Anonymity

Software Security

Winter 2023
Thompson Rivers University

Anonymity

- Anonymity: Concealing your identity
  - Anonymous communication on the Internet: The identity of the source and/or destination are concealed
- Anonymity is not confidentiality
  - Confidentiality hides the contents of the communication
  - Anonymity hides the identities of who is communicating with whom

Why Anonymity?

- Protection of privacy
- Avoiding online harassment
- Whistleblowing
- Activism work
- Accessing censored or restricted content
- Protecting against identity theft
- Avoiding discrimination
- Creative expression

Anonymity on the Internet

- Anonymity on the Internet is hard
  - Difficult, if not impossible, to achieve on your own
  - Packets contain the source IP address and destination IP address
- Anonymity is easier for attackers
  - An attacker can hack into someone else's computer and send communications from that computer
  - We assume honest users won't hack into other computers
- Main strategy for anonymity: Ask someone else to send messages for you

Proxies and VPNs

Proxies

- Alice wants to send a message to Bob
  - Bob shouldn't know the message is from Alice
  - An eavesdropper (Eve) cannot deduce that Alice is talking to Bob
- Proxy: A third party that relays our Internet traffic
  - Alice sends the message and the recipient (Bob) to the proxy, and the proxy forwards the message to Bob
  - The recipient's name (and optionally the message) is encrypted, so an eavesdropper does not see a packet with both Alice and Bob's identities in plaintext
  - Bob receives the message from the proxy, with no indication it came from Alice

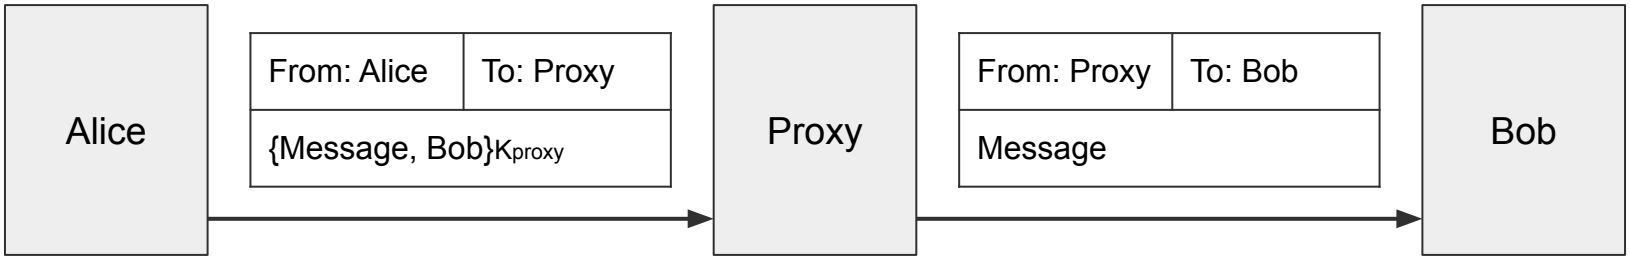
Alice
From: Alice
To: Proxy
{Message, Bob}Kproxy
Proxy
From: Proxy
To: Bob
Message
Bob

Virtual Private Networks (VPNs)

- VPNs create a secure and encrypted connection between your device and a remote network, allowing you to access resources on that network as if you were physically there.
- When you connect to a VPN, your device creates an encrypted tunnel between your device and the VPN server.
- The encryption used by VPNs helps protect your data and communications from interception and hacking.

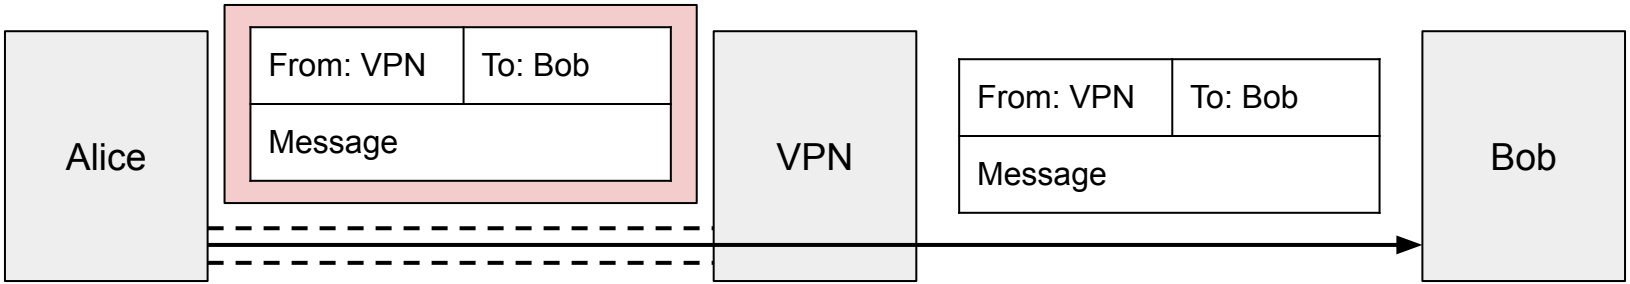
Alice
From: VPN
To: Bob
Message
VPN
From: VPN
To: Bob
Message
Bob

Proxies vs VPN

- Similar concept, but in VPN Alice directly sends packets as though coming from the VPN, wrapped in the VPN's layer of encryption
- Proxies operate at the application layer, while VPNs operate at the network layer

Proxies and VPNs: Issues

- Performance
  - Sending a packet requires additional hops across the network
- Cost
  - VPNs can cost $80 to $200 per year
- Trusting the proxy
  - The proxy can see the sender and recipient's identities
  - Attackers might convince the proxy to tell them about your identity

Real Use for VPNs

- Evading censorship
  - A local adversary can't see your traffic…
  - But the censor could just block VPN traffic instead
- Access control
  - Systems that only allow "internal" access or access from known networks…
  - The campus VPN is about solving this problem
- Separating some people from their money
  - Commercial VPN services:
    - Scare people into subscribing (and make it easy to do so)
    - Make it almost impossible to cancel!

