId;
  -- recieved nonce in phase 2
  responder_n: AgentId;
  -- received commitment value
  responder_c: CValue;
  -- link key of the pairing generated
  linkKey: boolean;
  -- verification token generated NC Phase 3
  vValue: VValue;
end;
```

5.2.2 Responder

In the responder model there are six distinct states defined by the enumeration ResponderStates. The first two states R_SLEEP and R_PHASEONE_DONE describe the responder’s interaction in Phase I of the protocol. R_WAIT_NONCE, R_NC_VERIF_SET, and R_PHASETWO_DONE model the responder’s interactions in phase II. Finally R_PAIREED models the responder’s interactions in phase III and IV.

Each responder record holds the current state of each pairing, which saves information between a pairing agent and itself, and contains the information received over the network such as the nonce and public key. Finally, the linkKey states whether or not the link key was generated.

```
Pairing: record
  state: ResponderStates;
```

```

initiator:      AgentId;
-- recieved public key in phase 1
initiator_pk:  AgentId;
-- received random value
initiator_r:   AgentId;
-- recieved nonce in phase 2
initiator_n:   AgentId;
-- received commitment value
initiator_c:   CValue;
-- link key of the pairing
linkKey:       boolean;
-- vvalue of the pairing
vValue:        VValue;
end;

Responder: record
  pairings: multiset[MaxInitiators] of Pairing;
end;

```

5.3 Modeling Messages

There are 4 major message types during the protocol pairing:

1. Public key messages (`M_PublicKey`) where the device sends its public key in the clear.
2. Nonce messages (`M_Nonce`) where the device sends a random nonce in the clear.
3. Commitment value message (`M_CommitValue`) where the device generates a commitment hash value based off the public keys and nonces used.
4. Exchange verification message (`M_ExchangeVerif`) where the device generates an exchange hash value based on the public keys, nonces, a random value, their IO capabilities, and the device addresses in use.

Given these states, the Message record is defined as follows:

```

Message : record
  -- message type
  mType:      MessageType;
  -- source of message
  source:     AgentId;
  -- intended destination of message
  dest:       AgentId;

  -- nonce from source to dest
  nonce:      AgentId;
  -- public key
  publickey:  AgentId;
  -- commit value (Phase 2)
  cValue:     CValue;
  -- exchange verification value (Phase 3)
  eValue:     EValue;
end;

```

Each message contains the values exchanged throughout the protocol as described in our overview of Bluetooth. Since the Murphi model checker do not support hash functions, we

simulate it by assuming that the cryptographic hash functions are collision resistant, therefore in order to generate the same hash value, the inputs must be equal. Thus we can safely model hash values as a record of their inputs, and when checking that the hash values are equal, our model simply verifies that all input values are equal as well. Below we present the commitment value, the exchange verification value, and the verification code value which are records of all their inputs in our model.

```

-- Commitment Value
CValue : record
  -- when created - pk of sender
  pk_send: AgentId;
  -- when created - pk of receiver
  pk_rcv:  AgentId;
  -- when created - nonce of sender
  n_send:  AgentId;
end;

-- Exchange Verification Value
EValue : record
  -- when created - pk of sender
  pk_send: AgentId;
  -- when created - pk of receiver
  pk_rcv:  AgentId;
  -- when created - nonce of sender
  n_send:  AgentId;
  -- when created - nonce of receiver
  n_rcv:   AgentId;
  -- when created
  -- random value generated by the receiver
  r_rcv:   AgentId;
end;

-- Verification Value
VValue : record
  -- public key of initiator
  pk_initiator: AgentId;
  -- public key of responder
  pk_responder: AgentId;
  -- nonce sent by initiator
  n_initiator:  AgentId;
  -- nonce sent by responder
  n_responder:  AgentId;
end;

```

5.4 Modeling Protocol Transitions

Transitions in the states of the initiator and responder are modeled using Murphi's rulesets. Each ruleset includes a preposition, which if satisfied, allows the programmer to define the next state. We model each state transition based on the transitions of the SSP protocol.

5.4.1 Initiator and Responder Transitions

As we can see in Fig 1, an initiator in the `I_SLEEP` state starts SSP by sending its public key to an Agent, which can be a responder or an intruder. Another ruleset is created to check for messages sent to each responder, in which a responder will reply with its own public key to the source, which could be the initiator or an intruder.

This process continues into phase II, in which we discuss the model for the JW association model. The other association models are similar enough that we will not discuss them to reduce redundancy. Rulesets are created in which the initiator and responder check the network for messages addressed to it that include the nonce and commitment value. These values are checked against received values by the initiator, and if they are confirmed, the protocol continues, otherwise, the message is removed from the network, and the branch of the Murphi search tree stops. We model the checking of the hashed values by checking that all input values are the same, if any are different then the hash would be different. In a similar fashion, phase III is modeled using rulesets for both initiator and responder in which messages with Exchange Verification values are sent across the network and verified by both initiator and responder. If everything verified, we set the initiator state to `I_PAIRED` and the responder state to `R_PAIRED`. Finally the `linkKey` boolean value is set to true on both the initiator and the responder to represent that the `linkKey` as been exchanged.

5.5 Modeling An Adversary

The adversarial model will follow the Dolev-Yao intruder model [4] in which the adversary is given the power to intercept, read, store, modify, and forward messages of its choosing. The main adversarial models present against Bluetooth are Man-in-the-Middle (MitM) attacks where an adversary is between the communications of two Bluetooth devices. By using the Dole-Yao mode, we can enumerate all states of interest [4].

5.5.1 Adversary Model

An intruder can intercept messages by taking messages from the network and adding it to its multiset of messages. The intruder record also has an array of `linkKeys` to represent the initiators and responders of which the intruder has the `linkKey`. The array of booleans `pk` and `nonce` represent whether the intruder has the public key and nonce from a certain agent. The `sent_pk` and `sent_n` array of booleans represent whether the intruder has sent its own public key and nonce to a certain agent.

