Leveraging Heterogeneous Programmable Data Planes for Security and Privacy of Cellular Networks, 5G & Beyond

JUNO-3 PI Meeting

April 3, 2023

K. K. Ramakrishnan

University of California, Riverside, CA

&

Timothy Wood

George Washington University, Washington D. C.

Toru Hasegawa, Yuki Koizumi and Junji Takemasa

Osaka University

&

Toshiaki Tanaka and Jun Kurihara

University of Hyogo

Importance of Programmable Data Planes for Cellular Performance & Security

- Cellular networks support a growing amount of traffic from mobile and Internet of Things (IoT) devices
  - Implementations moving to software-based environments: potential for increased vulnerability to security attacks, including violation of user privacy through eavesdropping
  - More slow and stealthy attacks: difficult to detect, need more memory and compute capacity
- Our project will use high speed programmable switches, SmartNICs and end-host servers supporting NFV to provide security monitoring and privacy protection solutions
  - Develop an efficient, high performance cellular network security solution
- Project builds on decade-long work on switching, SmartNIC and NFV work by PIs and collaboration between the PIs based in the US and Japan
- Monitoring: we will develop a collaborative filtering system for real-time monitoring of cellular traffic
  - Most of the traffic processed by high-speed programmable switches to extract coarse-grained metrics
  - Suspicious traffic redirected to programmable SmartNICs or the host for detailed forensics
- Privacy Protection: utilize P4 programmable switches for anonymization and privacy protection
  - Lightweight Anonymization at Terabit rates within the network layer with high-speed P4 switches
  - Use traffic morphing to handle fingerprinting attacks
- Societal Impact: Provide strong threat prevention and privacy preservation of cellular network users