Tor

Tor

- Idea: Send the packet through multiple proxies instead of just one proxy
- **Tor**: A network that uses multiple proxies (relays) to enable anonymous communications
  - Stands for The Onion Router
- Components of Tor
  - Tor network: A network of many Tor relays (proxies) for forwarding packets
  - Directory server: Lists all Tor relays and their public keys
  - Tor Browser: A web browser configured to connect to the Tor network (based on Firefox)
  - Tor onion services: Servers that can only be reached through the Tor network
  - Tor bridges: Tor relays that try to hide the fact that a user is connecting to the Tor network


Tor
Browser
Relay
Relay
Relay
Website

Tor Threat Model

- Security: Client anonymity and censorship resistance
  - Optional: Server anonymity with onion services
- Performance: Low(ish) latency (communication should be fast)
- Tor preserves anonymity against local adversaries
  - Example: An on-path attacker sees Alice send a message to a Tor relay, but not the final destination of the message
  - Example: The server should not know the identity of the client

Tor Circuits

- To communicate anonymously with a server, the Tor client forms a circuit consisting of 3 relays (by default)
  - Step 1: Query the directory server for a list of relays
  - Step 2: Choose 3 relays to form a Tor circuit
  - Step 3: Connect to the first relay, forming an end-to-end Tor connection
  - Step 4: Connect to the second relay through the first relay, forming an end-to-end Tor connection to the second relay
  - Step 5: Connect to the third relay through the second relay, forming an end-to-end Tor connection
  - Step 6: Connect to the web server
    - If the web server is using HTTPS, then an end-to-end TLS connection will be formed through the third relay
- Once a circuit is established messages are sent in fixed-size cells
  - Wrapped in multiple layers of encryption like an onion

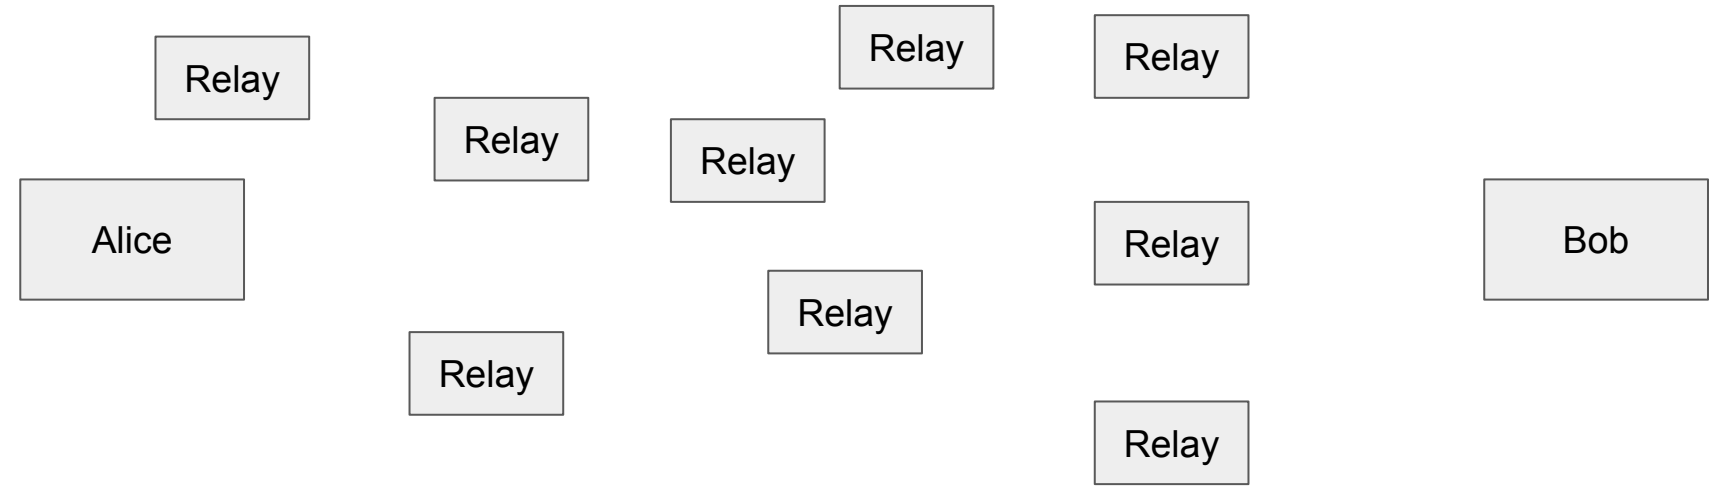

Suppose Alice wants to talk to Bob anonymously.