```
Intruder: record
  messages: multiset[MaxKnowledge] of Message;
  linkKeys: array[AgentId] of boolean;
  -- Do we know their PK
  pk: array[AgentId] of boolean;
  -- Did we send our PK to them
  sent_pk: array[AgentId] of boolean;
  -- Do we know their nonce
  nonce: array[AgentId] of boolean;
  -- Did we send our nonce
  sent_n: array[AgentId] of boolean;
end;
```

5.5.2 Intercept Messages

The adversarial model against our pairing protocol is mainly a Man-in-the-Middle (MitM) attack in which the adversary has control of the communication channel between the two users. In this case the intruder has three main courses of action: (i) intercept messages, (ii) pick a stored message

and forward it along without modification, (iii) or modify the messages before forwarding it.

```
-- intruder i intercepts messages
ruleset i: IntruderId do
  choose k: net do
    rule 100 'intruder intercepts messages'

      -- Pick a msg from network that is not
      -- intended for the intruder
      !ismember (net[k].source, IntruderId)

    ==>

    begin
      -- the message to intercept
      alias msg: net[k] do
        -- If we haven't seen the message, then add
        -- to intruder's messages
        if multisetcount(1:messages,
          int[i].messages[1] = msg
        ) = 0 then
          multisetadd (msg, int[i].messages);
        end;
      end;
      multisetremove (k, net);
    end;
  end;
end;
```

5.5.3 Forward without Modifications

In this model, if an intruder sees a message on the network that has never been seen, then the intruder records the message. The intruder later forwards this message to any possible agent, which is modeled through different branches of Murphi based on different values selected for the intruder rulesets. In most cases, the source is changed, otherwise the intruder would act as a passive network, in which case we do not actually need the adversary to forward the message. In this attack, an initiator pairs with an intruder using the public key of a responder.

Since the intruder simply forwards messages without changing the public keys, it cannot decode any messages sent between the initiator and responder once they're paired. However, since the initiator and responder are both paired with the intruder, the intruder can perform a denial-of-service attack by simply not forwarding the messages. This will clearly violate the initiator and responder authenticity invariant, since both think they are paired with the other, but are actually paired with an intruder.

```
-- intruder i sends recorded message
-- without modification
ruleset i: IntruderId do
  choose j: int[i].messages do
    ruleset k: AgentId do
      rule 100 'intruder sends recorded message'
        -- Pick a message stored by an intruder
        !ismember (k, IntruderId) &
        !(int[i].messages[j].source = k) &
```

```

    multisetcount (l:net, true) < NetworkSize
==>
-- Set destination to some AgentId,
-- source to us, and send it out
-- without further modifications
var
    outM: Message;
begin
    undefine outM;
    outM      := int[i].messages[j];
    outM.source := i;
    outM.dest  := k;
    multisetadd (outM, net);
end;
end;
end;
end;

```

5.5.4 Modify and Forward

In this adversary model, the intruder can send any message over the communication channel to the other devices. When the intruder receives a message from an initiator, the intruder modifies the public key to its own public key before forwarding it to the intended responder. Similarly, when the intended responder sends its public key to the intruder, the intruder replaces it with the intruder's own public key before forwarding it to the initiator. In this way, both initiator and responder will be paired with the intruder, but unlike the "intercept and forward" model above, the intruder's public key is used. This means that the intruder can decode all messages between the initiator and responder, and contains a link key with both the initiator and responder. This model causes both the authentication and secrecy invariants for both initiator and responder to fail since not only are initiator and responder paired with the intruder, but link keys are shared with the intruder. Note that this MitM attack only occurs in the JW association models since in NC and PK, users need to confirm the verification code before the protocol can continue.

```

-- intruder i modifies messages
-- and sends to network
ruleset i: IntruderId do
    choose j: int[i].messages do
        ruleset k: AgentId do
            rule 100 "intruder sends recorded
                message with its own info"
                -- Pick a message stored by an intruder
                !ismember (k, IntruderId) &
                !(int[i].messages[j].source = k) &
                multisetcount (l:net, true) < NetworkSize
==>
var
    outM: Message;
begin
    undefine outM;
    -- Modify the message depending on
    -- the message type:
    -- M_PublicKey, M_CommitValue,
    -- M_Nonce, or M_ExchangeVerif
    -- Send recorded message with modifications
    multisetadd (outM, net);

```

```

end;
end;
end;
end;

```

5.6 Security Properties and Invariants

In order to verify that the Bluetooth 4.0 Pairing Protocol is secure we must translate the security properties to invariants for the Murphi model. The invariants will be checked at every state being explored by the model checker, and if any invariants are violated, Murphi will return with the error. The following subsections discuss the security properties as well as the invariants we implement.

5.6.1 Authentication

Bluetooth 4.0 provides authentication by providing commitment codes, verification codes to users, and finally an exchange code. All three of these are done in phases 2 and 3 or the authentication steps in SSP. These steps are used to authenticate the two devices being paired. We define two authenticity invariants in our model such that no initiator is paired with an intruder and no responder is paired with an intruder if either of them are in the final paired states I_PAURED and R_PAURED. The invariants are as follows:

```

-- initiator authenticity
invariant "initiator correctly paired with
    good responder"
forall i: InitiatorId do
    ini[i].state = I_PAURED
    ->
    !ismember(ini[i].responder, IntruderId)
end;

-- responder authenticity
invariant "responders correctly paired with
    good initiator"
forall i: ResponderId do
    multisetcount(l:res[i].pairings,
        (res[i].pairings[l].state = R_PAURED &
         ismember(res[i].pairings[l].initiator,
                 IntruderId))
    ) = 0
end;

```

5.6.2 Confidentiality

Confidentiality in Bluetooth 4.0 is provided through the encryption of the communication channel once the link key is established. This link key is refreshed periodically to ensure the confidentiality of communication. However, confidentiality is only upheld once the link key is established, all other communication before this is in the clear. As a consequence, any eavesdropper can access the information being passed between the pairing devices. However, the information being passed is usually used for verification processes and therefore an adversary should not be able to gain all the information required to generate the paired link key. We thus represent this in our model by verifying that the adversary does not gain the link key of any user at any Murphi state. Once the pairing is complete, if the adversary has the link key, then the system has been compromised. Thus, we do not have to worry about the state of confidentiality after

the generation of the link key as we assume all cryptographic primitives such as encryption used in the communication are secure. The invariants in our model therefore check that once the initiators and responders are in the paired states that no intruder owns any of the honest agents' link key.

```

-- initiator confidentiality
invariant 'initiator link key is secret'
forall i: InitiatorId do
  ini[i].state = I_PAURED &
  ini[i].linkKey = true
  ->
  forall j: IntruderId do
    int[j].linkKeys[i] = false
  end
end;

-- responder confidentiality
invariant 'responder link key is secret'
forall j: IntruderId do
  forall i: ResponderId do
    multisetcount(l:res[i].pairings,
      (res[i].pairings[l].initiator = j &
       res[i].pairings[l].state = R_PAURED &
       res[i].pairings[l].linkKey = true &
       int[j].linkKeys[i] = true)
    ) = 0
  end
end;

```

5.6.3 Intention Preservation

Intention preservation is very similar to authenticity. However, the major difference is the intention of the initial pairing. The intention preservation invariant holds if and only if an initiator is paired with who the initiator originally intended to pair with, even if it's an intruder. For example if an initiator intends to pair with an intruder's device then intention preservation, unlike authenticity, still holds since at the end of the pairing the initiator is paired with the intruder's device. In order to verify this property we implement two invariants which check from the initiator's side and the responder's side of the protocol.

```

-- initiator intention preservation
-- every initiator device is paired with who
-- the device originally thought it
-- would be paired with
invariant 'initiator intention
  preservation pairing'
forall i: InitiatorId do
  ini[i].state = I_PAURED &
  ini[i].linkKey = true
  ->
  multisetcount(l:gpr,
    (gpr[l].initiator = i &
     pr[l].responder = ini[i].responder))
  >= 1
end;

