

Secure Messaging Protocols

Sebastian R. Verschoor

Offensive Technologies – Guest Lecture

10 April 2025



Disclaimer

These slides are optimized for didactic purposes. Primitives and protocols have been simplified, sometimes to the point where technically they are incorrect (and most likely insecure).

2 / 56

Outline

- Secure Messaging protocols
 - History
 - Secure messaging features
- Preliminaries
 - Attacker model
 - Symmetric Cryptography
 - Public Key Cryptography
- PGP
- OTR
- SCIMP
- Signal protocol
 - Signal application

3 / 56

History (1/3)

- ▶ 1991: Phil Zimmermann creates Pretty Good Privacy (PGP)
 - ▶ February 1993: US starts criminal investigation for “munitions export without a license”
 - ▶ 1995: PGP source code published as a physical book
 - ▶ US first amendment protects export of books
 - ▶ 1996: Criminal investigation was dropped, no charges were filed
- ▶ 2004: Nikita Borisov, Ian Goldberg and Eric Brewer create OTR
 - ▶ “Off-the-Record Communication, or, Why Not To Use PGP”
 - ▶ Has forward secrecy
 - ▶ Has deniability
 - ▶ Requires both parties online for setting up

4 / 56

History (2/3)

- ▶ 2011: Gary Belvin introduces SecureSMS; “OTR for SMS”
- ▶ 2012: SCIMP (Silent Circle instant messaging protocol)
 - ▶ By Vinnie Moscaritolo, Gary Belvin and Phil Zimmermann
 - ▶ SecureSMS for XMPP
 - ▶ I formally verified its security with ProVerif
- ▶ February 2014: Open Whisper Systems releases TextSecure v2
 - ▶ Asynchronous: allows offline initial user message
 - ▶ Later renamed to Signal
- ▶ May 2014: SC updates to SCIMP v2
 - ▶ Asynchronous: allows offline initial user message
- ▶ August 2015: SC releases code for SCIMP v2
 - ▶ Adds more inconsistencies between code and documentation
 - ▶ I find and report many security bugs in the code
- ▶ September 2015: SC discontinues SCIMP, switches to Signal based protocol

5 / 56

History (3/3)

- ▶ February 2014: Facebook (now Meta) acquires WhatsApp
- ▶ April 2014: Signal announces partnership with WhatsApp
- ▶ April 2016: WhatsApp completes integration of Signal protocol
- ▶ November 2016: Trevor Perrin and Moxie Marlinspike release official specification for the Signal protocol
- ▶ September 2023: Signal gets post-quantum confidentiality

6 / 56

Security and Privacy features

- ▶ Confidentiality
- ▶ Integrity
- ▶ Availability
- ▶ (Key) Authentication
- ▶ Forward Secrecy
- ▶ Post-Compromise Security (PCS)
- ▶ Deniability vs. Non-repudiation
- ▶ Transport Privacy

7 / 56

Trust Establishment


- ▶ Opportunistic encryption
- ▶ Public Key Infrastructure (PKI)
- ▶ Web-of-Trust (WoT)
- ▶ Trust-On-First-Use (TOFU)
- ▶ Fingerprint verification
 - ▶ Socialist Millionaire Protocol (SMP)
 - ▶ Short Authentication String (SAS)
 - ▶ Safety Numbers
- ▶ Key directory
- ▶ Key transparency
- ▶ Blockchain(?)

8 / 56

- ▶ User experience
- ▶ Multi-device
- ▶ Group chat
- ▶ File transfer
- ▶ Video-chat
- ▶ Backups
- ▶ (Formal) verification
- ▶ Implementation security
- ▶ Audits


