

High Performance Anonymity Proxy Framework

Andrew Mao, Kefei Dan Zhou
Advisors: Micah Sherr, Boon Thau Loo

Abstract

The prevalence of the world wide web has brought unparalleled convenience to our lives, but our reliance on the Internet has severe repercussions on our privacy. To access web services, we send packets revealing our computer's IP address and application data, allowing anyone to track our activities in cyberspace. Using anonymity networks, we forward encrypted packets through one or more nodes, thus hiding the source of the packet and protecting our identity.

Uses of Anonymity Networks

- Normal people:** protect their identities from marketers and identity thieves
- Military and law enforcement:** gather intelligence online without leaving traces and communicate securely with field agents
- Journalist and activists:** protect identity of sources

Existing Solutions

Among many anonymity services, Tor is the most widely used. Tor is a network of virtual tunnels that allows users to access the web anonymously.

Disadvantages of Tor and other existing solutions:

- Requires explicit application support
- Disregard applications' bandwidth/latency requirements
- Centralized directory server are vulnerable to failure or compromise
- Poor performance on applications with high interactivity (remote desktop) or high bandwidth (streaming video)

Application-Aware Anonymity

A³, designed by Micah Sherr, is a distributed low-latency network that provides anonymity between internal network nodes.

A³ combines a distributed hash table, network coordinate system, and metric-constrained path selection algorithm, and utilizes a datagram protocol as its transport mechanism over an onion route. While the datagram protocol is more efficient than a stream-oriented protocol over several hops, it also means that applications using TCP or connecting to network-external services cannot use A³.