Alice queries the directory server and chooses 3 relays

Tor Circuits

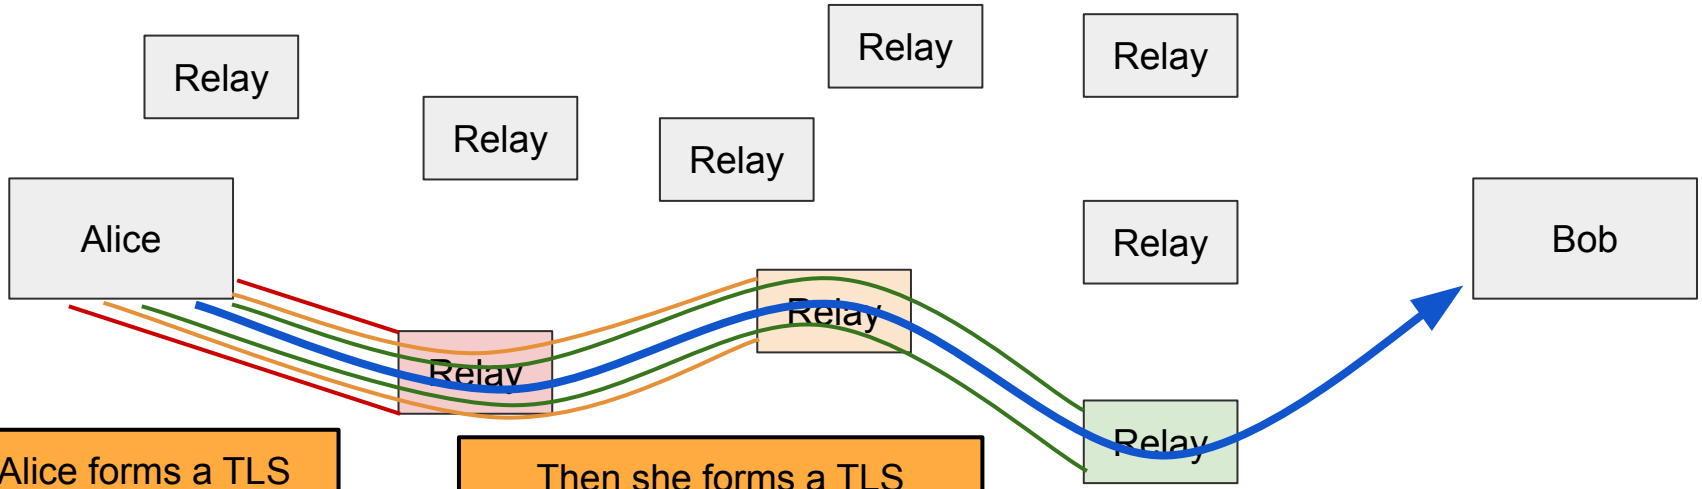

Alice forms a TLS connection with the entry node

Then she forms a TLS connection with the second node, through the first node

Notice: Relay 1 is only relaying TLS packets. It doesn't know the contents of the packets!

Then she forms a TLS connection with the exit node, through the second node

Finally, she connects to Bob (optionally forming a TLS connection with Bob)

Tor Circuits

- Function of the relays:
  - Perform Tor handshakes when requested
  - When receiving a cell, decrypt using the key obtained through the Tor protocol
  - If the destination of the cell is another relay, forward the cell to the next relay
  - If the destination of the cell is an external server, forward the cell to that server