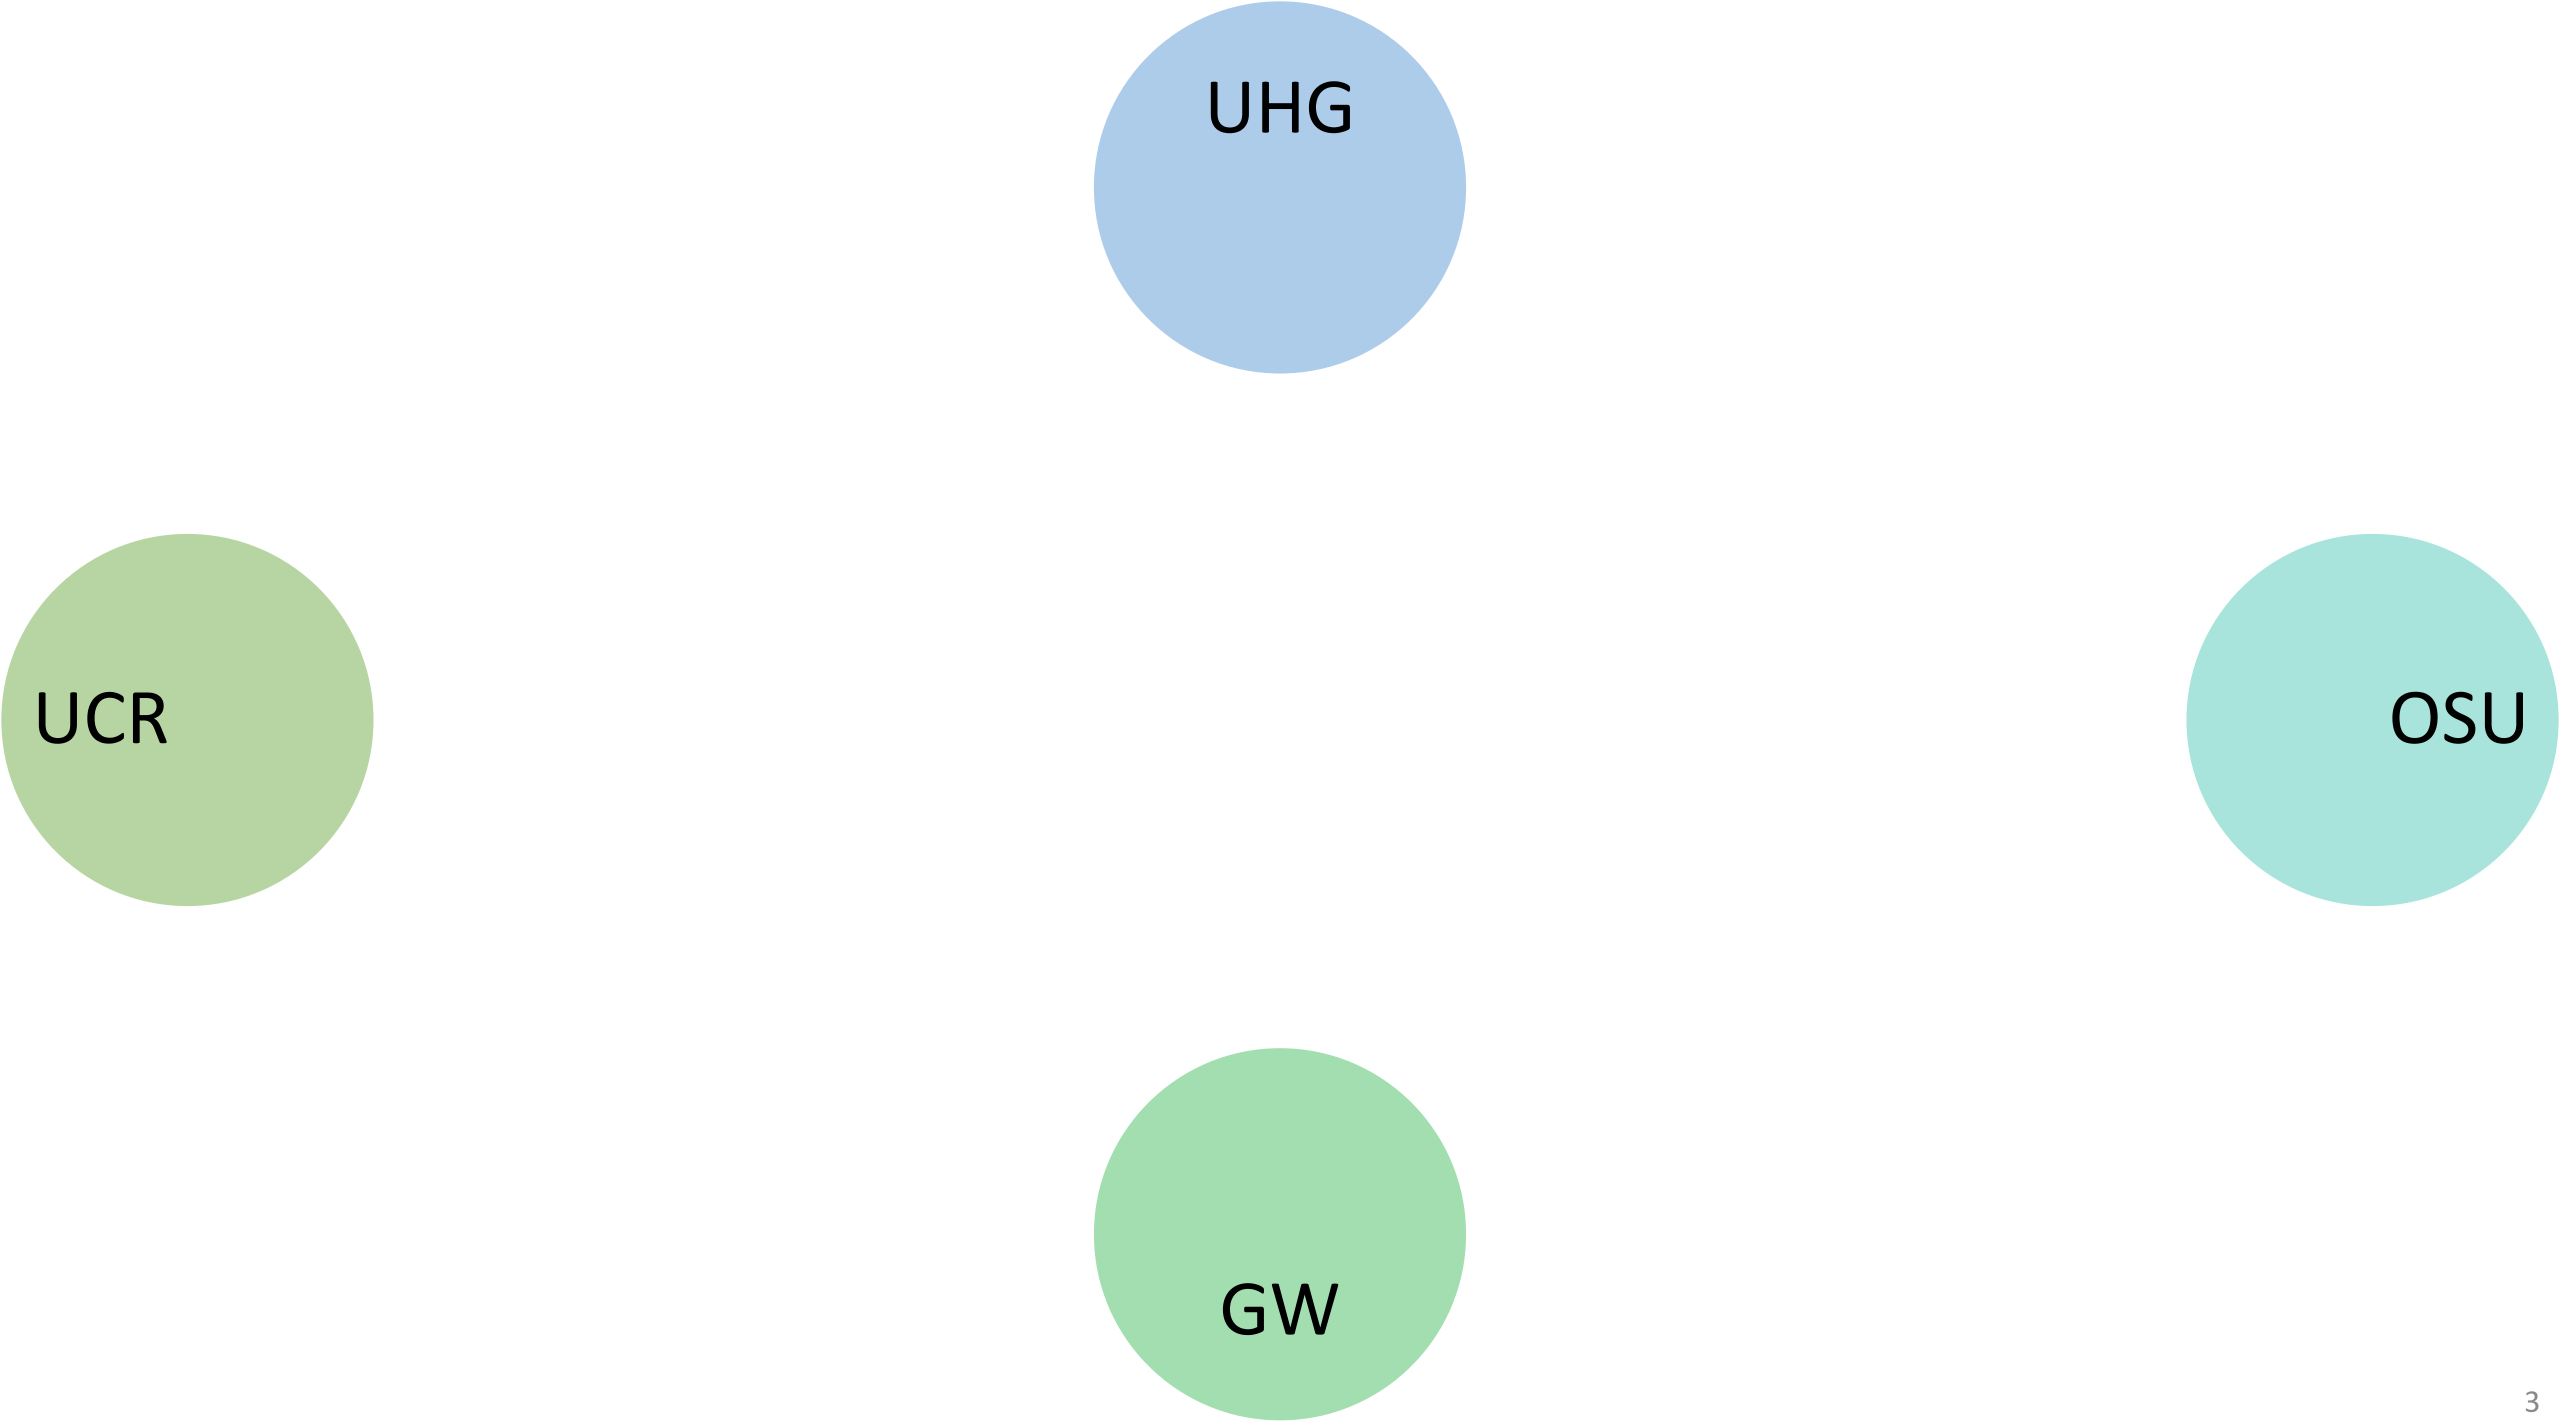
UHG
UCR
OSU
GW

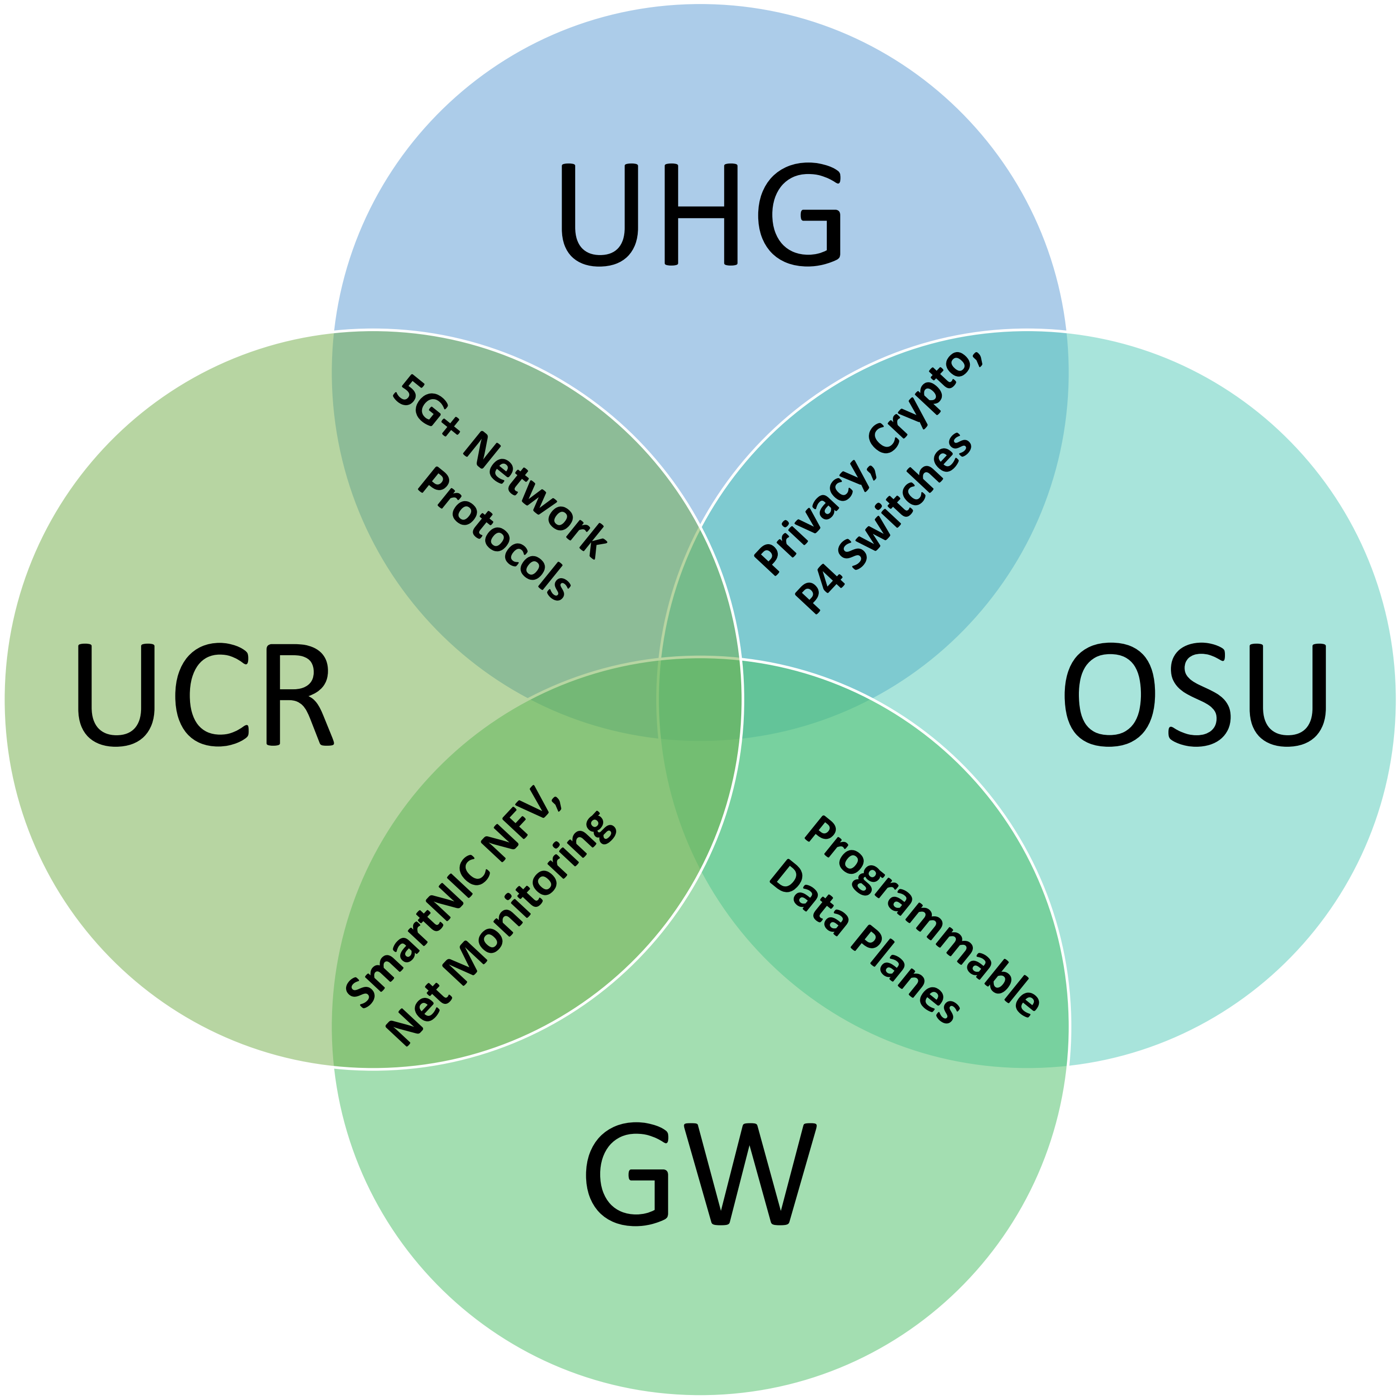
UHG
UCR
OSU
GW
5G+ Network Protocols
Privacy, Crypto, P4 Switches
SmartNIC NFV, Net Monitoring
Programmable Data Planes

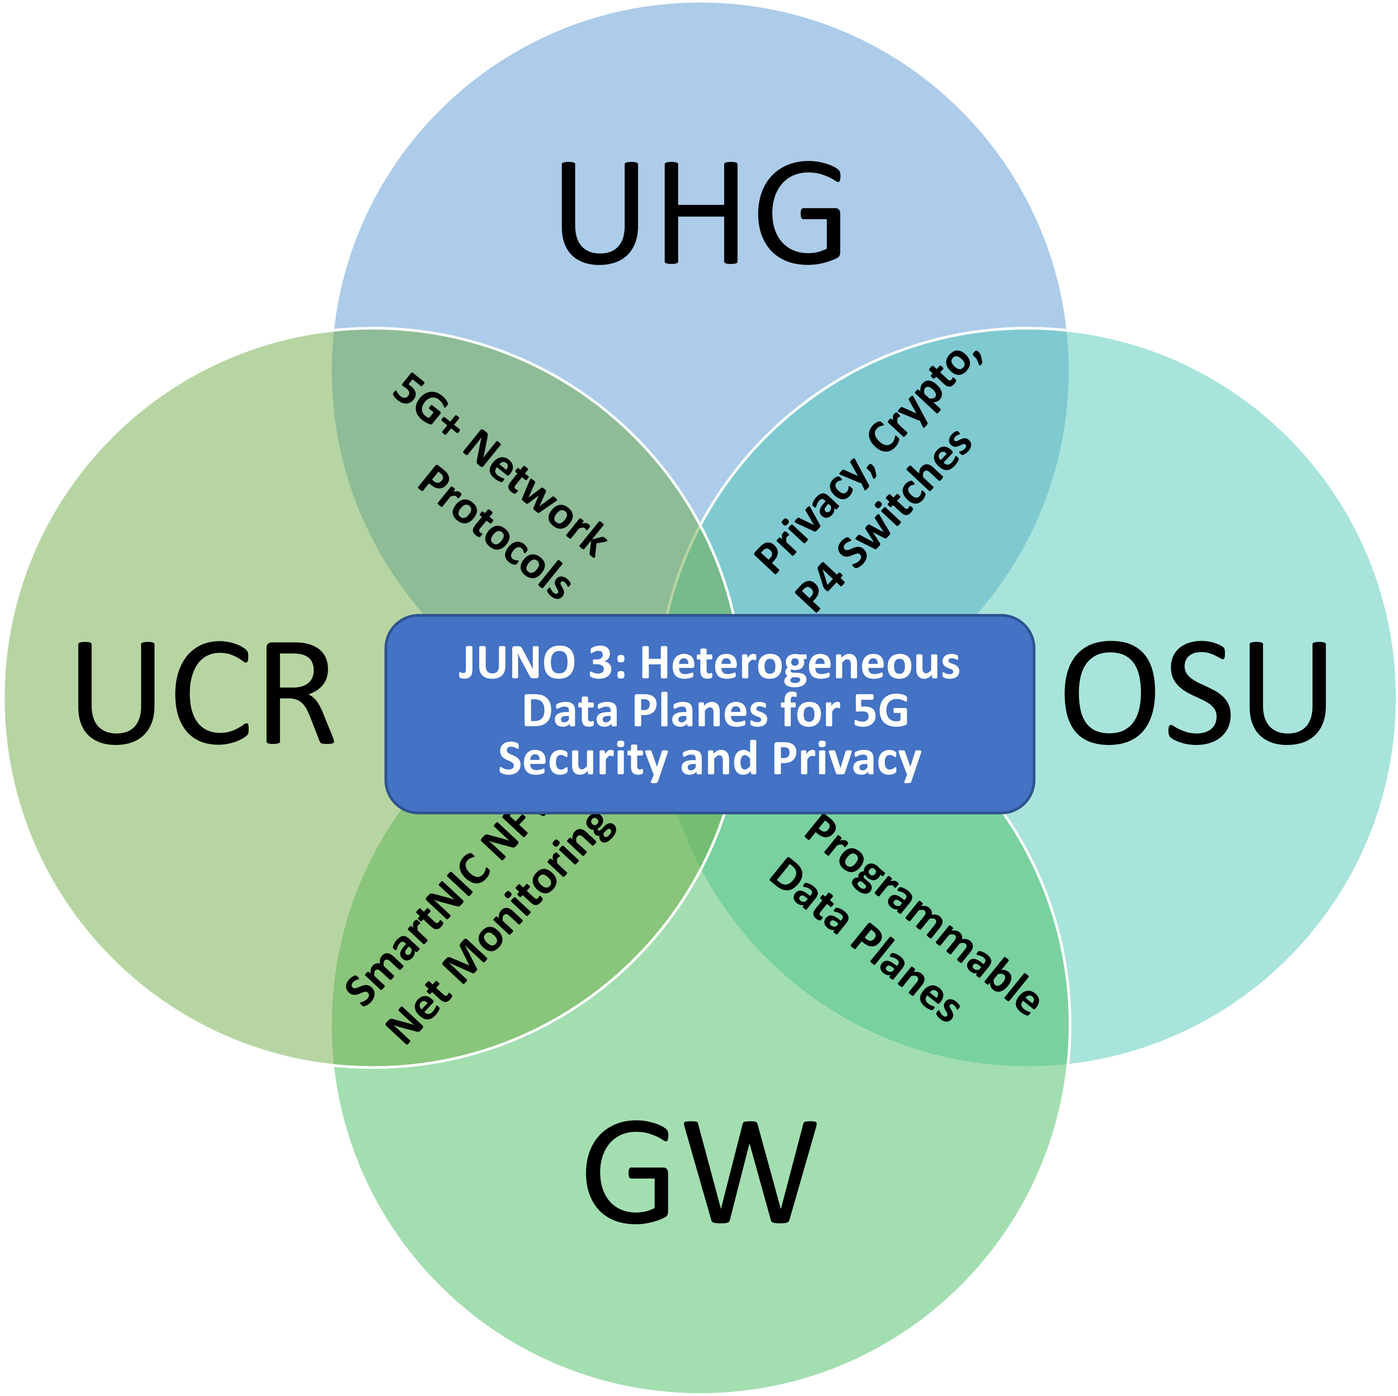
UHG
UCR
OSU
GW
5G+ Network Protocols
Privacy, Crypto, P4 Switches
SmartNIC NF
Net Monitoring
Programmable Data Planes
JUNO 3: Heterogeneous Data Planes for 5G Security and Privacy