- ▶ Briar
 - ▶ Discord (Dave)
 - ▶ Dust
 - ▶ Facebook Messenger
 - ▶ Google Allo
 - ▶ Google Chat
 - ▶ Google Messages
 - ▶ iMessage
 - ▶ irc
 - ▶ LINE
 - ▶ Matrix (Olm/Megolm)
 - ▶ Mattermost
 - ▶ Pond
 - ▶ QQ Mobile
 - ▶ Rocket.Chat
 - ▶ Session
 - ▶ SimpleX
 - ▶ Skype
 - ▶ Slack
 - ▶ SnapChat
 - ▶ Teams
 - ▶ Telegram (MTPROTO)
 - ▶ Threema
 - ▶ Viber
 - ▶ WeChat
 - ▶ WhatsApp
 - ▶ Wickr
 - ▶ Wire
 - ▶ X
 - ▶ XMPP (OMEMO)
 - ▶ Zoom
 - ▶ Zulip
 - ▶ ...
- many use the Signal protocol or a variant

End-to-End security (E2E, sometimes E2EE for encryption)
 ▶ all messages are handed to the adversary for delivery.¹




- Mallory has full control over all messages
- ▶ she may learn, change, inject, drop, reorder, and replay all messages
- Kerckhoffs principle
- ▶ Mallory knows everything except the key

¹For simplicity, I will omit Mallory from most diagrams for now.



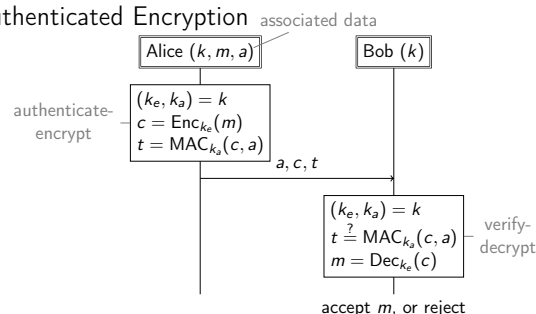
Provides confidentiality: Mallory learns nothing² about m

²almost nothing: I will not formalize this today



- Provides integrity: Mallory cannot³ change m
- Provides authenticity: Bob knows m was sent by Alice
- ▶ *implicit*: Bob assumes only Alice knows k

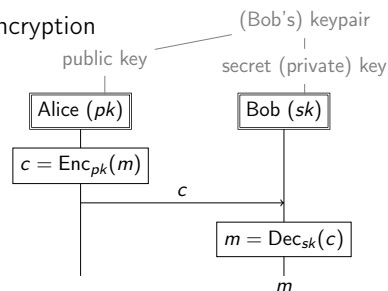
³almost cannot: again we omit the details



- ▶ Real-world AE often differs from encrypt-then-MAC
- ▶ Simplified notation:
 - ▶ $c = \text{AEAD}_k(m, a)$ c includes the tag
 - ▶ $c = \text{AEnc}_k(m)$ no a
 - ▶ $m = \text{VDec}_k(c, a)$ $m = \perp$ on rejection

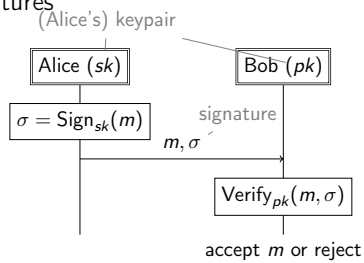
Strong assumption:

- ▶ Alice and Bob have *the same* secret key k



- ▶ Provides confidentiality
- ▶ Bob publishes pk , keeps sk secret
 - ▶ Anyone can encrypt, but only Bob can decrypt


Digital Signatures



- ▶ Provides integrity and authentication
- ▶ Alice publishes pk , keeps sk secret
 - ▶ Only Alice can sign, but anyone can verify
 - ▶ Provides non-repudiation: (m, σ) is proof for anyone that Alice sent m

17 / 56


Cryptographic Hash



- ▶ Let H be a hash function: $out = H(in)$
 - ▶ large input
 - ▶ small output (digest)
- ▶ Security
 - ▶ the output "seems random" different ways to formalize this
 - ▶ Mallory cannot compute out without knowing in
 - ▶ Mallory cannot compute in when given out
- ▶ Note there is no key involved
 - ▶ Mallory can compute H on inputs of her choice