Tor Circuits

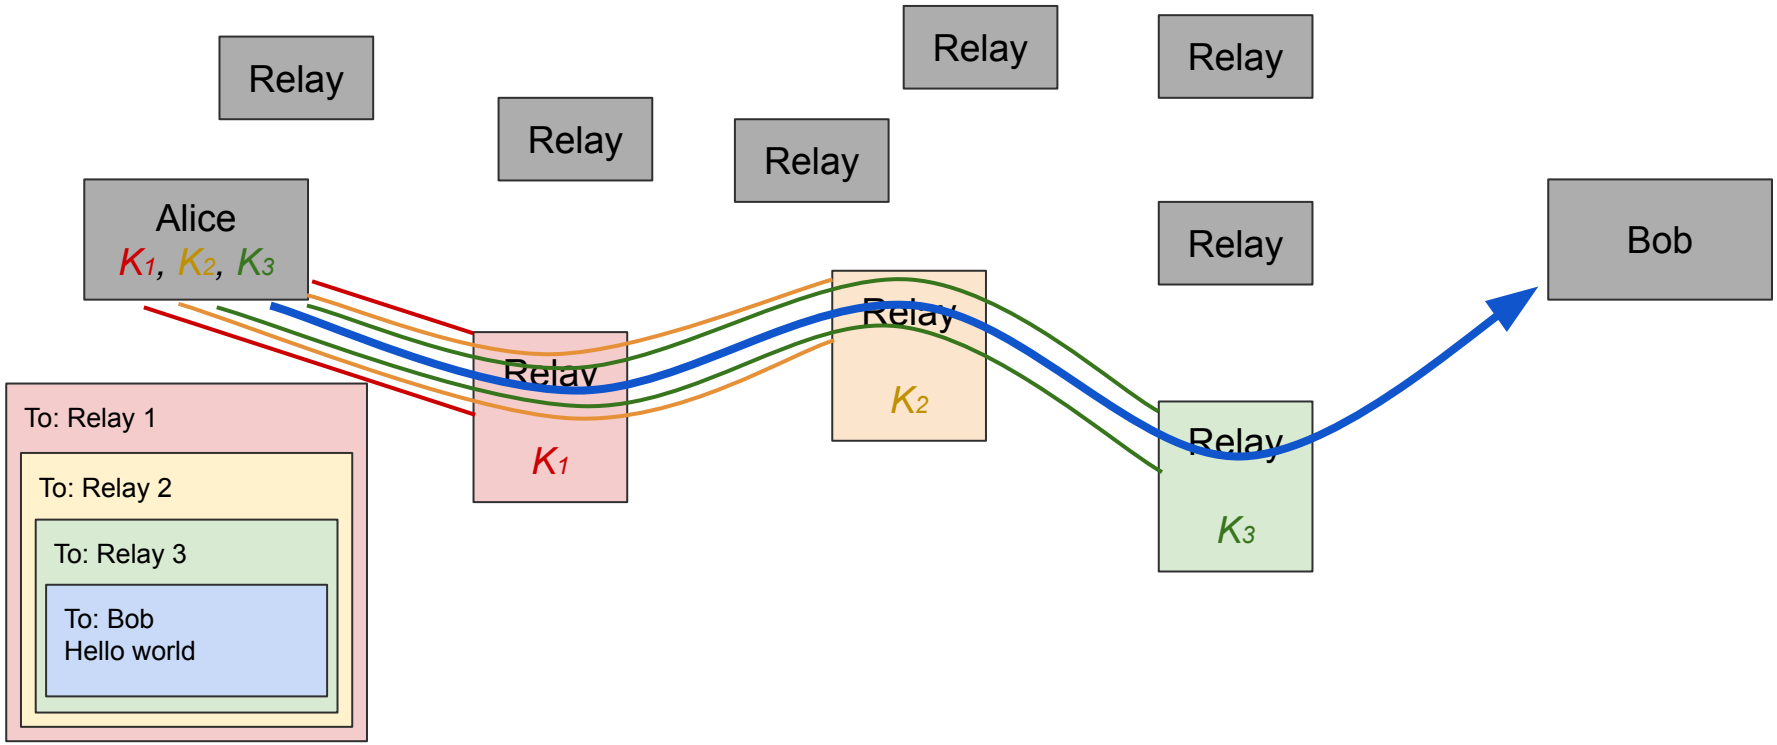

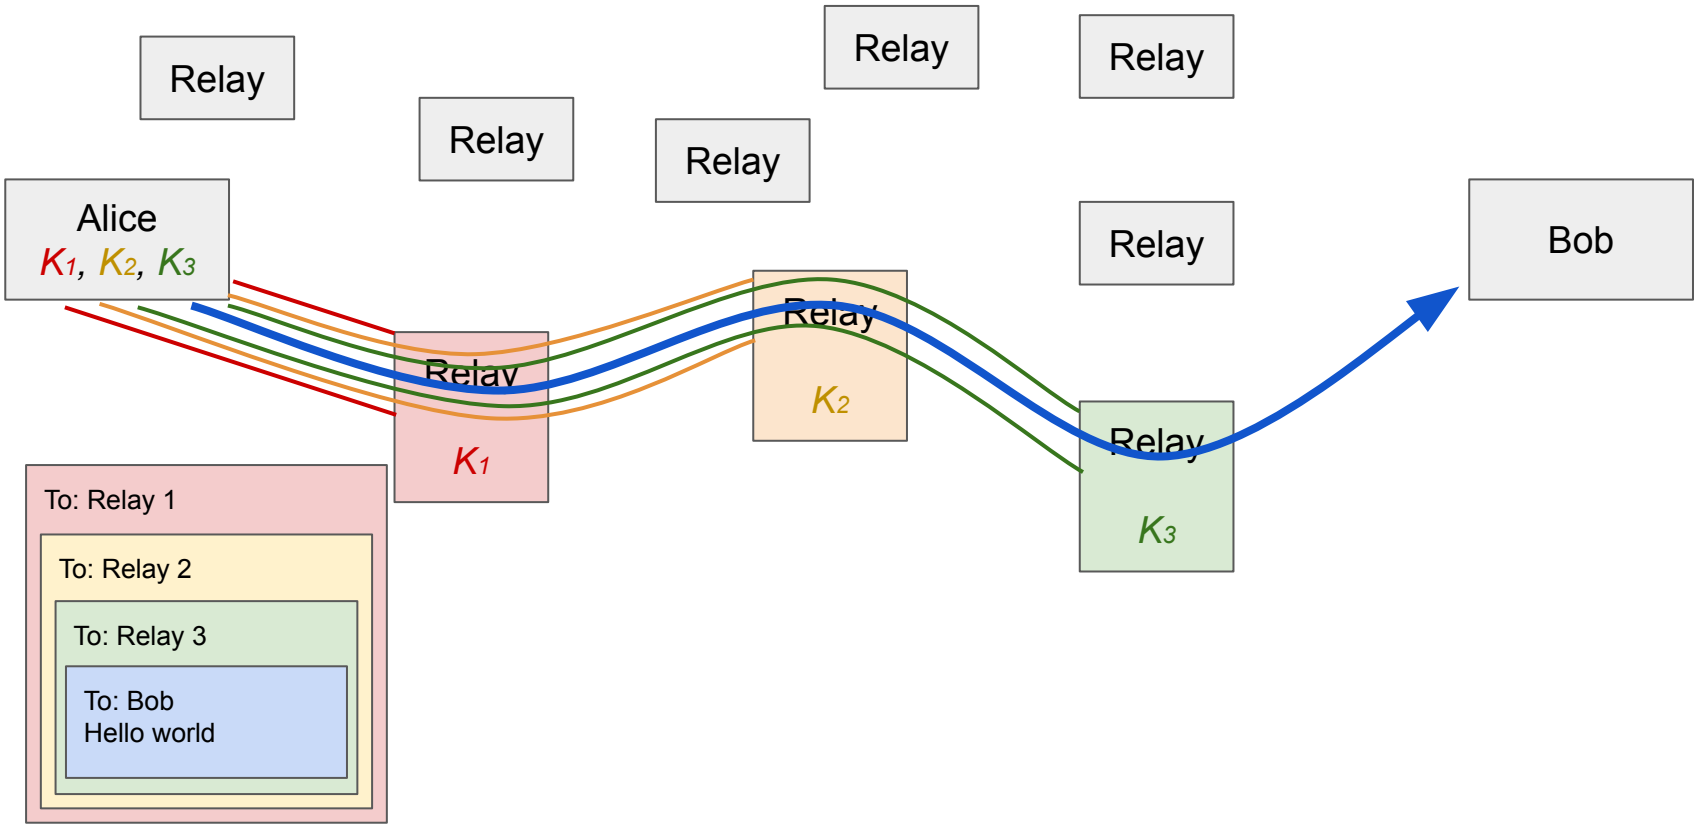
Relay
Relay
Relay
Relay
Alice
K1, K2, K3
Relay
Relay
Relay
Bob
Relay
K1
Relay
K2
Relay
K3
To: Relay 1
To: Relay 2
To: Relay 3
To: Bob
Hello world