-- responder intention preservation

```

```

-- every responder device is paired with who
-- the device originally thought it
-- would be paired with
invariant 'responder intention
  preservation pairing'
forall i: ResponderId do
  multisetcount(j:res[i].pairings,
    res[i].pairings[j].state = R_PAURED) >=
  1
  ->
  multisetcount(l:gpr,
    (multisetcount(j:res[i].pairings,
      res[i].pairings[j].state = R_PAURED &
      gpr[l].responder = i &
      res[i].pairings[j].initiator
        = gpr[l].initiator
    ) >= 1)
  ) >= 1
end;

```

5.7 Protocol Initial State

Our model begins with all states being cleared and the initiator and the responder both in the `I_SLEEP` and `R_SLEEP` states, respectively. All network and pairings are cleared, and the intruder only knows the information for itself and nothing else. The transition rules will then be fired off by the Murphi model checker that push these states to subsequent states, and the validity checks will only continue along plausible paths of execution. From this initial state, since the intruder does not have any information in the beginning it must learn the information through our adversary model, which we presented earlier.

6. ANALYSIS AND RESULTS

After running our Murphi model against our security invariants, our model discovered three major attacks, two of which have been discussed in previous works [15] and a novel attack which we will describe in the subsequent sections. The table in Figure 2 is a table of our model's state exploration results on finding these attacks. We subsequently increase the number of agents in our model and find that we cannot increase passed 2 initiators otherwise our state space increases tremendously that our systems do not have the computational power to complete the model verification. In our analysis, we used an Intel i7 quad 2.8Ghz 8GB RAM server machine.

6.1 Just Works Attack

In the Bluetooth 4.0 protocol using the JW association model our Murphi model verified an existing known MitM attack. A figure of this attack is diagrammed in Figure 3. In this attack the intruder intercepts a message from an initiator that was intended for an honest responder. The intruder then modifies the message with its own public key and sends this to the responder. The responder, not knowing that the intruder is initiating the pairing will respond back with its own public key in the exchange. The intruder then modifies this message and sends it to the initiator thus completing the public key exchange with both initiator and responder. From here the protocol continues as normal but since the JW association model does not have a verification check as in the NC model there is no detection of the MitM attack. There-

fore even though the responder and the initiator have different Diffie-Hellman keys, the pairing will still occur. The intruder in the end has thus gained full control over the connection and can read everything that the initiator and responder send to each other.

Our formal modeling found this attack which broke the authenticity, confidentiality, and intention preservation invariants. This means that in the JW association model an intruder has a communication pairing with both devices, has individual link keys for the individual pairings, as well as broke the intention preservation by intercepting and forcing both devices to pair with itself.

6.2 Numeric Comparison Attack I: The Impersonator