Overview of Proposed work

Task 1: Heterogeneous Data Plane Framework

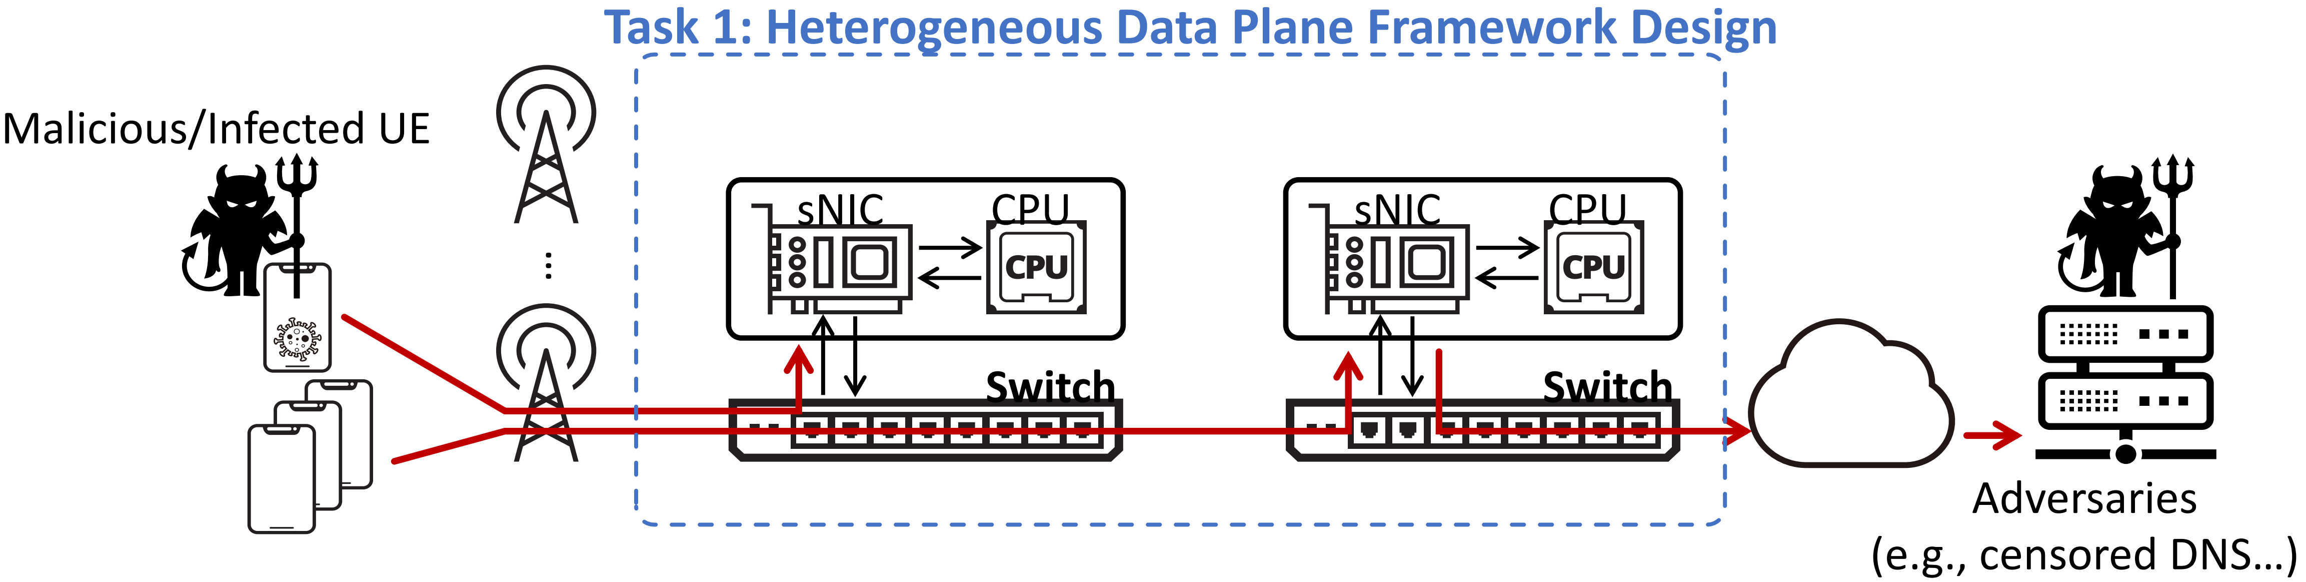

- Design a programmable data plane framework to leverage heterogeneity:
  - P4 Switches: High bandwidth (Tbps), limited memory (MBs)/programmability
  - SmartNICs: Moderate bandwidth (Gbps), moderate memory (GBs); Programmable
  - Host CPUs: Limited bandwidth, large memory (TBs), general purpose CPUs

Task 1: Heterogeneous Data Plane Framework

- Goal:
  - Coordinate protocols between various programmable network devices to overcome their limitations while optimizing their strengths
- Key techniques
  - **Cooperative flow filtering** and state caching across heterogeneous programmable network devices
  - **In-network ML inference:** get high throughput by using the pipelined processing of packets on P4 switches and slicing of GPUs on hosts
  - **Optimization algorithms** that use models of component capabilities to effectively determine which traffic monitoring modules to place on which types of data plane hardware

Task 2: Real-time Network Monitoring

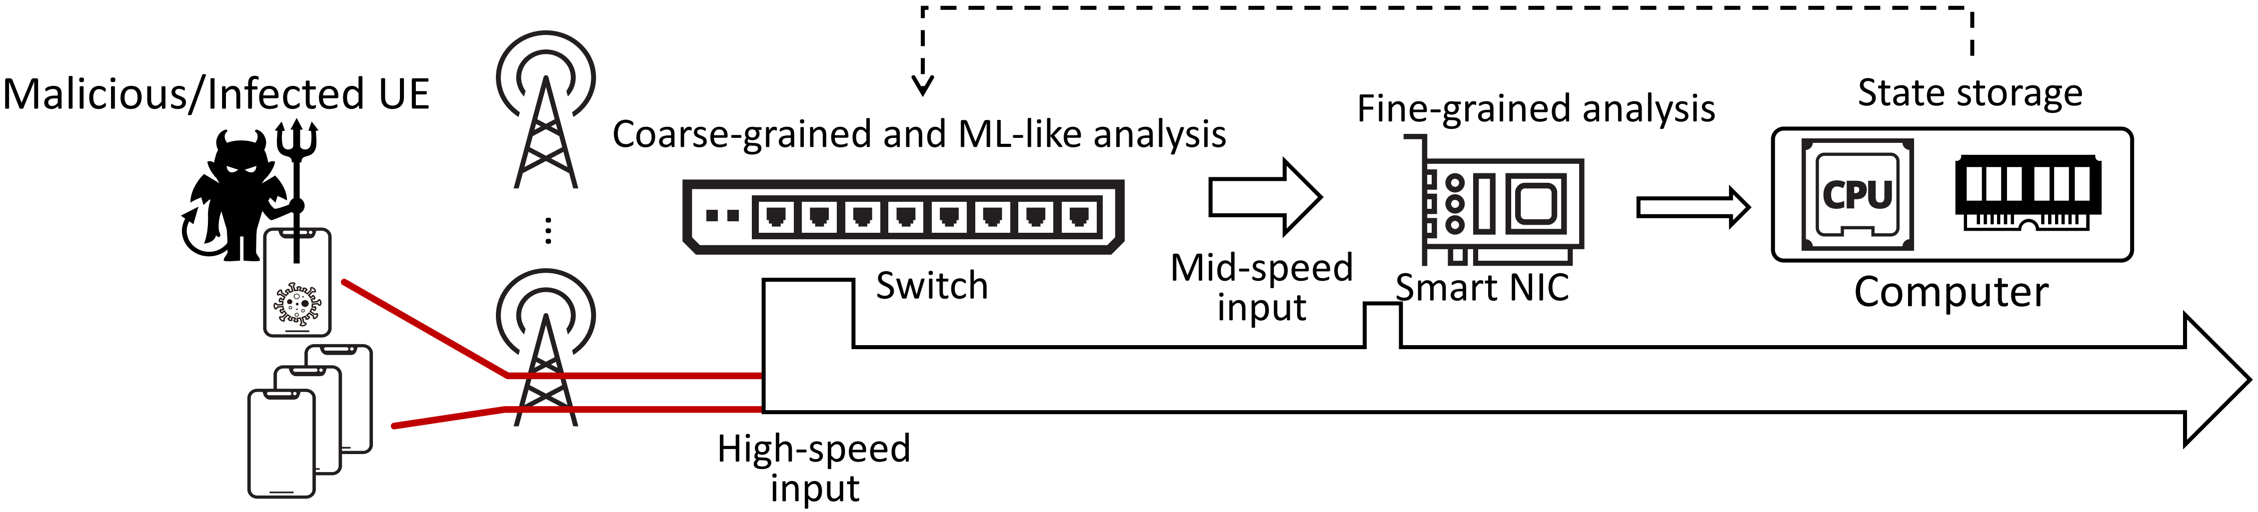

- Use heterogeneous data plane for real-time monitoring of 5G networks
  - 5G core based on multi-tier programmable data and control plane components
  - Line rate traffic monitoring for 5G resource management and anomaly detection

Task 2: Real-time Network Monitoring

- Goal: Leverage programmable dataplane at base stations and 5G core to enable real-time monitoring and security
- Applications:
  - **Resource management and mobility prediction** through analysis of control and data plane traffic in 5G core
  - **Protecting UEs** from both volumetric and slow attacks through UPF-based monitoring of control and data planes
  - **Securing 5G infrastructure** from rogue UEs and gNBs by aggregating monitoring data across devices