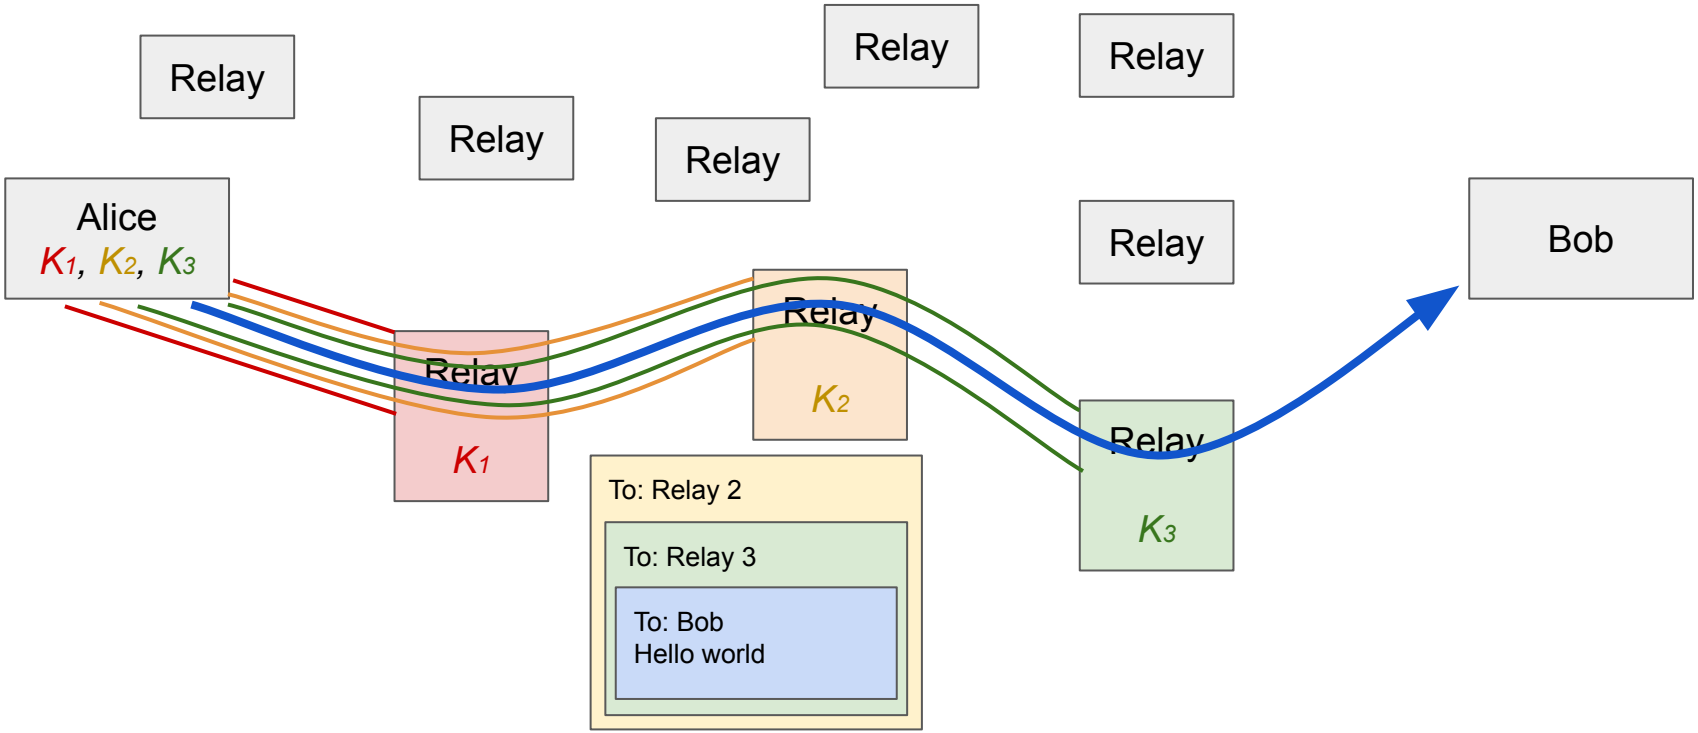
Relay
Relay
Relay
Relay
Alice
K1, K2, K3
Relay
Relay
Relay
Bob
Relay
K1
Relay
K2
Relay
K3
To: Relay 2
To: Relay 3
To: Bob
Hello world