In the last section, we saw that the JW association model was susceptible to MitM attacks, which is one motivation for the NC association model. In the NC model, the users of the devices is presented with a 6-digit verification codes to verify the connection. Thus when the intruder E modifies the messages and exchanges its own public key PK_E with both devices, the verification codes would be $V_A = g(PK_A, PK_E, N_A, N_E)$ and $V_B = g(PK_E, PK_B, N_E, N_B)$ on device A and B respectively. The users would only see Device A which shows V_a and Device B which shows V_b . Since we assume that g is a secure collision resistant hash function, $V_a \neq V_b$ because the inputs are different. Thus a prudent users would cancel the pairing and the MitM attack would be avoided. It is important to note that if there is user error involved and the users accept the verification codes despite the verification code displayed being different, then the MitM attack would succeed. Therefore the NC association model should protect against the MitM attack. However, as we will see, our model demonstrates that the NC model is in fact not secure from all MitM attacks.

To understand the “Impersonator” MitM attack, it is important to note that there are not any proper authentication checks to guarantee that a responder device should in fact be trusted. The SSP protocol relies on the built in exchanges to confirm a pairing, but does not make any extent to actually verify the identity and trustworthiness of the devices. This being said, our model discovered a MitM Impersonator attack in which the intruder impersonates an honest responder, and in turn tricks an honest initiator into pairing with it.

To motivate this attack we first discuss a real-world scenario. In our example, an adversary takes a rental car and replaces the display with its adversarial device or modifies the car in such a way that the device appears to be the actual on-board device. Now, when an honest user enters the car, they are tricked into believing that the device is an honest device when in fact it is an intruder. This scenario is similar to an ATM thief who places an external card reader over the actual card reader, tricking unsuspecting users into swiping their card and thereby compromising confidential data. This attack thus accomplishes an impersonation attack in which the user is tricked into forcing a pairing with the adversary. As the pairing proceeds, it is important to note that the user now sees the display of Device A as well as the impersonated device E, and thus the veri-

fication codes would be $V_A = g(PK_A, PK_E, N_A, N_E)$ and $V_E = g(PK_E, PK_A, N_A, N_E)$. Therefore the unsuspecting user will confirm the pairing and be paired with the intruder device. At this point the honest device responder does not need to be in the picture. However, if the intruder’s device doesn’t have access to certain things such as the on-board features, the intruder could optionally create a connection with the honest device responder and continue the impersonator. This attack is illustrated in Figure 4.

Our formal modeling verified the existence of this attack which breaks the authenticity and confidentiality invariants, but not the intention preservation invariant. Thus the intruder pairs with the honest devices and knows the link key of the initiator and possible responder. The intention preservation invariant does not fail because the initiator intended to pair with the intruder device since the intruder was impersonating an honest responder.

6.3 Numeric Comparison Attack II: The Proxy

In the NC association model, we already know a MitM attack exists. Another novel attack is discovered by our Murphi model that relies on a compromise of the communication channel between two honest paired devices. The attack is diagrammed in Figure 5, and we like to call a “Proxy” MitM attack. The attack consists of an intruder that simply forwards messages between an honest initiator and responder without modifications, thereby acting as a proxy. Once the pairing is made, although the Proxy intruder cannot eavesdrop any of the messages, it can perform a DoS attack by simply not forward messages along.