Task 3: Real-time Privacy Preservation/Attack Mitigation

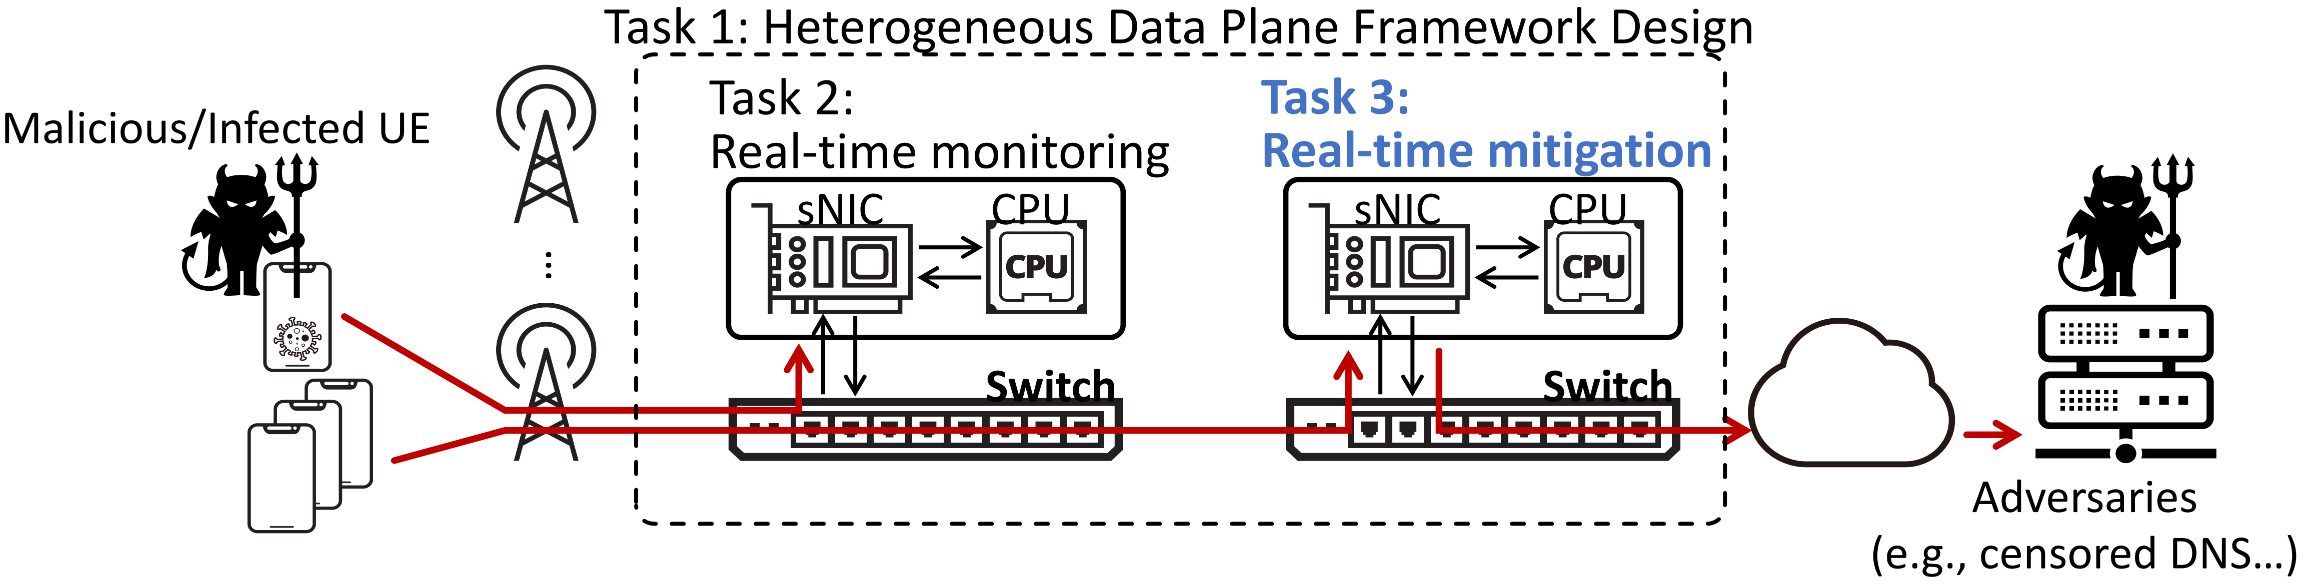

- Prevent attacks and preserve privacy for mobile users
  - Lightweight traffic encryption protocols to provide relationship anonymity
  - Traffic obfuscation techniques to prevent fingerprinting attacks

Task 3: Real-time Privacy Preservation/Attack Mitigation

- Goal: Privacy protection for users who access Web sites and various data
- Key techniques
  - **IP address obfuscation** through light weight anonymous routing protocol at the network layer and its implementation on a P4 switch
  - **Traffic morphing** to prevent fingerprinting attacks by inserting dummy packets and data
  - **DNS privacy protection** with a novel distributed and multiple-hop based approach of anonymization technique on DNS queries, and its design and implementation of its ‘lite-version’ on P4 switch

Task 4: Integration

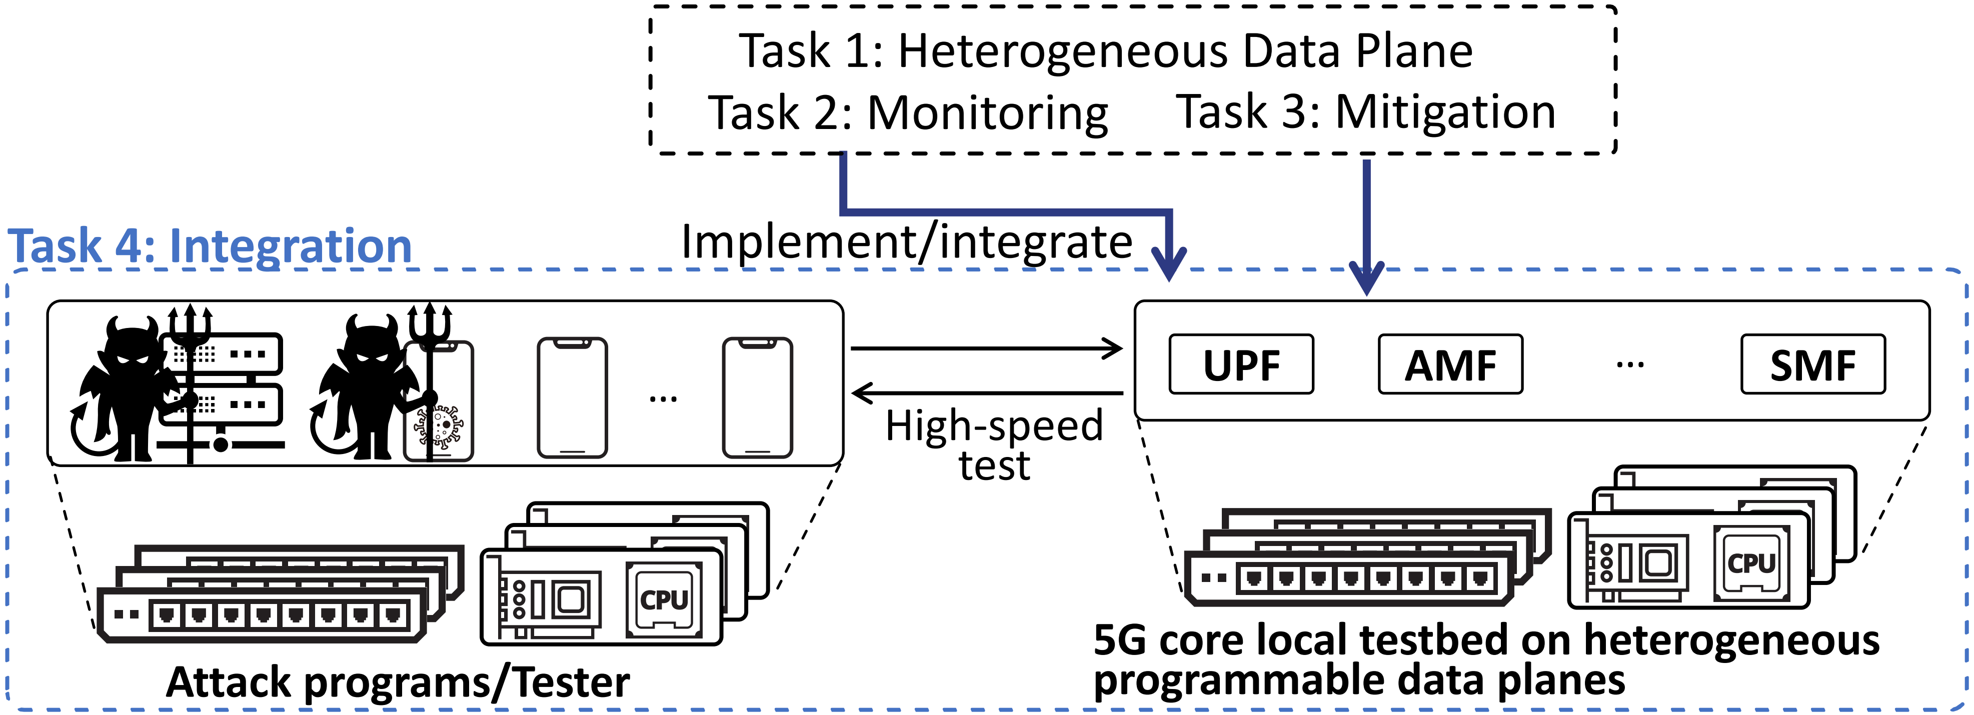

- Deploy a holistic system to study security and performance properties
  - Evaluate traffic monitoring and privacy preservation techniques on 5G testbed
  - Optimize the combination of hosts, P4 switches, and SmartNICs

Sample Projects

1. SmartNIC Accelerated Traffic Monitoring in the Cellular Core
2. Lightweight Anonymity Protocols for P4 Switches

Synergy: Traffic Monitoring in the Cellular Core

- 5G Core (5GC) is the heart of a 5G mobile network
  - Establishes reliable, secure connectivity to network for end users
  - Mobility management, authentication/authorization, and policy management
- The 5GC is comprised of several control plane and dataplane NFs.
- The User Plane Function (UPF) is the primary dataplane component of the 5GC

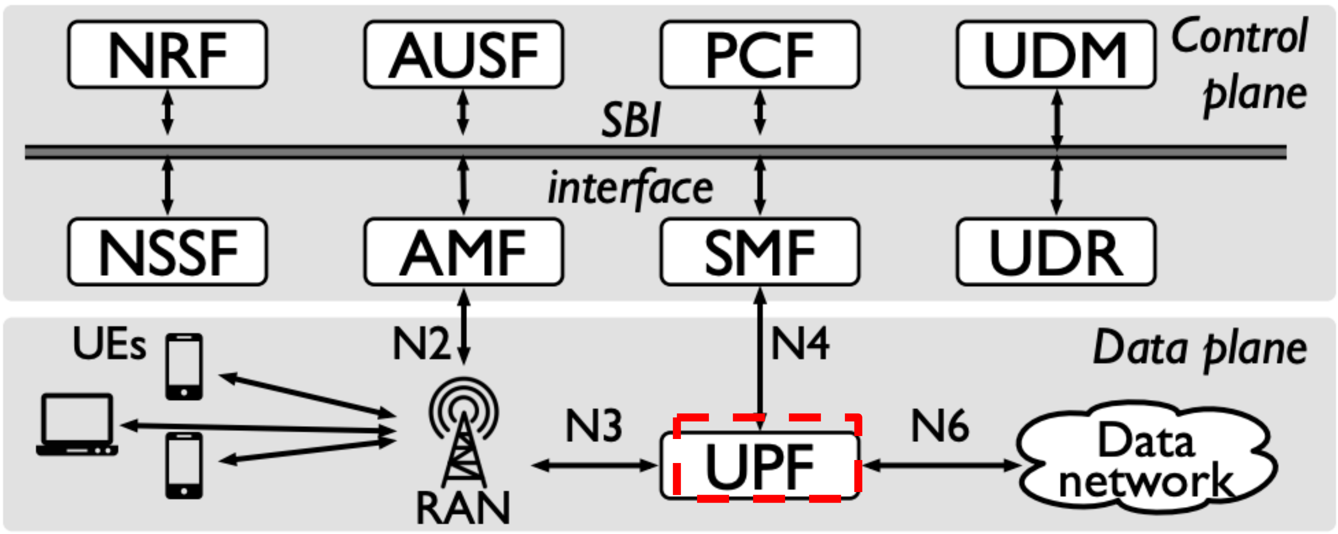

User Plane Function (UPF) : The 5G Dataplane

- Interconnect point between the mobile infrastructure and the data network.
- Implements complex rules for forwarding and tunneling.
  - Processes packets belonging to different sessions , priorities, including shaping and policing.
  - Must have high throughput and low latency→ spurious retransmissions and also packet loss.