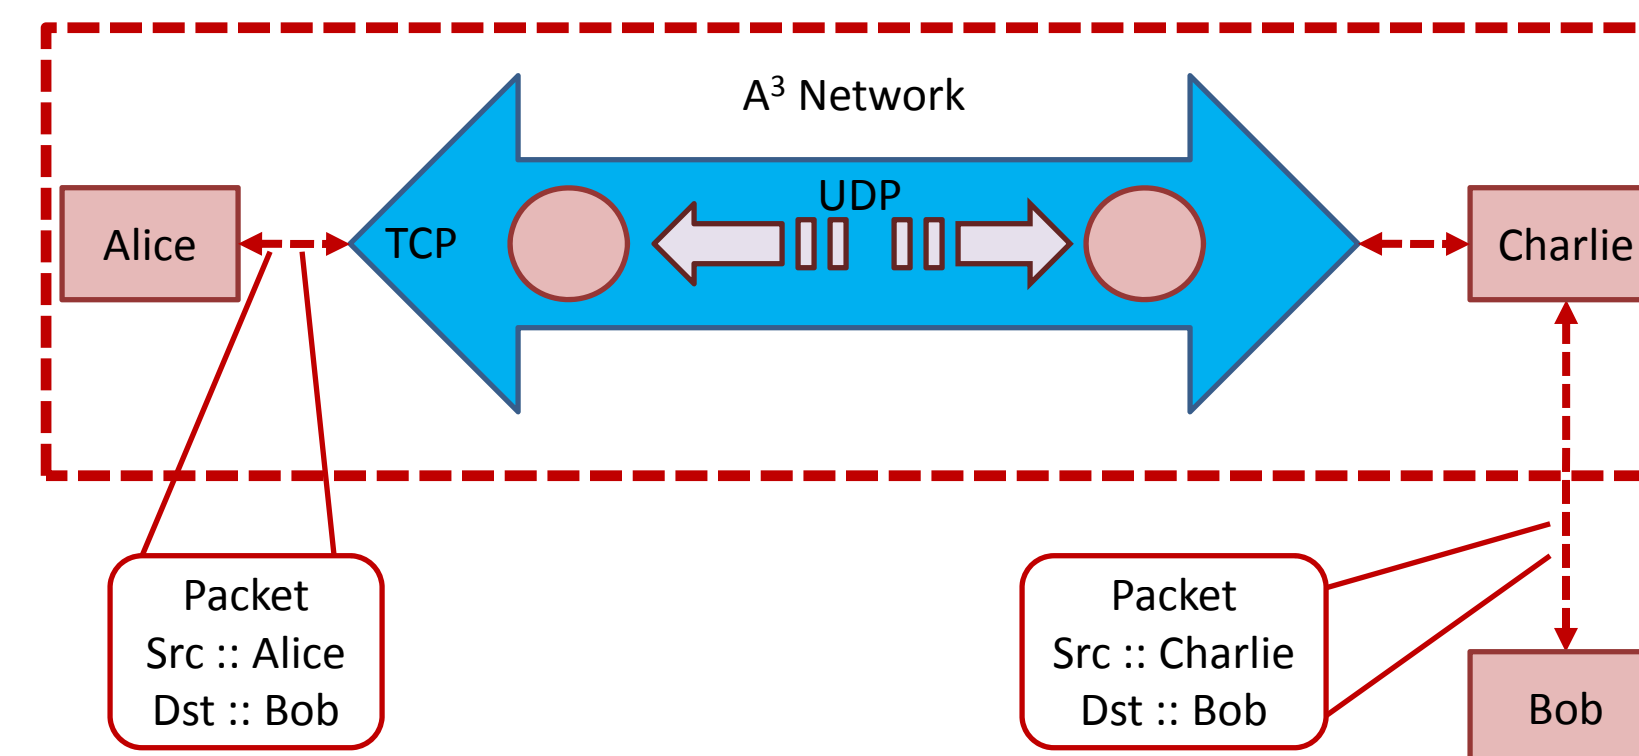
System Architecture

The high-performance proxy framework is designed to allow users to transparently use the features of A³ and other similar specialized networks. The main challenge is to create a low-overhead architecture using portable technologies and supporting all standard Internet protocols.