The intruder sets up a communication channel such that the honest responder sends its messages destined for the intruder device rather than the initiator device. In the beginning the initiator sends its public key destined for Device B. However, the intruder intercepts this message and drops the packet from the network. Although the Bluetooth protocol is a broadcast protocol we are assuming a strong adversary model and therefore we assume that the intruder has some power over the network to drop packets or at least disrupt the communication such that the Device B would not interpret the message sent from Device A. After intercepting the message from Device A, the intruder changes the source of the packet and forwards it with no other modifications to Device B. Device B then responds with its public key but with the message destined for the intruder. The intruder in turn forwards this message to Device A and the protocol follows as standard but with the intruder making sure that Device B believes that honest initiator’s address is the intruder’s address.

Our formal modeling discovered the existence of this attack which only breaks authenticity for the responder. Confidentiality and intention preservation invariants are held because the pairing still occurs between Device A and Device B, and the link key is secret since the intruder has not compromised the public key exchange. Authenticity fails because Device B is actually communicating with the intruder rather than Device A, this means that Device B requires the intruder in order to complete the communication with Device A. At this point the intruder controls Device B’s communication chan-

nels. Although this is not as critical as the other verified attacks, it is significant because the Bluetooth SSP protocol should ensure that not only are devices paired with the devices they originally intended to be pair with, but the initiator and responder should also be communicating directly without a MitM attack such as this one.

7. DISCUSSION AND FIXES

We have discussed the discovered attacks found by our Murphi model. In the following sections, we discuss their implications and security risks as well as potential recommendations for fixes to better secure the Bluetooth 4.0 Pairing Protocol.

7.1 Just Works Attack

The JW association model is known to not provide protection against MitM attacks and is only present to allow for convenience for devices such as mice, keyboard, or headsets. For other devices, the NC association model is recommended for pairing, such as for smartphones or laptops to help mitigate MitM attacks. Also for the general usages of Bluetooth, the range of the protocol is between 5-30 meters and therefore an adversary needs to be fairly close physically in order to mount the attack. In this case, the attack could be impractical.

7.2 Numeric Comparison Attack I: The Impersonator

The NC association model is presented by the Core Bluetooth Specifications [12] to protect against MitM attacks with high probability. However, our model verifies the previously discovered attack by Phan and Mingard [15] of the Impersonator MitM attack. This attack is a very critical compromise of Bluetooth SSP because the NC association model is presumed to protect against MitM attacks.

Such an attack is difficult to fix, especially for a protocol that attempts to be extremely energy efficient and lightweight. One way to protect against such MitM attacks is to establish a trusted authority similar to the Certificate Authorities that are used on the Internet. Even then impersonation is difficult to defend against because there are many attack vectors and self signing that can exist such as impersonating trusted entities in the verification process. Also, the Bluetooth device would then need to shave an eternal communication channel in order to properly verify, which could significantly decrease its efficiency and increase energy usage. Furthermore, having a centralized trust authority would mean that device must be connected to the Internet in order to be paired, which in many cases is impractical and possibly defeat the purpose of the ad-hoc feature of Bluetooth.