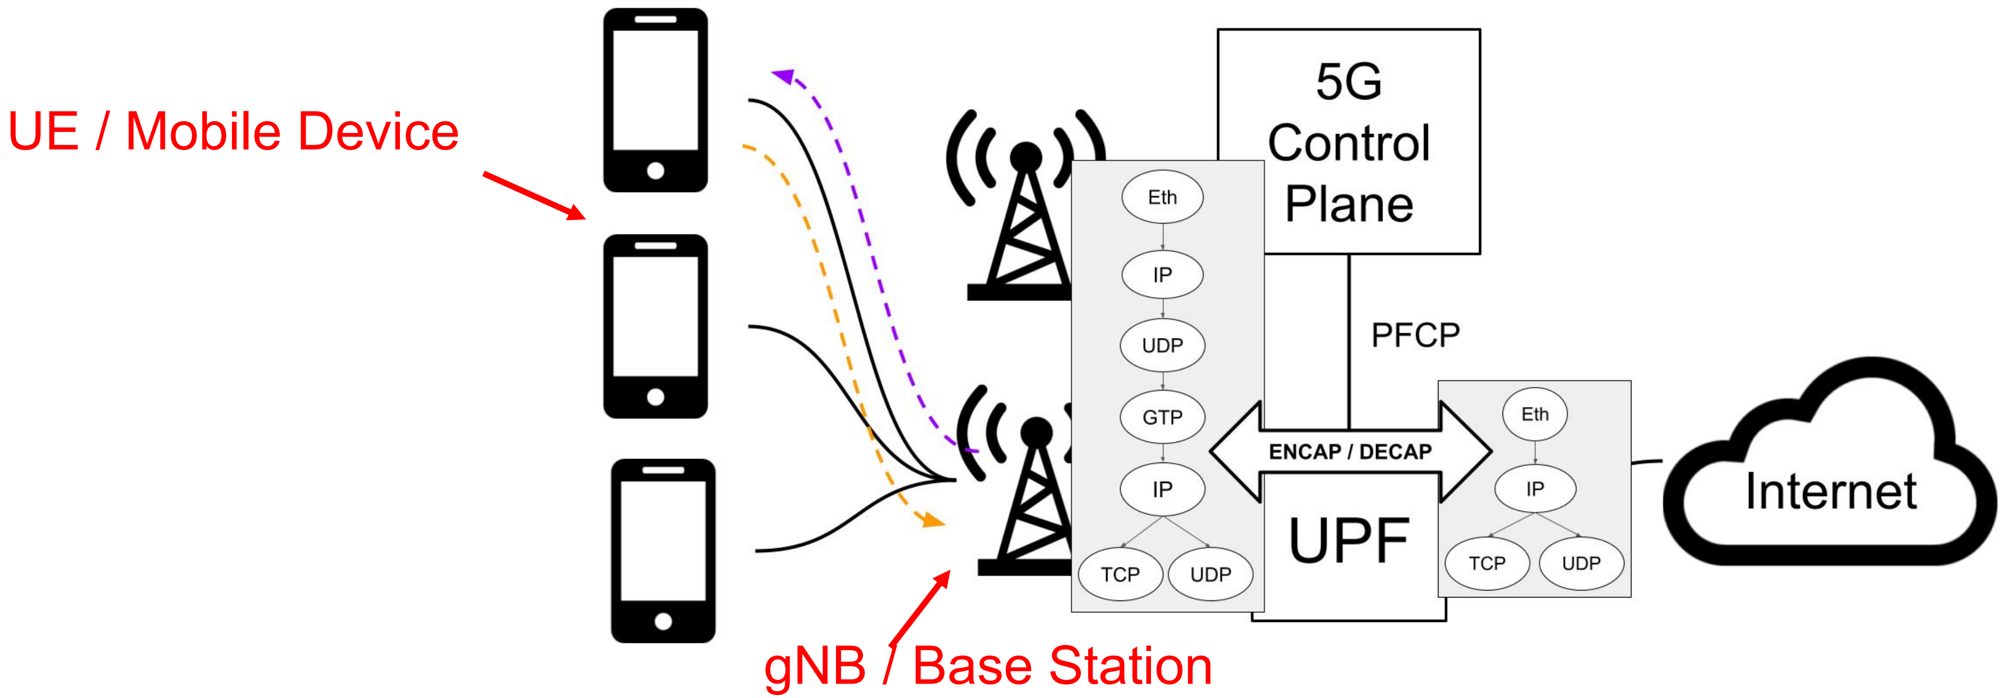

UPF: Perform flow-state dependent processing

- Has to be aware of the idle and active transitions (Paging) → Save energy.
- Radio association change from one base station to another → Mobility.
- Lots of protocol msg exchanges that change the dataplane
  - Change must be affected quickly → Packets can be forwarded

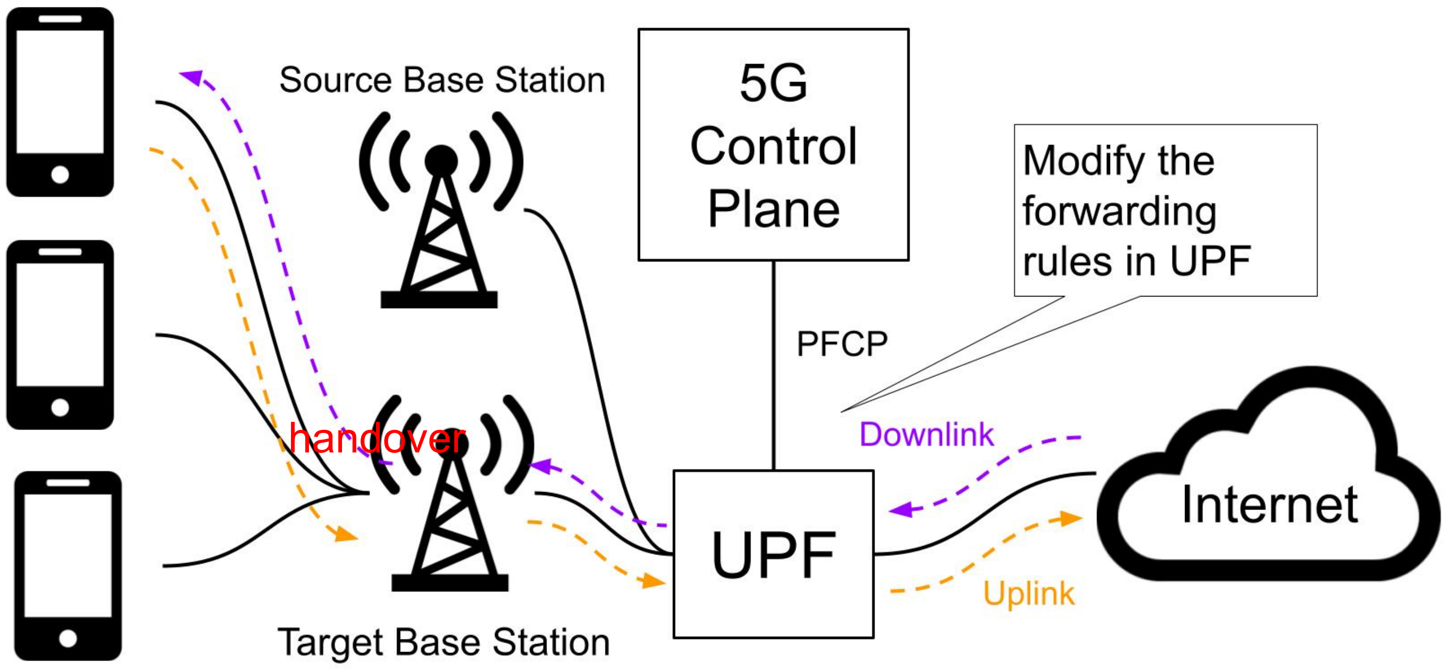

Improving Control/Data Plane throughput and latency

Require programmability: Rules have to be rapidly and frequently modified

- Rules change when mobile device goes through handover, attach/detach, idle-active transition.
- Programmability of the SmartNIC can be tremendous advantage
- Take advantage of the tight coupling by using a SmartNIC accelerator.

- **With a performant UPF, we can now enhance functionality: monitor the network; mobility prediction**
  - Prepopulate the state in 5GC to accelerate the control plane event → **Quicker handover.**
  - Enabler: Monitor control plane traffic in the SmartNIC → Predict Mobility → Prepopulate State

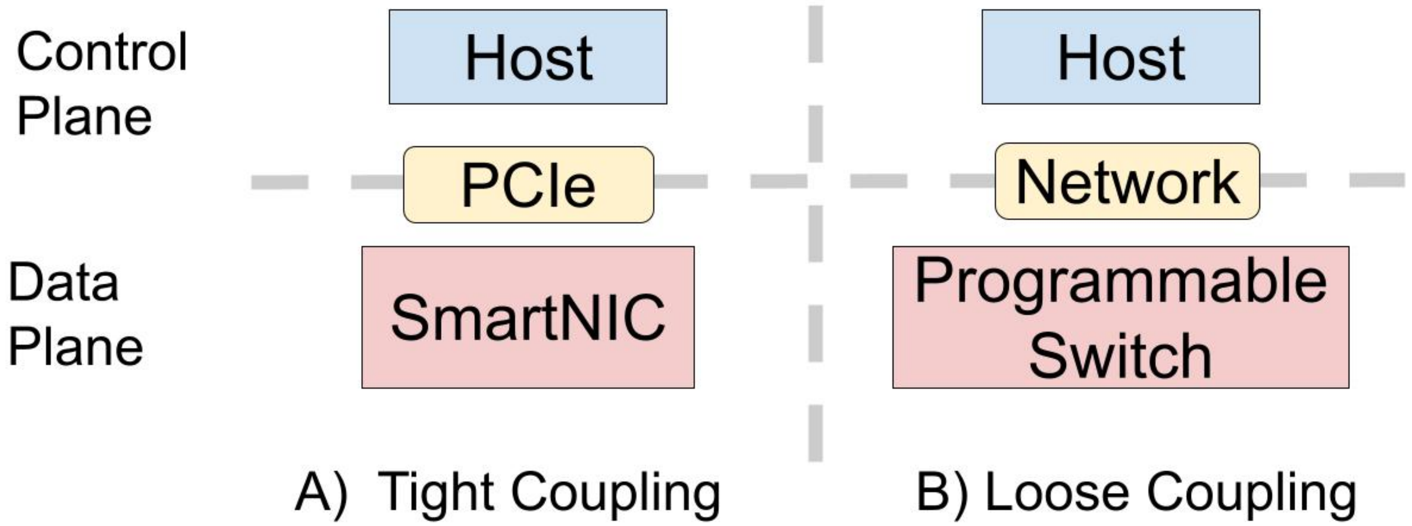

Synergy: P4 Programmable SmartNICs Primer

- We implemented Synergy on a Netronome Agilio LX 2x40GbE.
- 96 P4/ Micro-C programmable flow processing cores
  - Parallelism for packet forwarding processing.
- Memory Hierarchy (SRAM + DRAM)
  - Large memories (DRAM), slow to access → Buffering.
  - Smaller memories, bounded access latency → Forwarding.
- Dynamically push rules into the SmartNIC
  - Pushing rules in response to handovers / paging.
  - Takes advantage of the tight coupling with the host

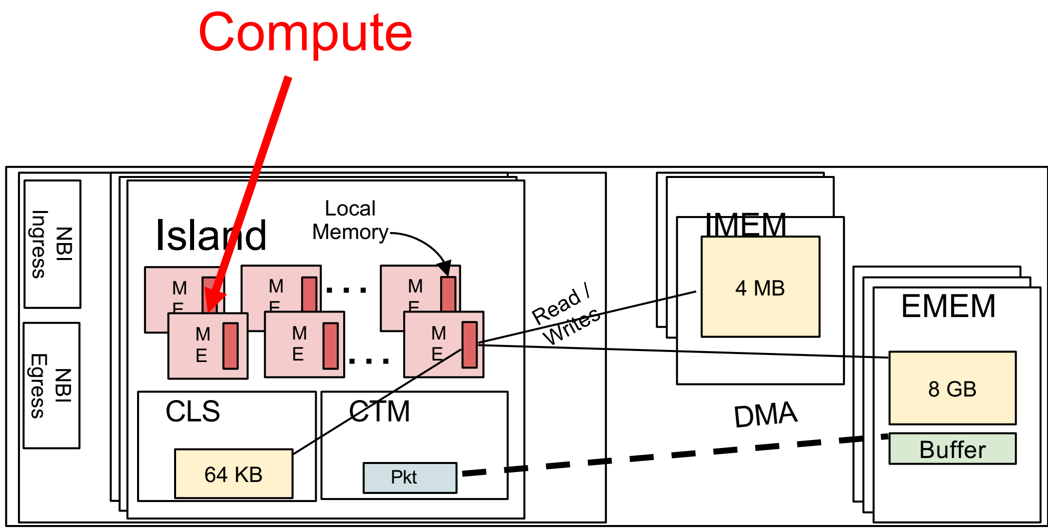

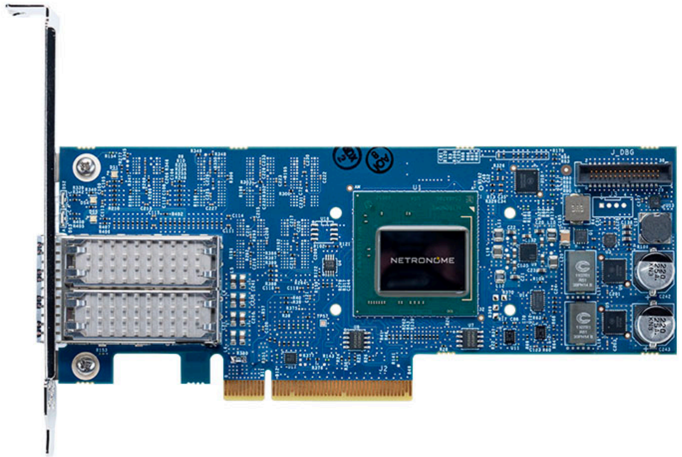

Application - Monitoring and Mobility prediction

- Carryout monitoring in the SmartNIC to do mobility predictions.
- Input: Mobile Device's Location report that the AMF requests; Output: UE's next base station.
- Model: Recurrent neural network model [1]
- **SUMO-based vehicular mobility dataset:**
  - 700k vehicle trips across 247 gNBs. [2]
  - At each point in time, the UE connects with the gNB providing the best communication conditions, measured in terms of path loss.