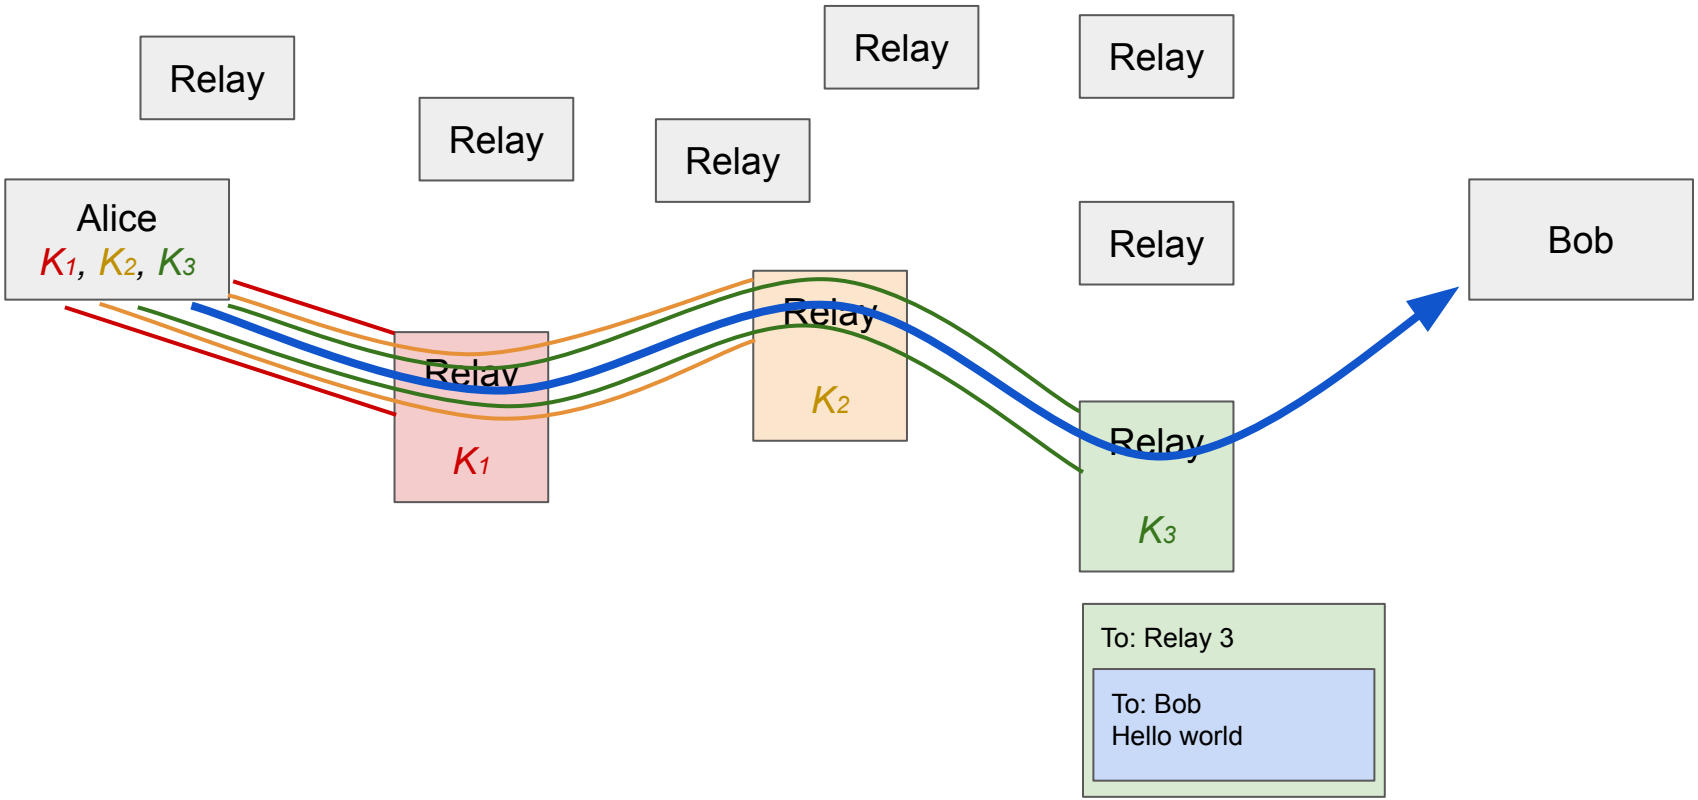
Relay
Relay
Relay
Relay
Relay
Alice
K1, K2, K3
Relay
K1
Relay
K2
Relay
K3
Bob
To: Relay 3
To: Bob
Hello world