These attacks would be mostly successful in situations when a user must pair with a device that he or she does not own. Because Bluetooth does not have centralized trust such as web Certificate Authorities, SSP cannot guarantee trust of devices. Most users and Bluetooth vendors do not consider this because they expect most usage to be between devices actually owned by the user. However, as we have demonstrated in the car rental example, such assumptions about trust are not always true.

7.3 Numeric Comparison Attack II: The Proxy

Although the Impersonator attack is more critical than the Proxy attack in the NC model, the Proxy MitM is still significant since an honest device believes its communication is being sent to another honest device through a trusted network when in fact it is not. Although the communication will be encrypted in this scenario, the responder device relies on the intruder device to communicate properly with the other device. Giving an intruder such control can lead to denial of service attacks and allow them an easier interface to possible eavesdropping on the traffic between the two users. Similar to the other attacks on the protocol, the intruder must have their device within range for this to be plausible.

Like the NC Impersonator MitM attack, the NC Proxy MitM attack is difficult to fix because of the inherent mutability of the underlying networking paradigm. To help us find a fix for this adversary, we need to create a way to make Bluetooth addresses immutable. If we can make these addresses immutable, then an adversary would not be able to spoof its address to trick an honest initiator or responder to complete the pairing through it. However, it isn't always clear exactly how this can be done, and it also depends on the underlying hardware and radio frequency generations. Finding a fix for this attack could be possible lead-ins for future work.

8. CONCLUSIONS

Since its conception in 1994, Bluetooth has quickly risen in popularity and its use and applications have become widespread. As a consequence, its ubiquitous nature has made Bluetooth's SSP security an important problem.

In our paper, using the Dolev-Yao [4] adversarial model, we presented and discussed a formal model using Murphi to exhaustively verify the security of Bluetooth's SSP. In our findings, we verify two previously known attacks on Bluetooth [15]. In the JW association model, a MitM attack is found where an attacker can easily pair with both initiator and responder, thereby allowing it to eavesdrop and modify any messages sent across the paired devices. Bluetooth designed the NC association model specifically to combat this MitM attack. However, as we have seen, their efforts fall short of preventing all MitM attacks from Bluetooth SSP. In the NC association model, we verified the Impersonator MitM attack informally reasoned by Phan and Mingard [15].

In addition, we have found a novel attack on Bluetooth's supposedly secure NC association model: the Proxy MitM attack. In this attack, an adversary allows an initiator to pair with a responder by intercepting the key exchange messages, but forwarding these messages along without modification, except for the source. As a result, at the completion of the pairing process, although the intruder does not have the link key, it has successfully forced the initiator and responder to communicate through the intruder as a proxy. In this way, the intruder compromises the authenticity of the Bluetooth SSP security. Because the adversary now controls all communication between the initiator and responder, successful communication between the initiator and responder depends solely on the intruder's likelihood to proxy messages between them. Thus, the intruder can easily perform

a denial of service attack.