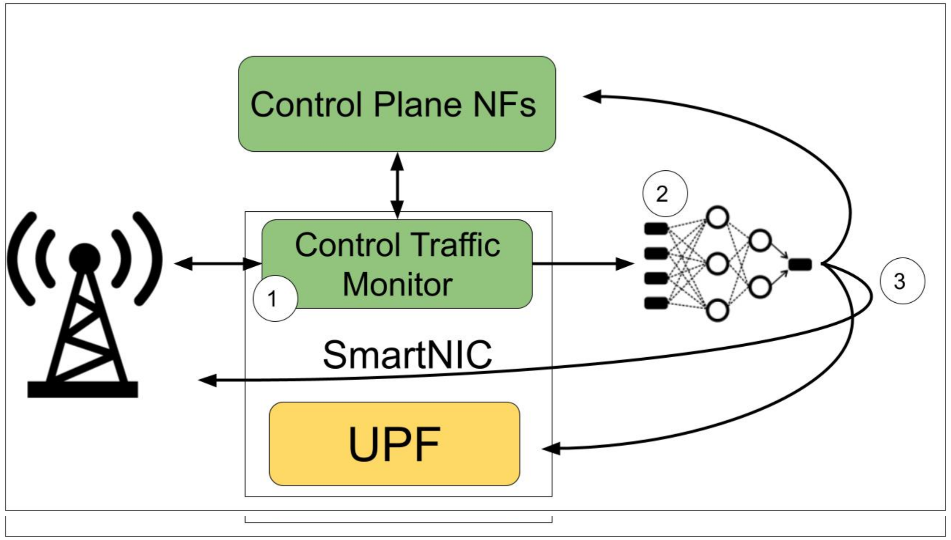

Benefit of state pre-population

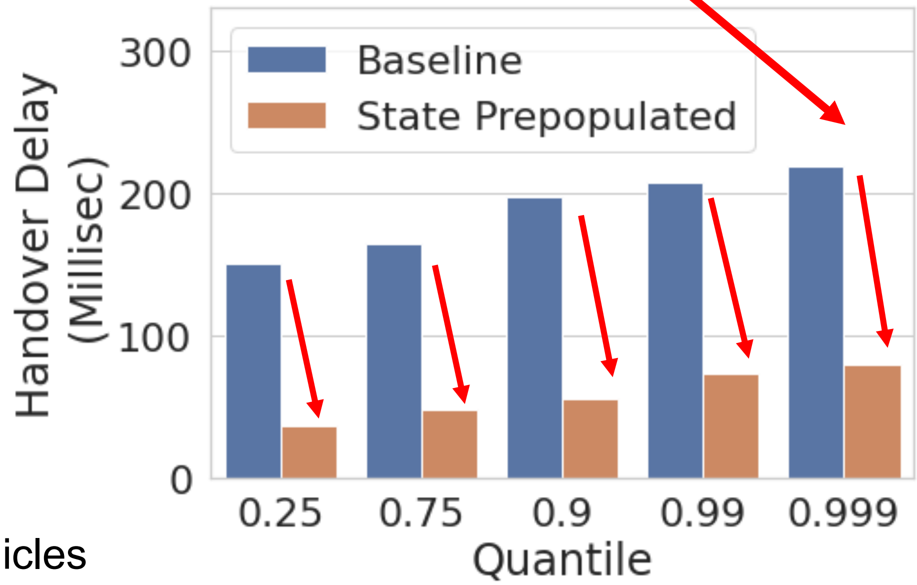

[1] Mobility aware and dynamic migration of MEC services for the internet of vehicles
[2] ttp://kolntrace.project.citi-lab.fr/

Benefit of Mobility Prediction and Tight coupling

- Not possible to always have correct predictions
- Predict mobility with lookahead ~5 sec
  - Lookahead > 5 sec: Poor accuracy
  - Lookahead < 5 sec: Too many predictions
- Correct predictions reduce handover delays

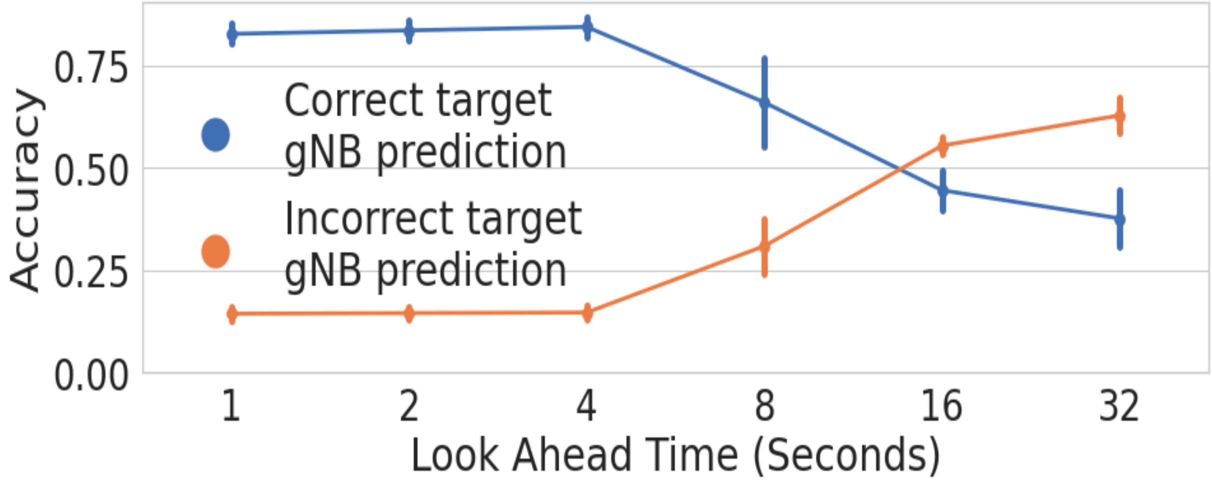

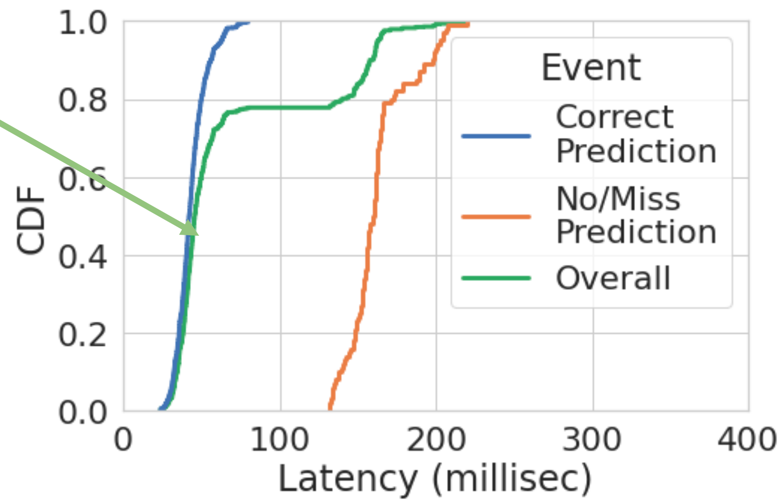

Benefit of Mobility Prediction and Tight coupling

- Low programming latency → Lower Handover Delays
- Programmable Switches have been emulated here.
- Constraint with Programmable Switch: ~1200 rules/sec[1]

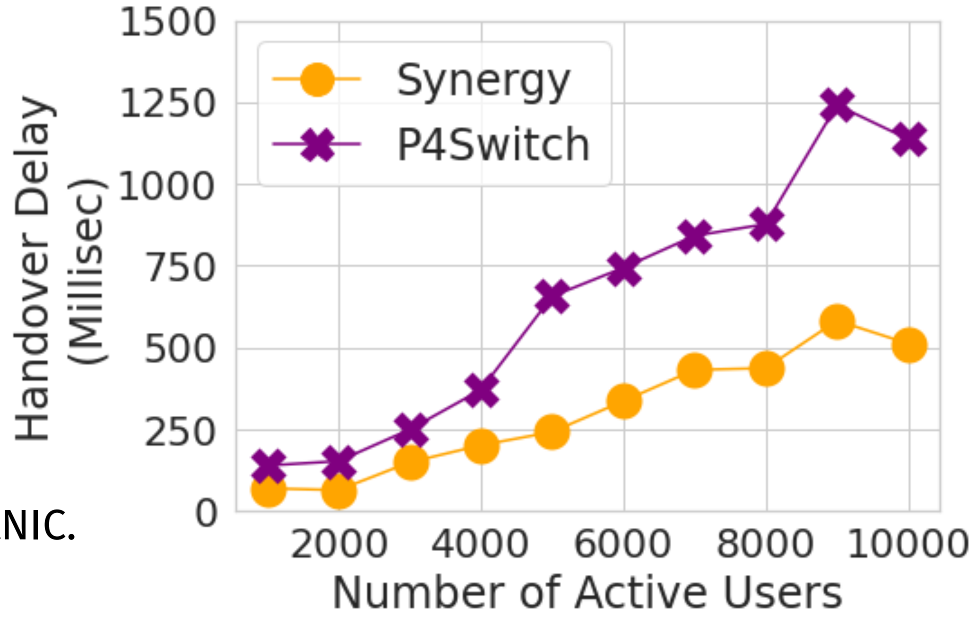

Summary

- Synergy accelerates the UPF throughput and reduces latency with the SmartNIC.
- SmartNIC-based UPF buffers packets:
  - Speed up control-to-data plane interaction
  - Reduces the impact on Handovers and Paging
- Handover latency reduced by monitoring and mobility prediction to prepopulate state.

[1] NetWarden: Mitigating network covert channels while preserving performance (NDSS'20)

Background: Tor vs. Lightweight Anonymity Protocol in the Network Layer

- Tor assumes global adversary
  - Tor relays encrypt both headers and payloads to prevent a global adversary from correlating packets intercepted at multiple places

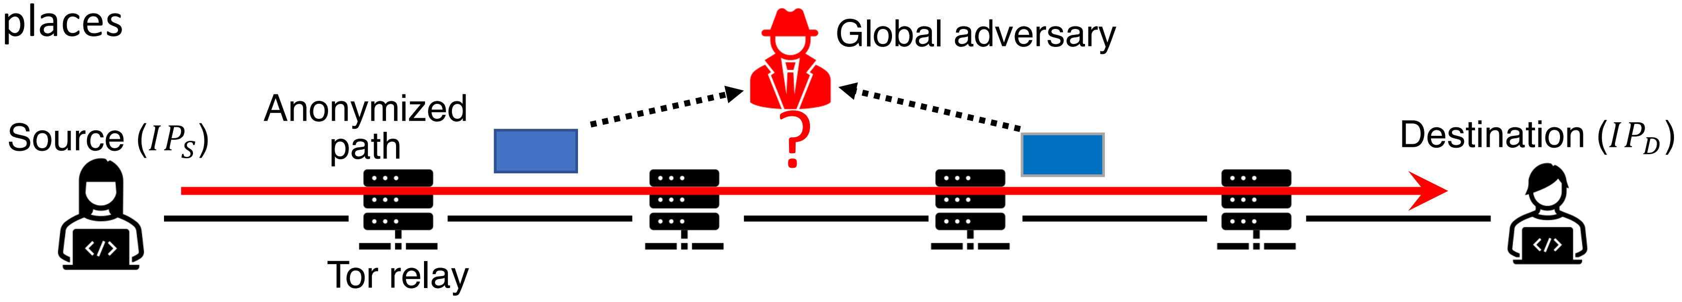

- Lightweight anonymity protocol assumes weak adversary
  - Routers obscure either source or destination address to prevent a local adversary from correlating source and destination addresses from packets intercepted at one place (relationship anonymity)
  - Advantages
    - High-speed forwarding owing to encryption of only headers
    - Short path length owing to the underlying IP routing path