18 / 56


Digital signatures with hashes



19 / 56

Public Key Cryptography

- ▶ No large messages
 - ▶ PKE has limited size of messages
 - ▶ PKE is relatively slow
 - ▶ How does Alice know that pk belongs to Bob?
- Mallory-in-the-Middle (MitM) attack:



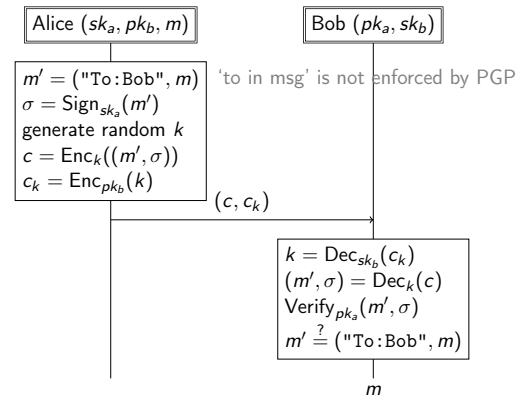
20 / 56

PKC: key/user authentication

- ▶ Not just an issue for secure messaging protocols
- ▶ SSH uses Trust On First Use (TOFU)
 - ▶ asks user to verify public key fingerprint (its hash) on first login
 - ▶ how to verify this, the protocol does not say
 - ▶ once accepted, it will silently keep accepting until the key changes
- ▶ TLS uses certificates and a PKI
 - ▶ you connect to a server
 - ▶ server presents a certificate: "this public key belongs to this website"
 - ▶ the certificate is signed by an authority
 - ▶ you as a user trust the authority (right?)
- ▶ PGP, OTR, SCIMP and Signal all use different methods

21 / 56

Pretty Good Privacy



22 / 56

PGP: key authentication

- ▶ PGP uses the Web of Trust (WoT)
- ▶ You meet in person, then sign each others key
 - ▶ yes there are (were?) key signing parties!
- ▶ Everyone publishes the signed keys
- ▶ If you get a key you don't know, you check if it's signed by someone you trust

Everyone loves this system and it scales great in practice!

23 / 56

PGP

- ▶ PGP uses hybrid encryption
 - ▶ faster + larger messages
- ▶ PGP has many options
 - ▶ I am 80% sure the above is the sign-and-encrypt option
 - ▶ Without recipient ID, Bob could re-encrypt to others
 - ▶ not part of the PGP specification
 - ▶ Complexity leads to bad user experience, which leads to loss of security
- ▶ Two major issues
 - ▶ If a private key leaks, all messages leak (past and future)
 - ▶ Bob can publish (m, σ) as evidence that Alice said m

24 / 56

Off-the-Record (OTR)




Design philosophy: A secure online conversation should be more like a private in-person conversation

- ▶ Forward secrecy
 - ▶ leaking long-term keys should not reveal information about old messages
 - ▶ "key erasure property"
- ▶ Deniability (informal)
 - ▶ leaking a secure online conversation should not leak any more information than leaking a plain-text conversation would
 - ▶ has many subtly different *mathematical* formalizations
 - ▶ unclear if this affects *legal* deniability

25 / 56

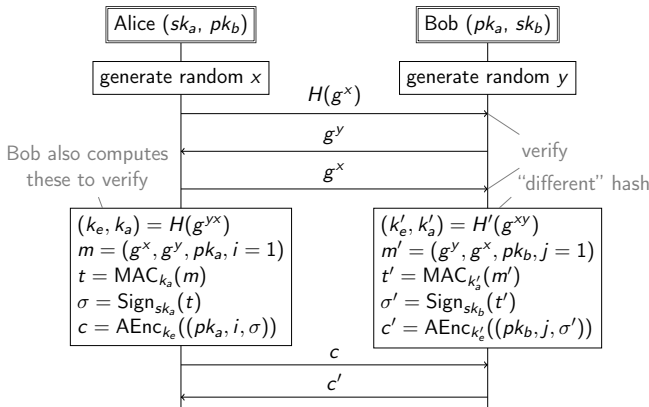
Diffie-Hellman Key Exchange



- ▶ correct: $g^{xy} = g^{yx}$
- ▶ security: follows from the Diffie-Hellman assumption:
 - ▶ Given g^x and g^y , it's hard to compute g^{xy}