Through our formal analysis, we hope to shed light on Bluetooth SSP and to further advance the efforts made to ensure its security.

9. REFERENCES

- [1] J. Bau and J. Mitchell. Security modeling and analysis. *Security Privacy, IEEE*, 9(3):18–25, May-June.
- [2] R. Chang and V. Shmatikov. Formal analysis of authentication in bluetooth device pairing. In *Proceedings of LICS/ICALP workshop on foundations of computer security and automated reasoning for security protocol analysis (FCS-ARSPA â07)*, July.
- [3] W. Diffie, P. C. Van Oorschot, and M. J. Wiener. Authentication and authenticated key exchanges. *Des. Codes Cryptography*, 2(2):107–125, June 1992.
- [4] D. Dolev and A. C. Yao. On the security of public key protocols. *Information Theory, IEEE Transactions*, 29(2):198–208, March.
- [5] J. P. Dunning. Taming the blue beast: A survey of bluetooth- based threats. *Security & Privacy, IEEE*, 2010.
- [6] K. Haataja and P. Toivanen. Practical man-in-the-middle attacks against bluetooth secure simple pairing. In *Wireless Communications, Networking and Mobile Computing, 2008. WiCOM 08. 4th International Conference on*, pages 1–5, October 2008.
- [7] S. Just, M. & Vaudenay. Authenticated multi-party key agreement. In *In Advances in Cryptology, Asiacrypt 96, LNCS 1163*, pages 36–49, 1996.
- [8] C. Kuo, J. Walker, and A. Perrig. Low-cost manufacturing, usability, and security: An analysis of bluetooth simple pairing and wi-fi protected setup. In *Proceedings of International Conference on Usable Security (USEC 07)*, pages 325–340, 2007.
- [9] A. Lindell. Attacks on password pairing in bluetooth v2.1. In *CSI â08, Maryland*, 2008.
- [10] A. Lindell. Bluetooth v2.1â a new security infrastructure and new vulnerabilities. In *BlackHat Briefings, Las Vegas.*, 2008.
- [11] Bluetooth, SIG. Bluetooth core specification v1.0 - 2.1+edr. <https://www.bluetooth.org/Technical/Specifications/adopted.htm>, 1997-2007.
- [12] Bluetooth, SIG. Bluetooth core specification v4.0.30. <https://www.bluetooth.org/Technical/Specifications/adopted.htm>, June 2010.
- [13] A. J. Menezes, P. C. Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. BocaRaton: CRC Press, 1997.
- [14] P. Mingard. Elliptic curve based diffie-hellman in bluetooth v2.1. Master of computer science semester project thesis, EPFL, Switzerland, 2007.
- [15] R.-W. Phan and P. Mingard. Analyzing the secure simple pairing in bluetooth v4.0. *Wireless Personal Communications*, 64:719–737, 2012.
- [16] S. Sandhya and K. Devi. Analysis of bluetooth threats and v4.0 security features. In *Computing, Communication and Applications (ICCCA), 2012 International Conference*, pages 1–4, Feb.
- [17] M. Shrivastava. Analysis of security risks in bluetooth. In *IJCAR*, volume 1, pages 88–95, December.

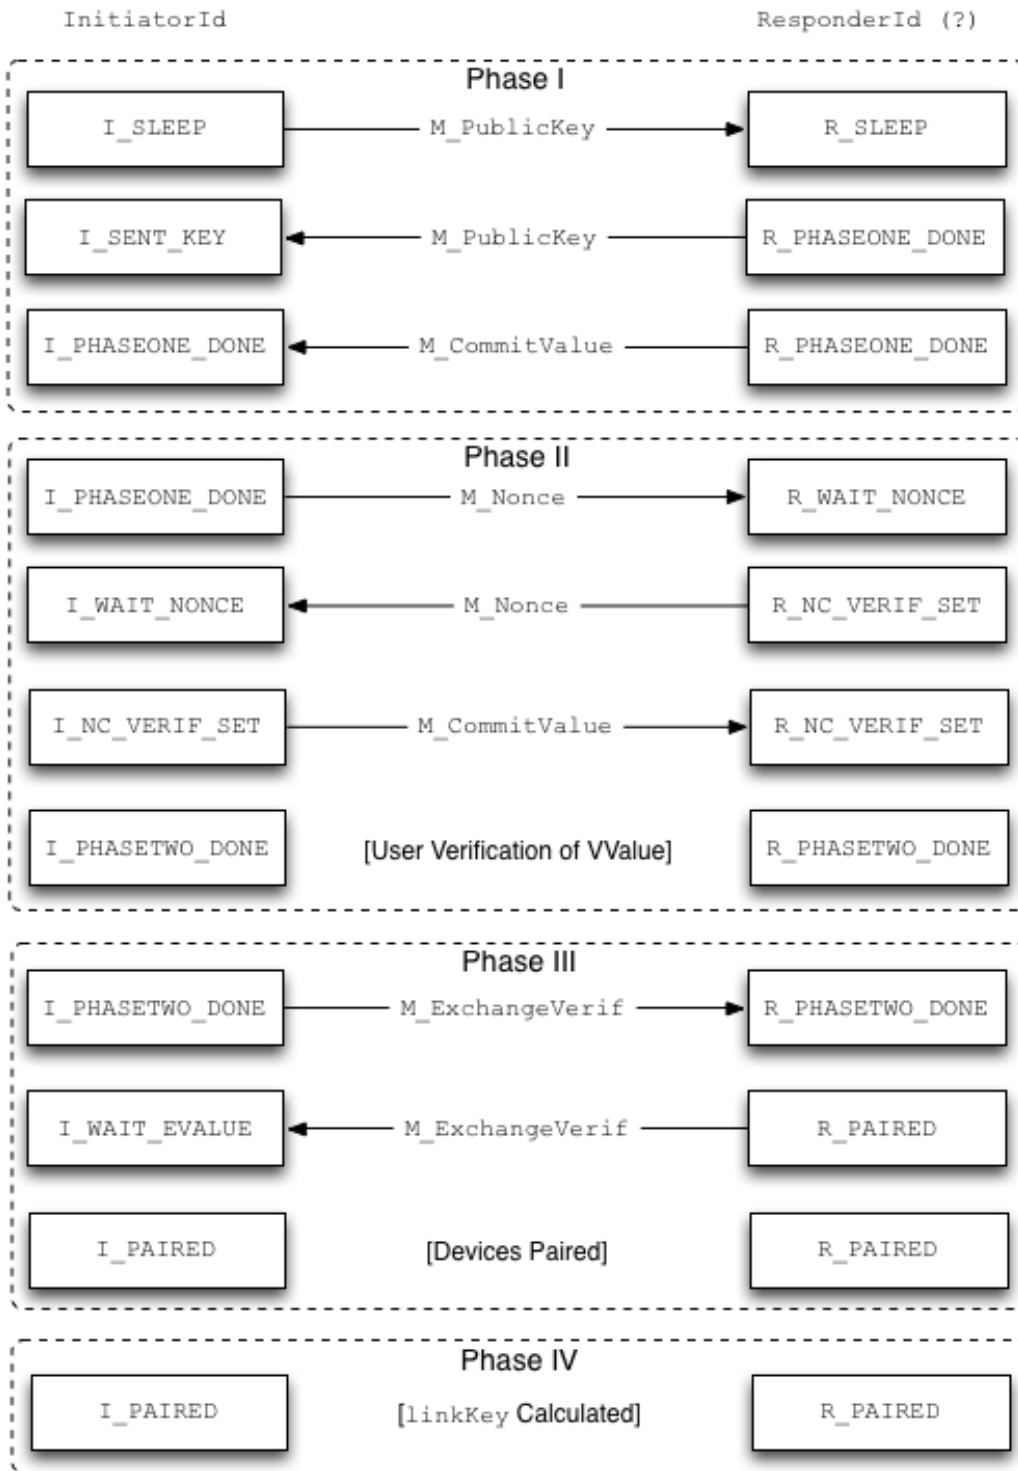


Figure 1: A diagram describing the transitions of the initiator and responder, and the messages that are passed across during SSP. Created by the author.

	JW	JW*	NC	NC*	NC*	NC*
NumInitiators	1	1	1	1	2	3
NumResponders	1	1	1	1	1	1
NumIntruders	1	1	1	1	1	1
NetworkSize	10	10	10	10	10	10
MaxKnowledge	10	10	10	10	10	10
MaxPairings	10	10	10	10	10	10
States	3,408	1,705	3,149	18,326	182,350	N/A
Time	1.3 sec	1.25 sec	1.28 sec	2.06 sec	10.54 sec	N/A
Attack	✓	✓	✓	✓	✓	N/A

Figure 2: Table of Murphi state exploration results when searching for attacks. Created by the author. * denote the models that are tested when the initiator cannot initiate a connection with an intruder device.

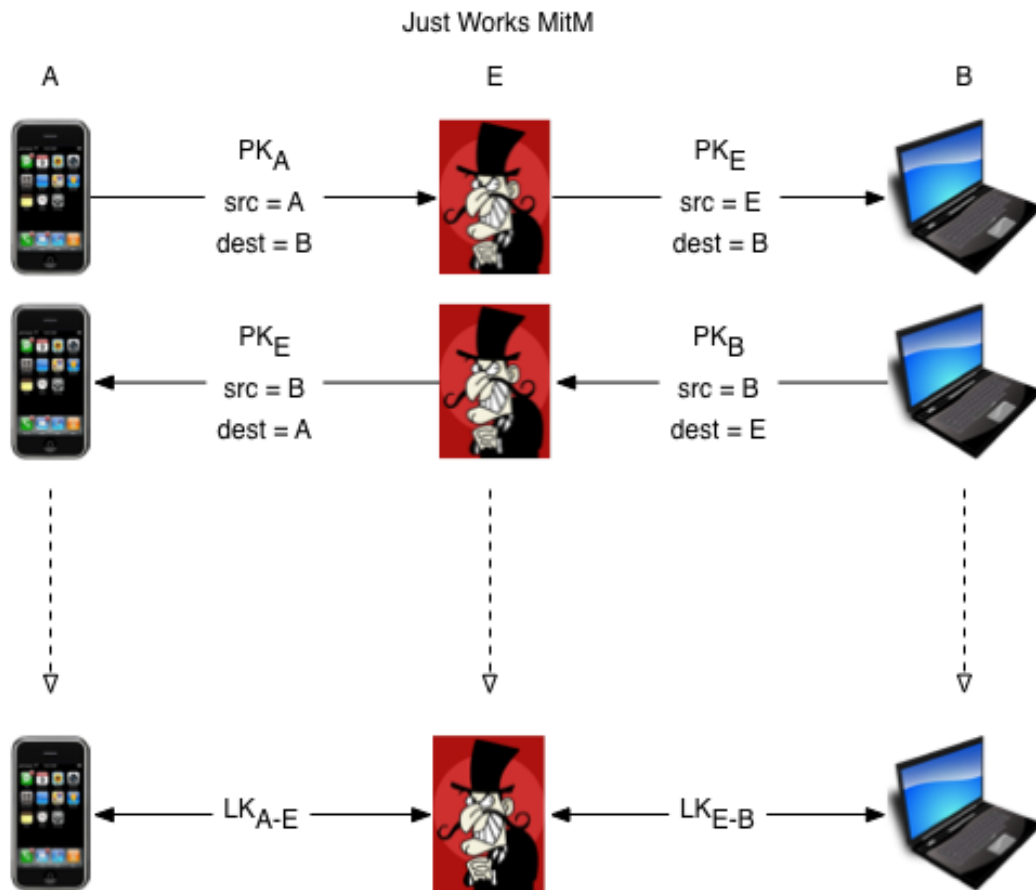


Figure 3: A diagram describing the MitM attack for the Just Works association model. Created by the author.

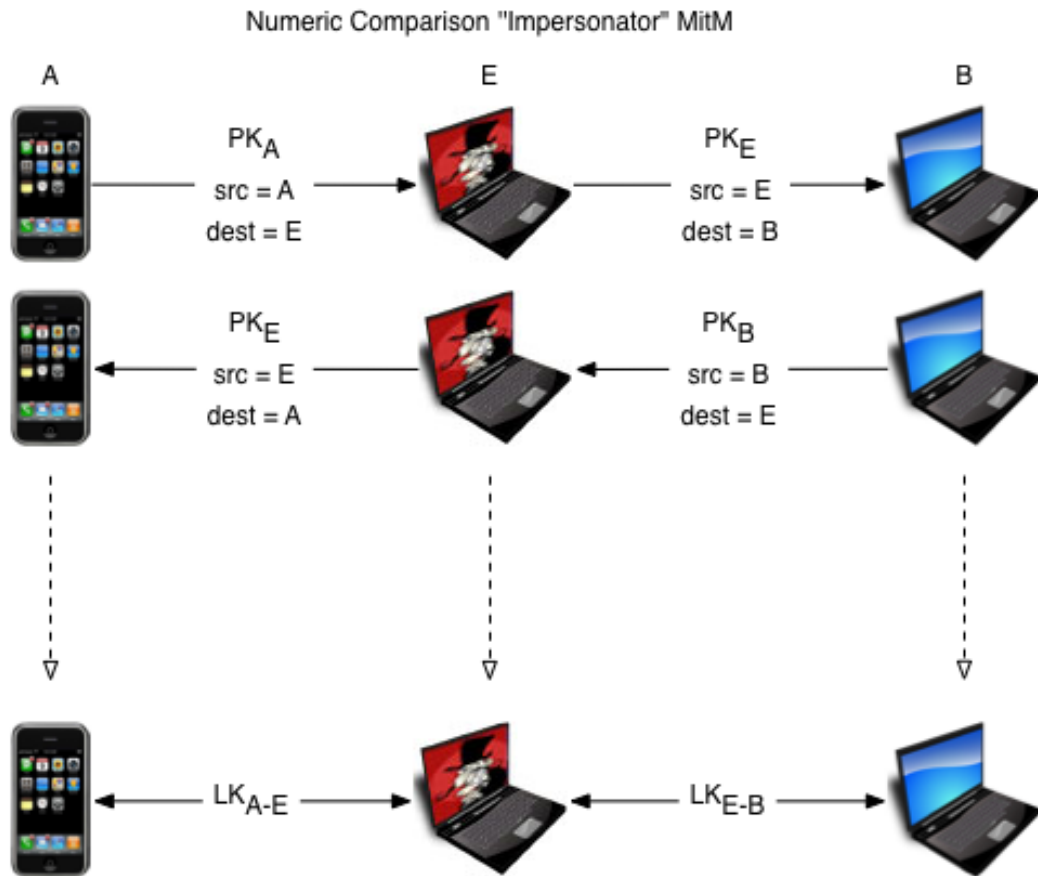


Figure 4: A diagram describing the "Impersonator" MitM attack for the Numeric Comparison association model. Created by the author.

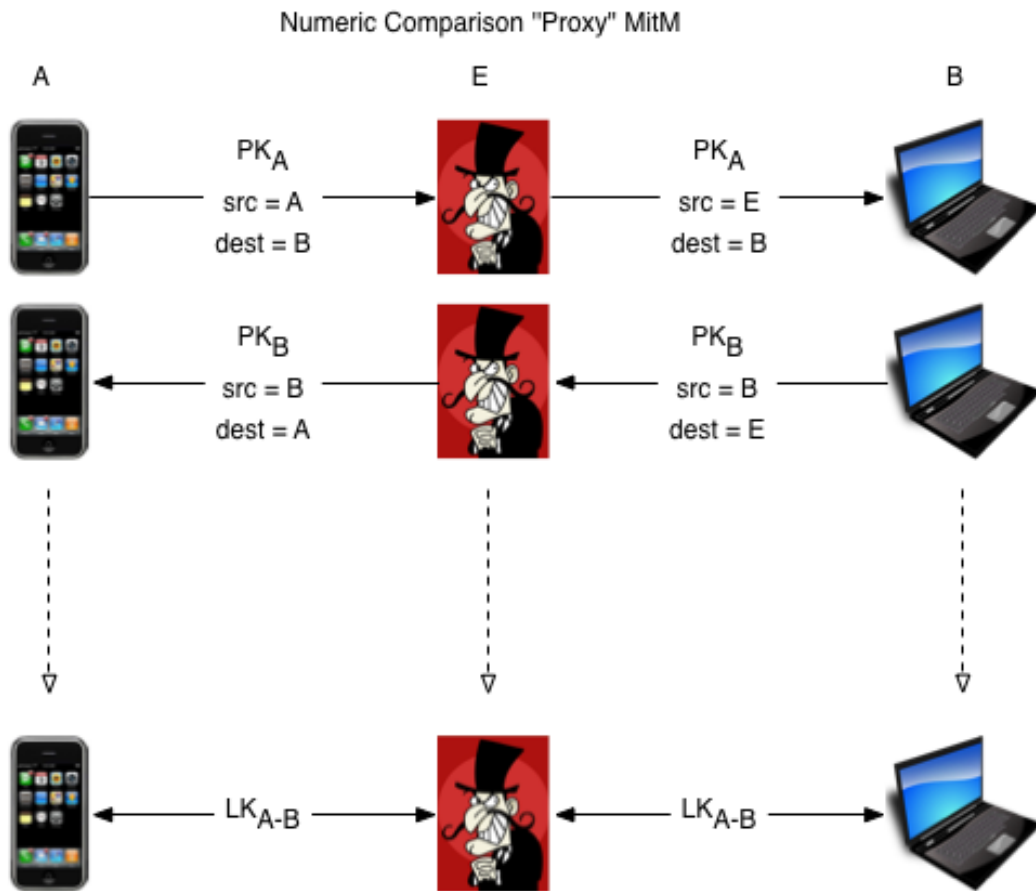


Figure 5: A diagram describing the "Proxy" MitM attack for the Numeric Comparison association model. Created by the author.