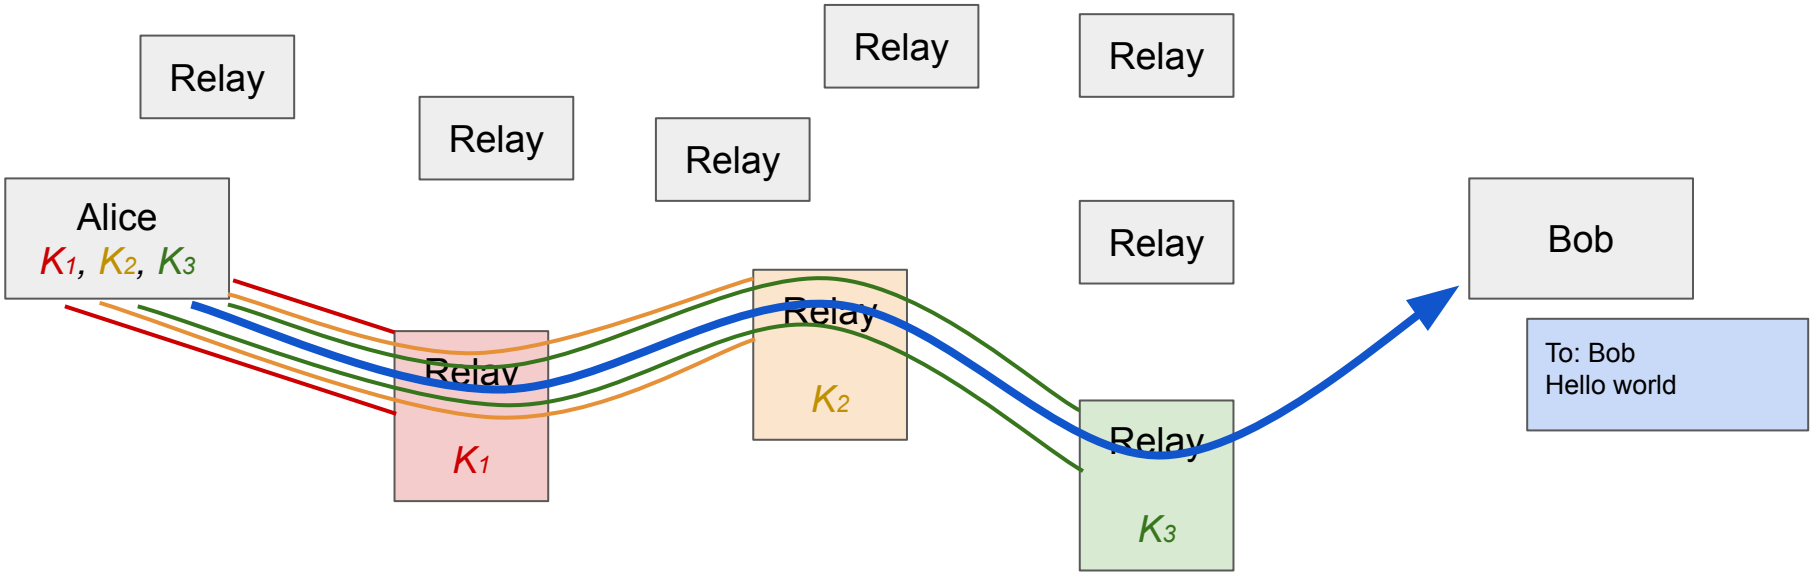
Relay
Relay
Relay
Relay
Alice
K1, K2, K3
Relay
Relay
Bob
Relay
K1
Relay
K2
Relay
K3
To: Bob
Hello world