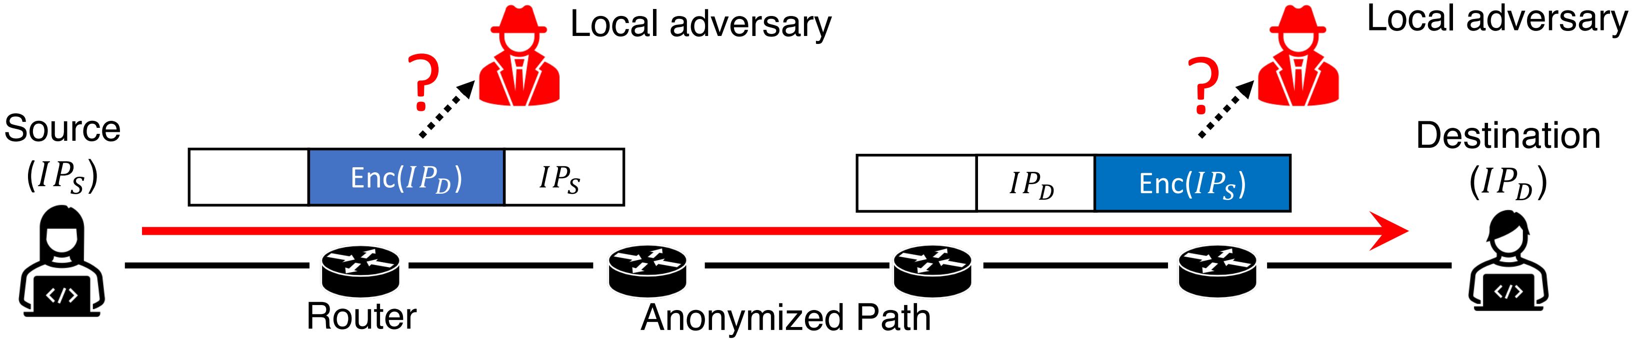

[4] Yutaro Yoshinaka, Junji Takemasa, Yuki Koizumi, Toru Hasegawa, “Design and analysis of lightweight anonymity protocol for host- and AS-level anonymity ,” Computer Networks, Volume 222, 109559-109559, Feb. 2023.

Problems and Goals

- Problems
  - Existing implementations only target **software switches**
    - Terabit speed, required for backbone routers, is unachievable on software switches
      - PHI [Chen 2017] achieves **120Gbps**
      - HORNET [Chen 2015], a protocol performing onion routing, achieves **93.5Gbps**
- Goals
  - Implementing lightweight anonymity protocols on **programmable switches** for anonymous communication at **terabit-class speed**

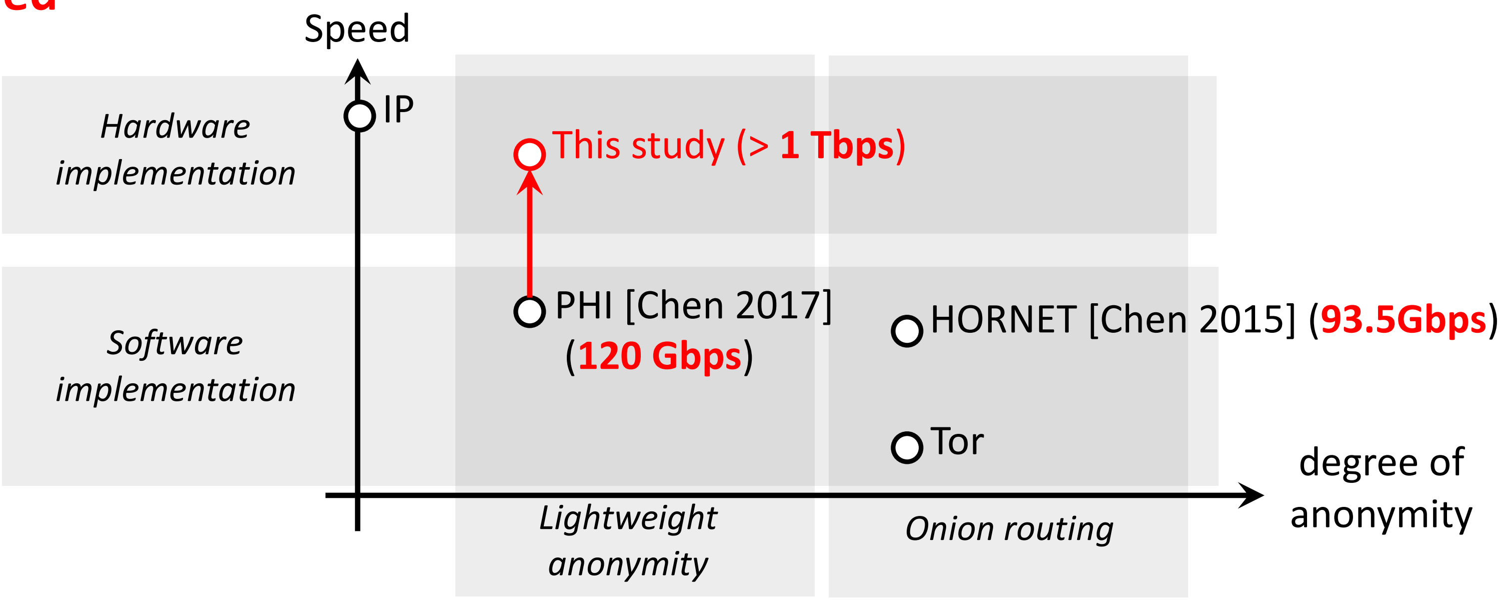

[Chen 2017] Chen, Chen, and Adrian Perrig. "PHI: Path-hidden lightweight anonymity protocol at network layer." *Proc. Priv. Enhancing Technol.* 2017.1 (2017): 100-117.
[Chen 2015] Chen, Chen, et al. "HORNET: High-speed onion routing at the network layer." Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. 2015.

Lightweight Anonymity Protocol (AP)

- Path setup phase
  - Routers create a forwarding state list in the onward trip to the destination
    - Pairs of predecessor and successor information (forwarding states) determined by IP routing
    - Encrypted with routers' secret keys and stored in packet headers
  - Destination replies with completed forwarding state list in the return trip

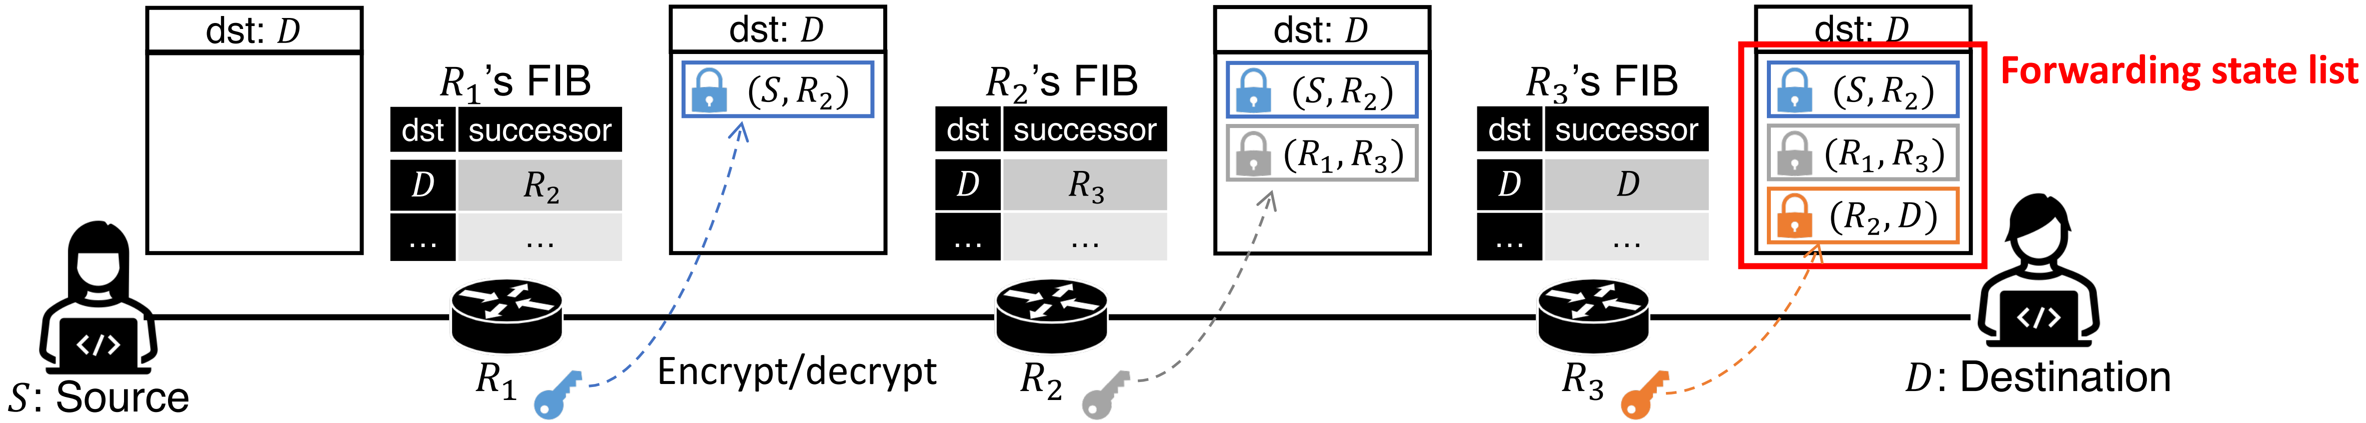

- Data transmission phase
  - Forwarding state list is recorded in packet headers
  - Routers retrieve and decrypt their own forwarding state for forwarding
  - Routers are stateless (states are kept in packet headers)

Design Challenges and Solutions

1. Cryptosystem
   - It must be able to be executed by a **limited instruction set** and a **limited number of stages**
   - Solution: **One-time pad cipher**
     - Suitable for encrypting on switches because encryption and decryption are equivalent
     - Half SipHash (keyed hash function) only uses addition, XOR, and bit rotation

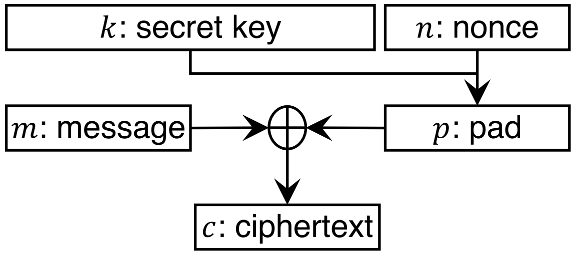

2. Data structure for accommodating a forwarding state list
   - It must satisfy anonymity requirements (e.g., **not leaking length of path**) and must be **lightweight to process** and not require extra computation
   - Solution: Rotatable array
     - The forwarding state list is **rotated after accessed**

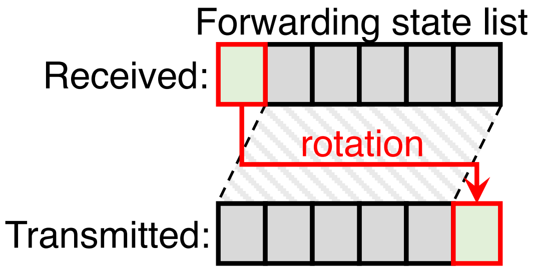

3. IP routing support
   - It must handle **normal IP packets** as well for incremental deployment
   - It must keep **IPv4 full route** (500K entries)
4. Layout of the program on the switch's pipeline
   - Must implement the above functions with satisfying **Tofino's constraints**
   - Must minimize the number of recirculations for **terabit-class speed forwarding**

Layout of the Program on Switch Pipeline

- Layout constraints
  - A switch ASIC has 2 12-stage pipelines (ingress and egress pipeline)
  - Generating a one-time pad requires **32 stages**
  - Looking up FIB requires **12 stages**
  - Egress ports must be determined in the **ingress pipeline**
- Resulting Layout: We implement the functions with **one recirculation**
  - IP packet processing