26 / 56

OTR: Authenticated Key Exchange



27 / 56

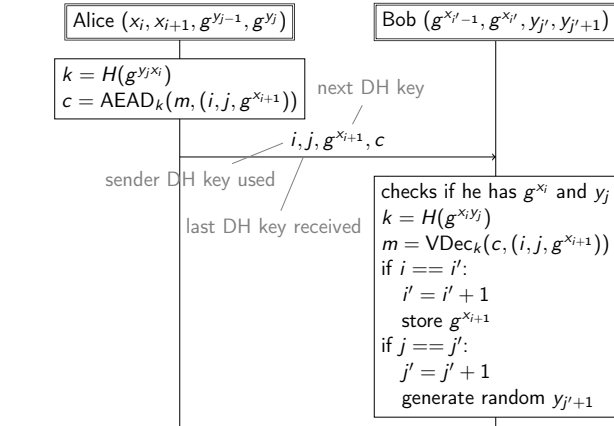
OTR



- ▶ Deniability:
 - ▶ Parties only ever sign public values (g^x, g^y, pk)
 - ▶ No proof of conversation contents
 - ▶ Parties only ever sign *their own* public key
 - ▶ No proof of intent to communicate with other party
- ▶ Forward secrecy:
 - ▶ Securely delete x once we are done with it
 - ▶ (x, g^x) is called an ephemeral keypair
 - ▶ No information left on the device to recompute k_e
 - ▶ Discard k_e once we are done with it
 - ▶ ... but when is that?

28 / 56

OTR: Sending Data Message



29 / 56

OTR



- ▶ Forward Secrecy:
 - ▶ with every *reply* we remove an old key
 - ▶ old keys cannot be derived from previous ones
 - ▶ one-sided conversations don't move forward
 - ▶ could be fixed with heartbeat messages
- ▶ Post Compromise Security:
 - ▶ if Mallory steals your keys she can read your messages
 - ▶ once you generated a new DH key, she no longer has access
 - ▶ (this is not true if she actively maintains a MitM attack)
- ▶ Can handle missing a message, however
 - ▶ cannot handle out of order messages
 - ▶ storing old keys would compromise FS

30 / 56

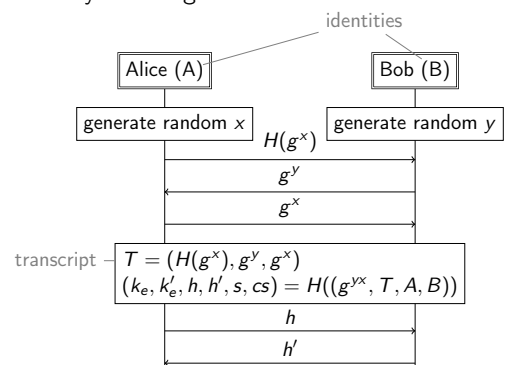
OTR: key authentication



- Option 1:
- ▶ Users can see the used public key fingerprints
 - ▶ Verify these out-of-band
- Option 2:
- ▶ Users are assumed to share some secret that Mallory doesn't know
 - ▶ Hashes used public keys and the secret together
 - ▶ Compare if they are the same using a zero-knowledge protocol
 - ▶ this all happens in-band!
- Usability studies show issues with both

31 / 56

SCIMP: Key Exchange



32 / 56

SCIMP: (key) authentication



- ▶ There are no public keys, so key authentication is impossible
- ▶ In fact, the key exchange is **unauthenticated**
- ▶ To authenticate the (already established) session:
 - ▶ s is the "short authentication string"
 - ▶ Alice and Bob **must** compare s out-of-band
- ▶ But it also means we have good deniability