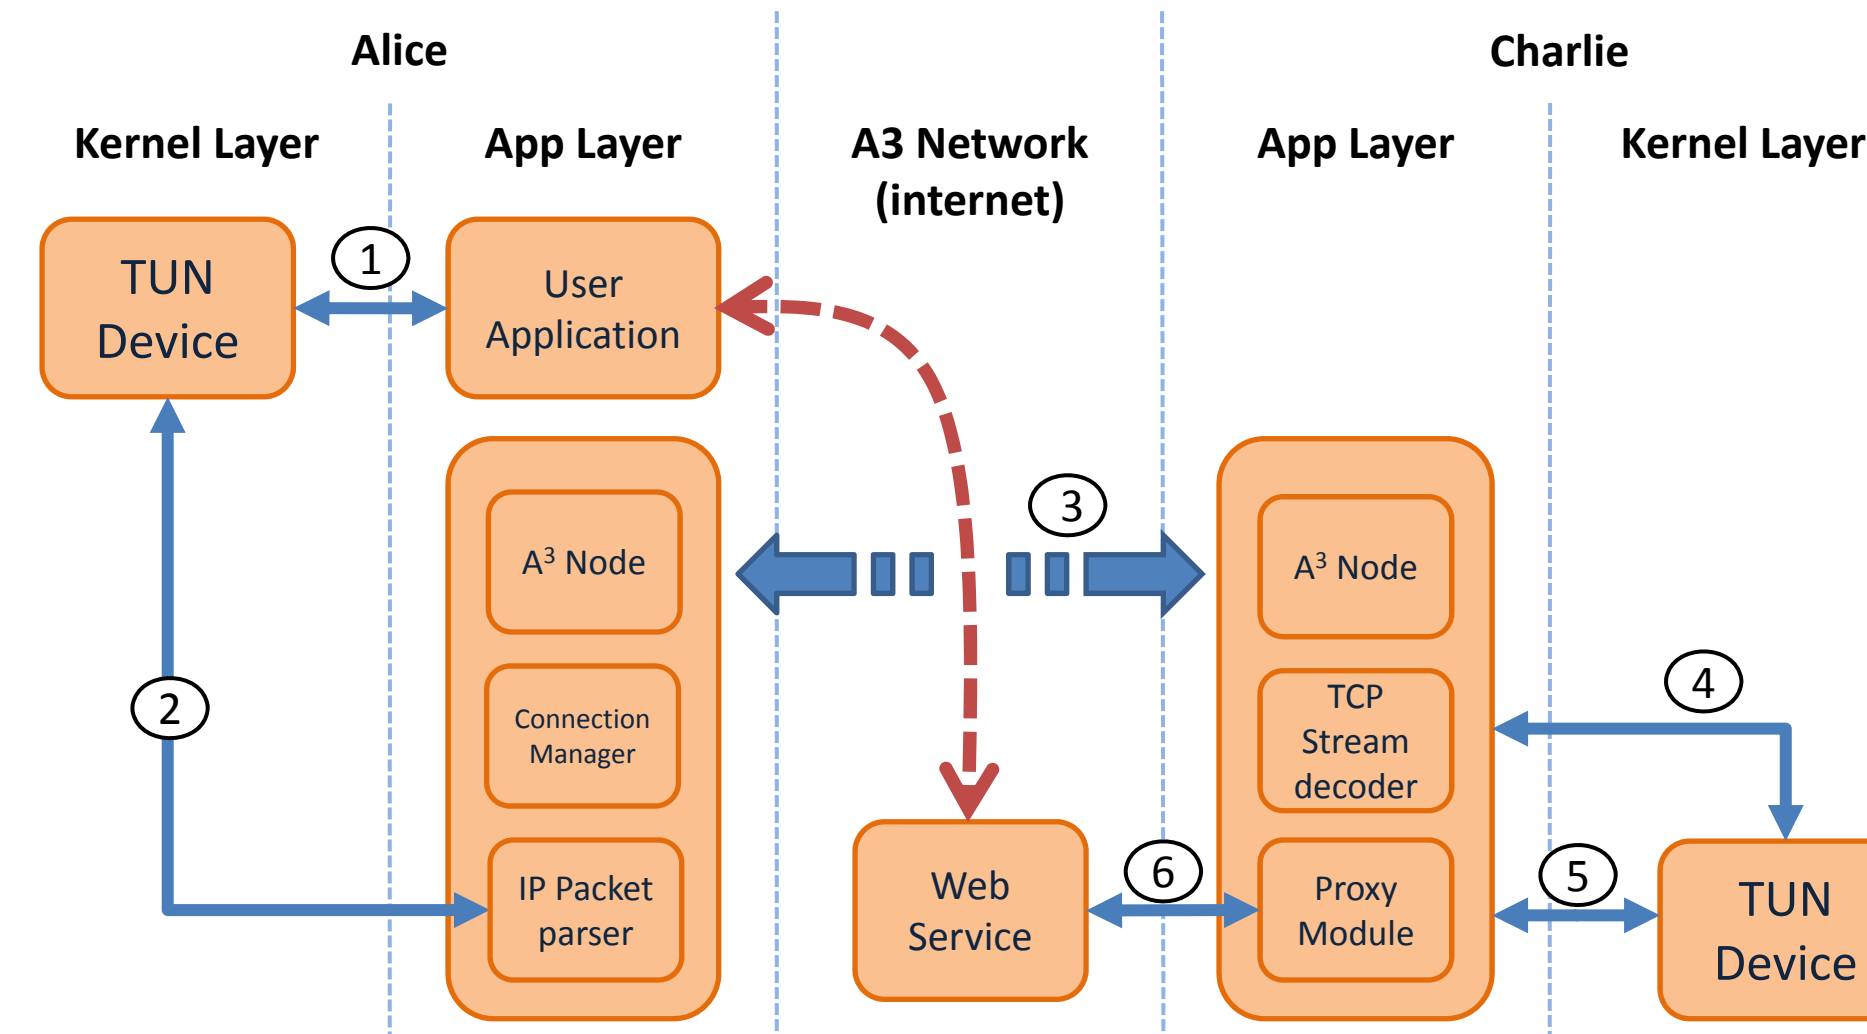
Design goals

- Allow any network application to make use of the A³ network
- Allow connections to reach destinations external to the network
- Augment the path selection algorithm to account for the last out-of-network hop

TCP over UDP protocol with anonymity
(novel implementation)



System Architecture



Design Advantages

With this anonymity framework, any application using normal internet protocols can transparently access an anonymity system to reach any location internal or external to the network. In this example, we implemented the anonymity framework over the A³ network, making A³ available to the masses.

The prototype anonymity framework, implemented in Python, showed minimal CPU overhead and was able to reach a transfer speed of 10 Mbps without any algorithm optimizations. In addition, we implemented a geolocation module to enhance the existing path selection algorithm by selecting an exit node geographically closest to our destination.

Other Advantages


- Applications don't need explicit proxy (i.e. SOCKS) support
- Portable design – most platforms support virtual network devices
- Compact - utilizes kernel implementation of TCP/UDP protocol
- Forward and backward compatible with core IP protocols
- Preserves anonymity of A³, with packet rewriting so that all virtual network devices can use arbitrary IP addresses

Future Work

The anonymity framework can also be used in a reverse direction, enabling anonymous services to originate from within the A³ network. By using the distributed hash table, public hosts can replicate the services of an internal network server without compromising the server's identity.

Deployment

Standard Proxy



A³ Network with anonymity framework