Tor Circuits

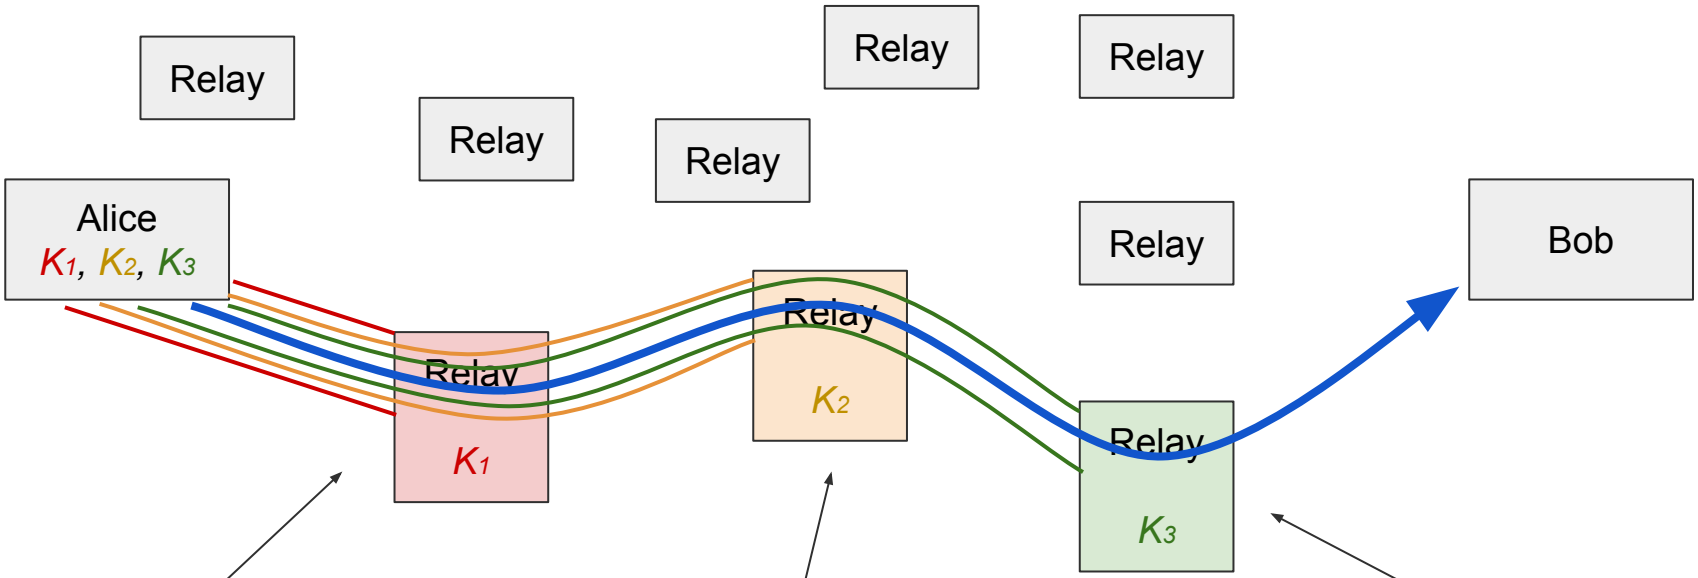

Relay 1 knows that Alice is using Tor but *not* who Alice is talking to

Relay 2 knows nothing (other than that someone is using Tor)

Relay 3 knows that someone is talking to Bob, but not who

Tor Exit Nodes

- Notice: The exit node can see the message and the recipient
  - Without collusion, the exit node doesn’t know the sender
- The exit node is a man-in-the-middle attacker
  - If the user is not using TLS to connect to the end host (using HTTP), the exit node can see and modify the traffic
  - If the user is using TLS (using HTTPS), the exit node cannot see or tamper with the contents of the traffic

Tor Exit Nodes in Practice

- Administrators of Tor exit nodes often receive abuse complaints
  - Users complain to the exit node
  - Users complain to the Internet service provider (ISP), which complains to the exit node
- As a result, most Tor relays choose to only be entry or intermediate nodes, not exit nodes
  - Exit node bandwidth is the bottleneck in Tor, not internal bandwidth
- Many Tor exit nodes are actively malicious
  - Wikileaks started by snooping traffic on a Tor exit node they ran
  - Sometimes, exit nodes will rewrite Bitcoin addresses or other cryptocurrency-related URLs!

a researcher who ran an exit node for research…
got a visit from the FBI!

Tor Weaknesses: Timing Attacks

- A network attacker who has a full (global) view of the network can learn that Alice and Bob are talking
  - Exploit a timing attack: Observe when Alice sends a message, when Bob receives a message, and link the two together
- Global adversaries are *outside* of Tor's threat model and are not defended against
  - Tor only defends against local adversaries with partial views of the network
  - Timing attacks could be defended against by delaying the timing of packets, but this violates Tor's performance goal

Tor Weaknesses: Collusion

- **Collusion**: Multiple nodes working together and sharing information
  - Collusion is adversarial (dishonest) behavior
  - Honest nodes should never share information with other proxies
  - If all nodes in the circuit collude, anonymity is broken
  - If at least one nodes in the circuit is honest, anonymity is preserved
- It is easy to form some amount of colluding nodes
  - An attacker can create hundreds nodes in the Tor network to increase the chance that your circuit consists entirely of the attacker's nodes!
- The more nodes we use, the more confident we are that they are not all colluding
  - It's much harder for 10 nodes to collude than for 2 nodes to collude
  - 3 nodes is generally considered good enough for industrial-grade security and is the default

Tor Weaknesses: Collusion

- Defense: **Guard nodes**
- Guard nodes must have a high reputation and must have existed for a long time
- Clients will always use a guard node as the entry node (by default) and the same guard node is used for a long period of time
  - Attackers' nodes are unlikely to become guard nodes
  - Because clients use the same guard nodes for a long period of time, there is only a low chance that the client will switch to an attacker's guard node

Tor Weaknesses: Distinguishable Traffic

- Tor does not hide the fact that you are using Tor
  - Example: A local adversary can see that you are sending packets to a Tor relay
  - Anonymity only works in a crowd
- Example: A TRU student sends an anonymous threatening message using Tor. The administrators notice that only one student on the TRU network is using Tor!
- Every Tor browser should be configured similarly, so network adversaries cannot distinguish any patterns in the packets
- Tor browsers should resist tracking (e.g. no tracking cookies)

Tor Weaknesses: Distinguishable Traffic

- Defense: **Tor bridges**
  - Notice: Attackers can tell you are using Tor because they can see you are connecting to an entry node
  - Lists of entry nodes are publicly available
- **Tor bridges** are entry nodes that are not available on any public list
  - Users request bridges from a separate directory, which will only give a few bridges to the user
  - There is no publicly available list of all bridges!

Tor Weaknesses: Distinguishable Traffic

- Censors can no longer block Tor based on IP addresses, but they can still distinguish traffic that looks like Tor traffic from normal traffic
- Defense: **Pluggable transports**
  - Pluggable transports change the appearance of the client's traffic to the entry node (only for bridges)
  - Obfuscates the encrypted traffic to make it "look" more like normal web traffic

- achieving true anonymity online is difficult, if not impossible, and there is always some risk of being identified or tracked.
- achieving anonymity depends on your specific needs and threat model
- Use a reputable and trustworthy VPN
- Use Tor
- Use a privacy-focused browser