33 / 56

SCIMP Symmetric Ratchet

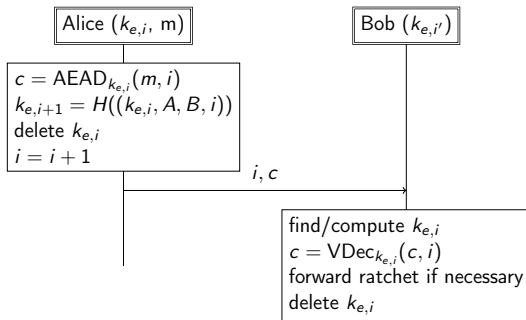


$$k_{e,1} \xrightarrow{H} k_{e,2} \xrightarrow{H} k_{e,3} \xrightarrow{H} k_{e,4} \xrightarrow{H} k_{e,5}$$

- ▶ symmetric key ratchet: $k_{e,i+1} = H(k_{e,i})$
- ▶ Send index i alongside ciphertext
- ▶ For example, Bob
 - ▶ has $k_{e,1}$, gets $i = 1$ (in order)
 - ▶ decrypts with $k_{e,1}$
 - ▶ computes $k_{e,2} = H(k_{e,1})$
 - ▶ deletes $k_{e,1}$
 - ▶ gets $i = 4$ (messages skipped)
 - ▶ computes $k_{e,5} = H(H(H(k_{e,2})))$
 - ▶ decrypts with $k_{e,4}$, then deletes $k_{e,4}$
 - ▶ stores $k_{e,2}$, $k_{e,3}$, and $k_{e,5}$
 - ▶ gets $i = 3$ (out-of-order message)
 - ▶ decrypts with $k_{e,3}$
 - ▶ deletes $k_{e,3}$
 - ▶ $k_{e,2}$ threatens forward secrecy of msgs ≥ 2

34 / 56

SCIMP: Sending Data Messages



35 / 56

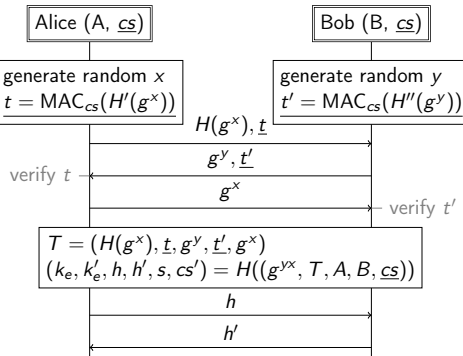
SCIMP



- ▶ Forward secrecy
 - ▶ We forward the ratchet on each *sent* message
 - ▶ But stealing old keys also leaks newer keys
 - ▶ If a key is too old (> 32 hashes old) it is removed

36 / 56

SCIMP: Rekeying



37 / 56

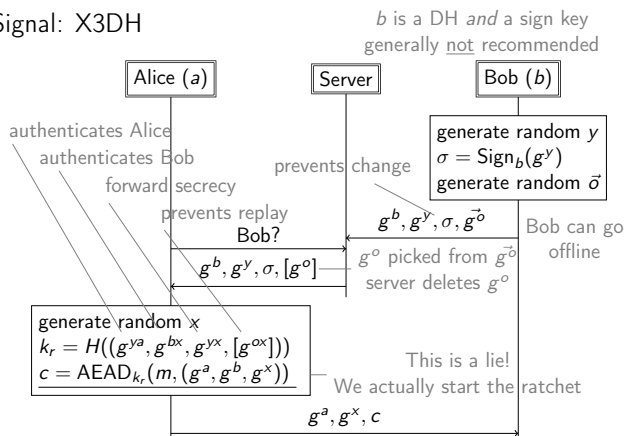
SCIMP



- ▶ Future secrecy
 - ▶ rekeying mixes in new ephemeral keys
- ▶ Unspecified when rekeying should happen
- ▶ Store oldest unused receive key
 - ▶ in case out-of-order messages arrive
 - ▶ compromise between usability and forward secrecy
- ▶ On invalid t : all state is deleted
 - ▶ no longer authenticated!
 - ▶ Mallory can easily desynchronize

38 / 56

Signal: X3DH



39 / 56

Signal: X3DH



- ▶ Asynchronous
 - ▶ Bob's handshake is independent of who wants to contact him
 - ▶ Bob does not have to be online
- ▶ All messages are delivered via the server
 - ▶ Mallory may control the server
- ▶ Alice encrypts a data message with her first message(!)


40 / 56

Signal: Double Ratchet



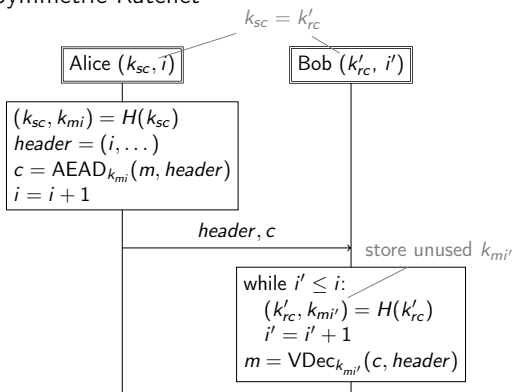
- Combine
 - the symmetrical ratchet (from SCIMP)
 - but split the chain key from the message key
 - the Diffie-Hellman ratchet (from OTR)
 - but require storage of fewer DH secrets

Signal: Symmetric Ratchet




- Parties only ever store a single k_{sc}
- $(k_{sc}, k_{mi}) = H(k_{sc})$
- If Bob has $i = 0$ and receives $i = 3$
 - he iterates H three times
 - he overwrites k_{sc} with the new value
 - he stores k_{m1}, k_{m2}
 - he uses k_{m3} to decrypt
 - old k_m does not impact other keys
 - thus it does not threaten forward secrecy of other keys

Signal: Symmetric Ratchet

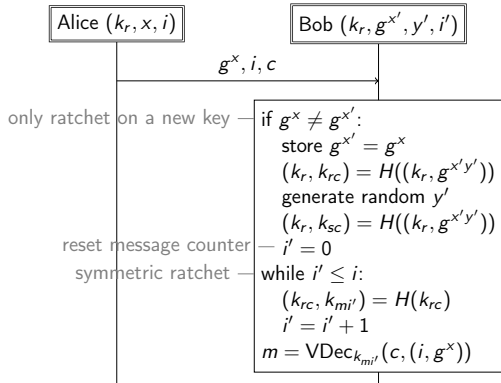


Signal: Diffie-Hellman Ratchet



- Bob has k_r, y_0
- Alice sends a new g^{x_0} :
 - Bob computes $(g^{x_0})^{y_0}$
 - Bob ratchets: next receiving chain key k_{rc}
 - Bob generates new random y_1 and computes $g^{x_0 y_1}$
 - Bob ratchets: next sending chain key k_{sc}
 - Bob sends g^{y_1} to Alice
- When Alice sends g^{x_1} , Bob ratchets again (twice)
- Old values are deleted

Signal: Diffie-Hellman ratchet

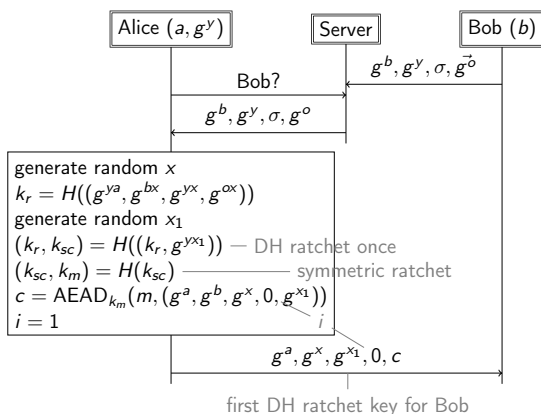


Signal: Diffie-Hellman ratchet



- Compare against OTR:
 - Bob uses g^x immediately
 - Bob verifies authenticity of g^x through associated data: $VDec_{k_{mi}}(\cdot, (i, g^x, \dots))$
 - but k_{mi} is derived from g^x (!)
 - this turns out to be secure in this context, but this is not at all obvious (to me)
 - store only one g^x and one y per peer
- What if you missed a message before ratcheting?
 - header includes a value i_p : the total number of messages sent with the previous send chain key
 - compute all missed $k_{mi'}$ before starting the DH ratchet

Signal: X3DH + ratchet start



Signal: key authentication



- Safety numbers:
 - these are hashes of the public key (+ some other values)
 - users should compare these out-of-band

Signal Protocol



What security is not provided by this protocol?


- ▶ Mallory can block all messages
- ▶ Server may send the wrong g^b :
 - ▶ MitM, if Alice and Bob don't check the safety number
- ▶ If there's no g^o in original message (server ran out or is malicious)
 - ▶ Messages can be replayed to Bob
 - ▶ Reduced forward secrecy, until Bob refreshes g^y
- ▶ Key Compromise Impersonation
- ▶ Unknown Key Share

49 / 56

Signal: Key Compromise Impersonation




- ▶ Expected: If Mallory steals b , she can impersonate Bob to others
- ▶ KCI attack: If Mallory steals y (a key from Bob), she can impersonate others to Bob



50 / 56

Signal: Unknown Key Share



- ▶ Alice thinks she is talking to Mallory
- ▶ Mallory substituted her own keys with those of Bob
- ▶ Alice is actually talking to Bob

51 / 56

Signal: multiple devices



- Setup
- ▶ Desktop displays QR
 - ▶ address
 - ▶ ephemeral public key
 - ▶ Phone scans QR, encrypts to device's ephemeral key
 - ▶ identity key pair
 - ▶ account info
 - ▶ linking token
 - ▶ Desktop registers with server as new device

Sending messages

- ▶ Encrypt the message to each device of the user
- ▶ Encrypt the message to each other device of yourself

Phishing

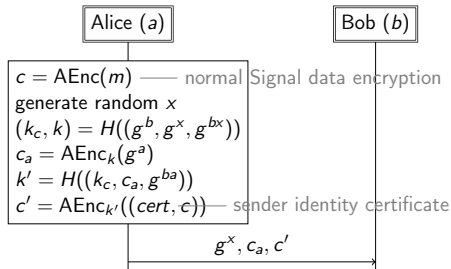
- ▶ Mallory sends device link QR, disguised as group invite QR
- ▶ user scans and doesn't read the pop-up message
- ▶ Mallory can now read along and/or impersonate

52 / 56

Signal: sealed sender



Hide sender from metadata



- ▶ Delivered to server over one-way authenticated channel
- ▶ Cautious senders should use TOR/VPN to hide their ip address

53 / 56

Signal: attachments



- ▶ generate a random key k
- ▶ encrypt the file using k
- ▶ upload the encrypted file to the file server
- ▶ send k and address to recipient *over a pairwise Signal session*
- ▶ only need to re-encrypt k and address for other recipients

54 / 56

Signal WhatsApp: groups



Each group member:

- ▶ generates a sending chain key k_{sc}
- ▶ generates an ephemeral signing key pair (sk, pk)
- ▶ sends k_{sc}, pk to each group member *over a pairwise Signal session*
- ▶ if anyone leaves the group: delete k_{sc} (and sk ?)

For each message m

- ▶ ratchet forward: $(k_{sc}, k_m) = H(k_{sc})$
- ▶ encrypt: $c = Enc_{k_m}(m)$
- ▶ authenticate: $\sigma = Sign_{sk}(c)$
- ▶ send (c, σ) to the server
- ▶ server forward (c, σ) to all group members

55 / 56

References



- ▶ 2015, Unger et al. SoK: Secure Messaging <https://doi.org/10.1109/SP.2015.22>
 - ▶ PGP: RFC 9580
 - ▶ OTR: <https://otr.cypherpunks.ca/Protocol-v3-4.0.0.html>
 - ▶ SCIMP: <https://ia.cr/2016/703>
 - ▶ Signal: <https://signal.org/docs/>
- Slides will be made available on my website zeroknowledge.me

56 / 56