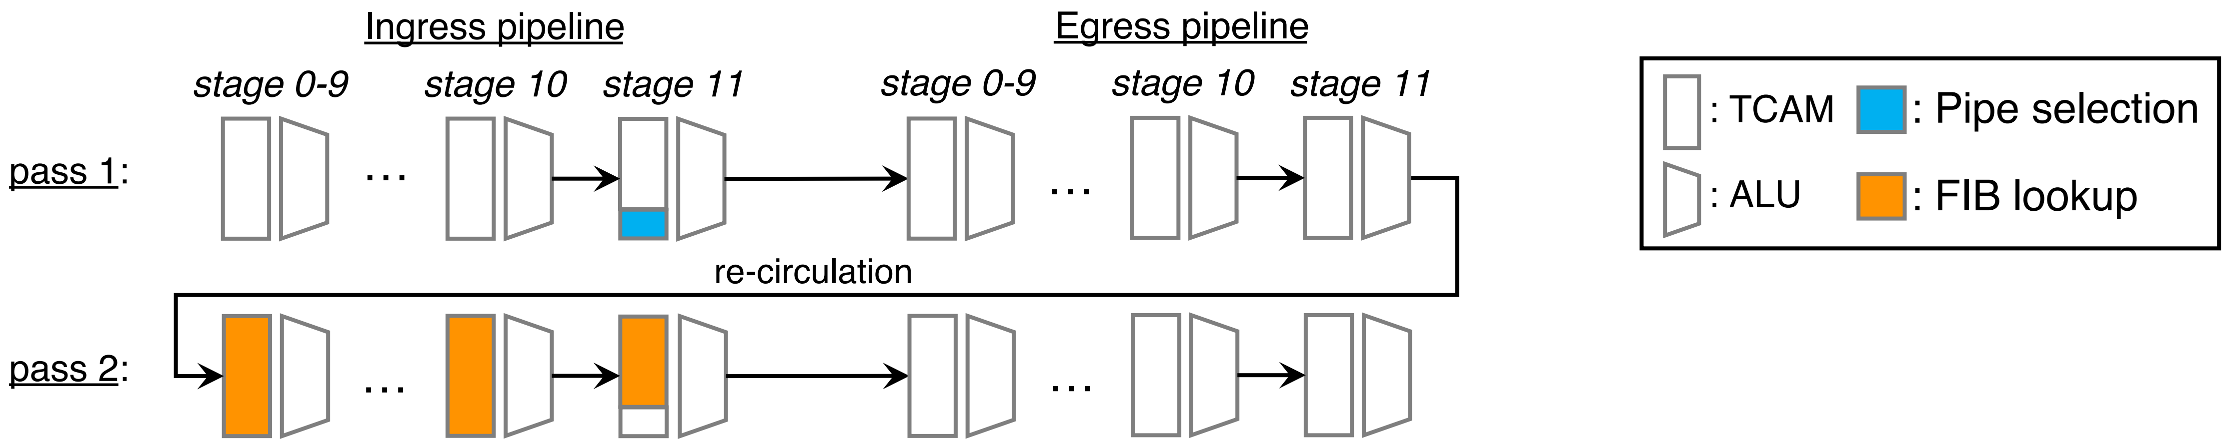

Layout of the Program on Switch Pipeline

- Anonymity packet processing (path setup phase: encryption)

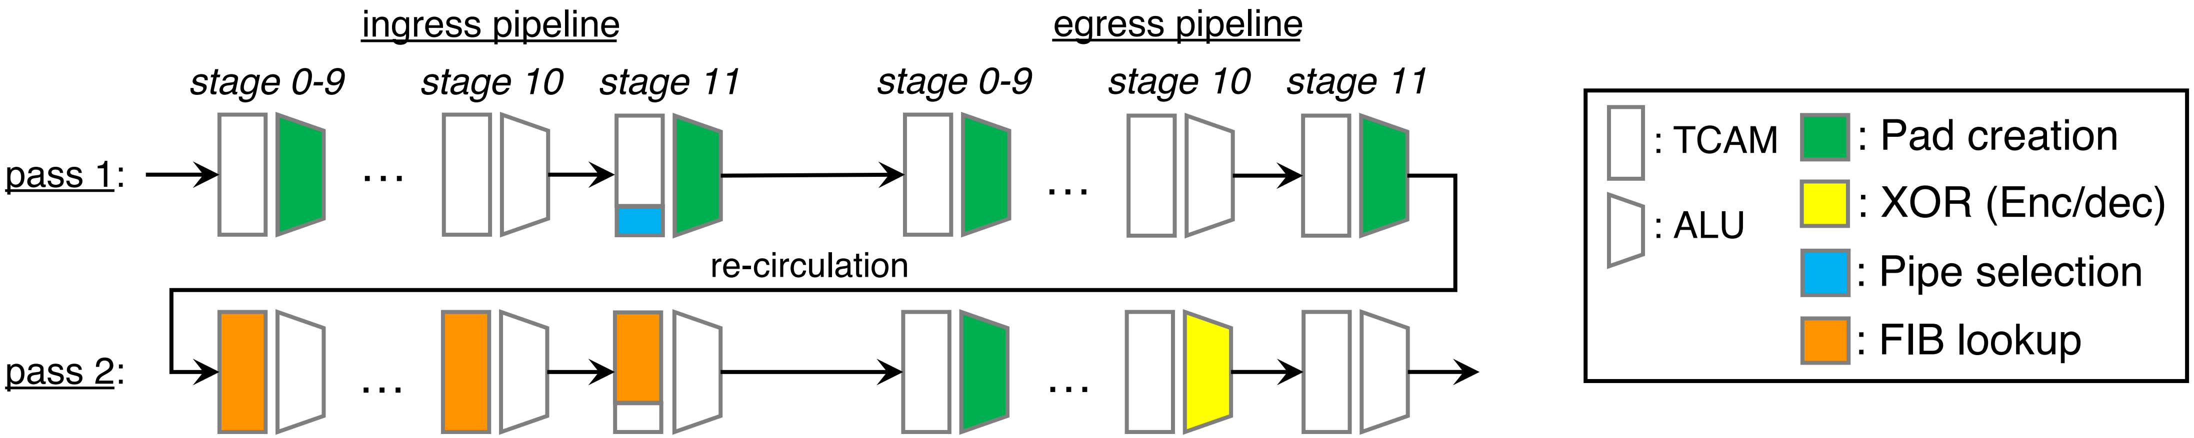

- Anonymity packet processing (data transmission phase: decryption

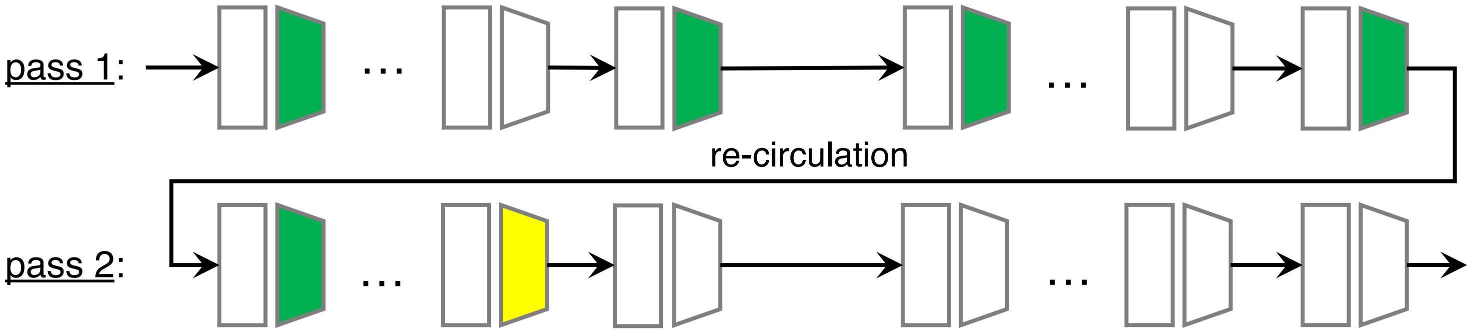

- Integrated pipeline layout

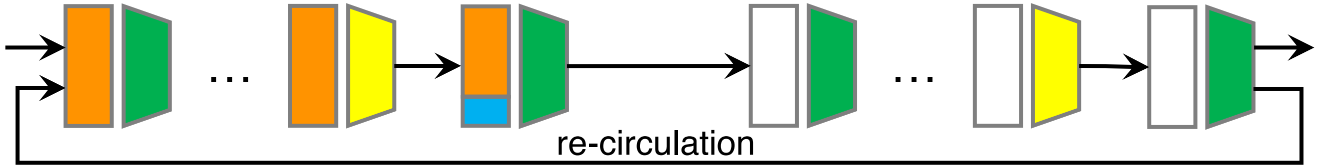

Evaluation and Future Directions

- Evaluation method
  - We implement designed pipeline using a Tofino switch (Wedge 100BF-32X **32 100 Gbps ports**)
  - Ideal forwarding speed is **1.2 Tbps**
    - 200 Gbps traffic is injected and amplified six times inside the switch
  - 2 ports are used for input, 12 are used for recirculation, and 12 are used for output
- Result: 1.16 Tbps forwarding speed

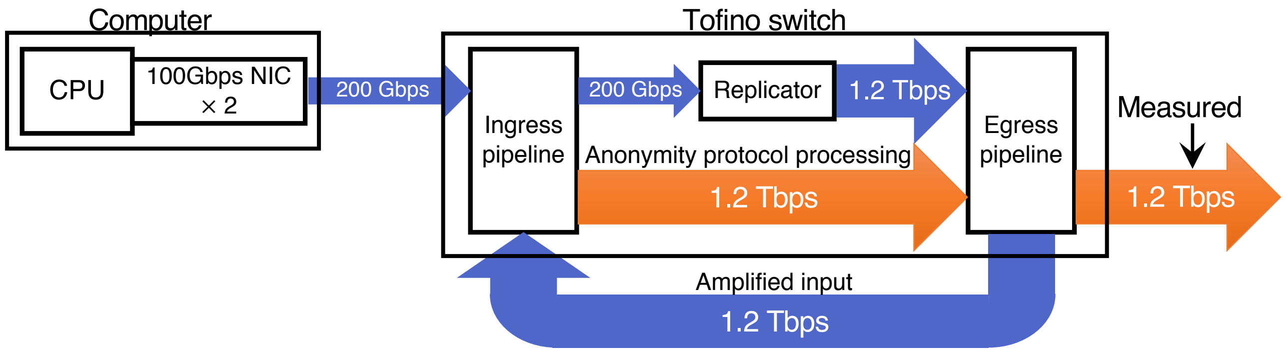

Experiment setup

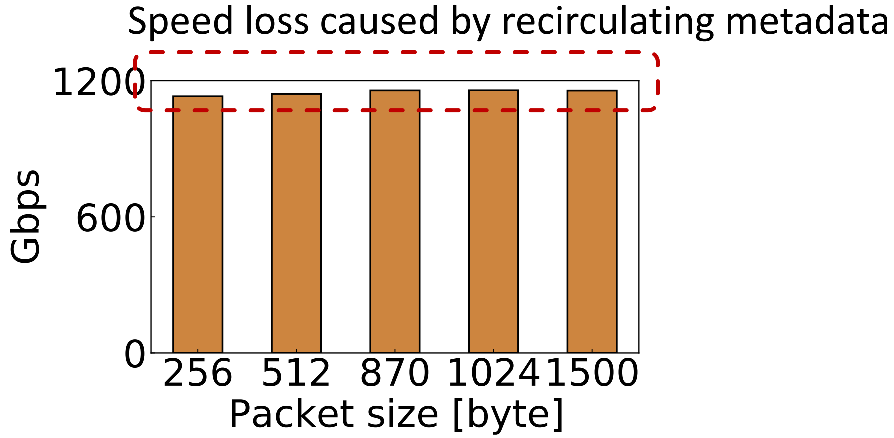

Forwarding speed result

- Ongoing research
  - Message authentication code (MAC) generation/verification **against attacks on path integrity**
  - **Stronger cryptosystem** like one-time pad with ShipHash instead of Half SipHash
- **Future research plan:**
  - **Smart NICs and switches for collaborative mitigation against active attacks**

Collaborative Smart NIC/Switch Secure and High-speed Anonymity Protocol

- Background: Anonymity protocols are **weak against active attacks**
  - Attackers can identify that a victim uses a certain router for anonymous communication
    - Injecting traffic in special patterns (like periodic bursts) toward the victim
    - Probing the perturbation caused by the injected traffic
- Approach: Smart NICs and switches collaboratively monitor and prevent the active attacks
  - Programmable switches
    - Coarse-grained analysis of active attacks
    - Traffic morphing to mitigate the attack
  - Smart NICs
    - High-speed and fine-grained monitoring of active attacks against anonymity protocol

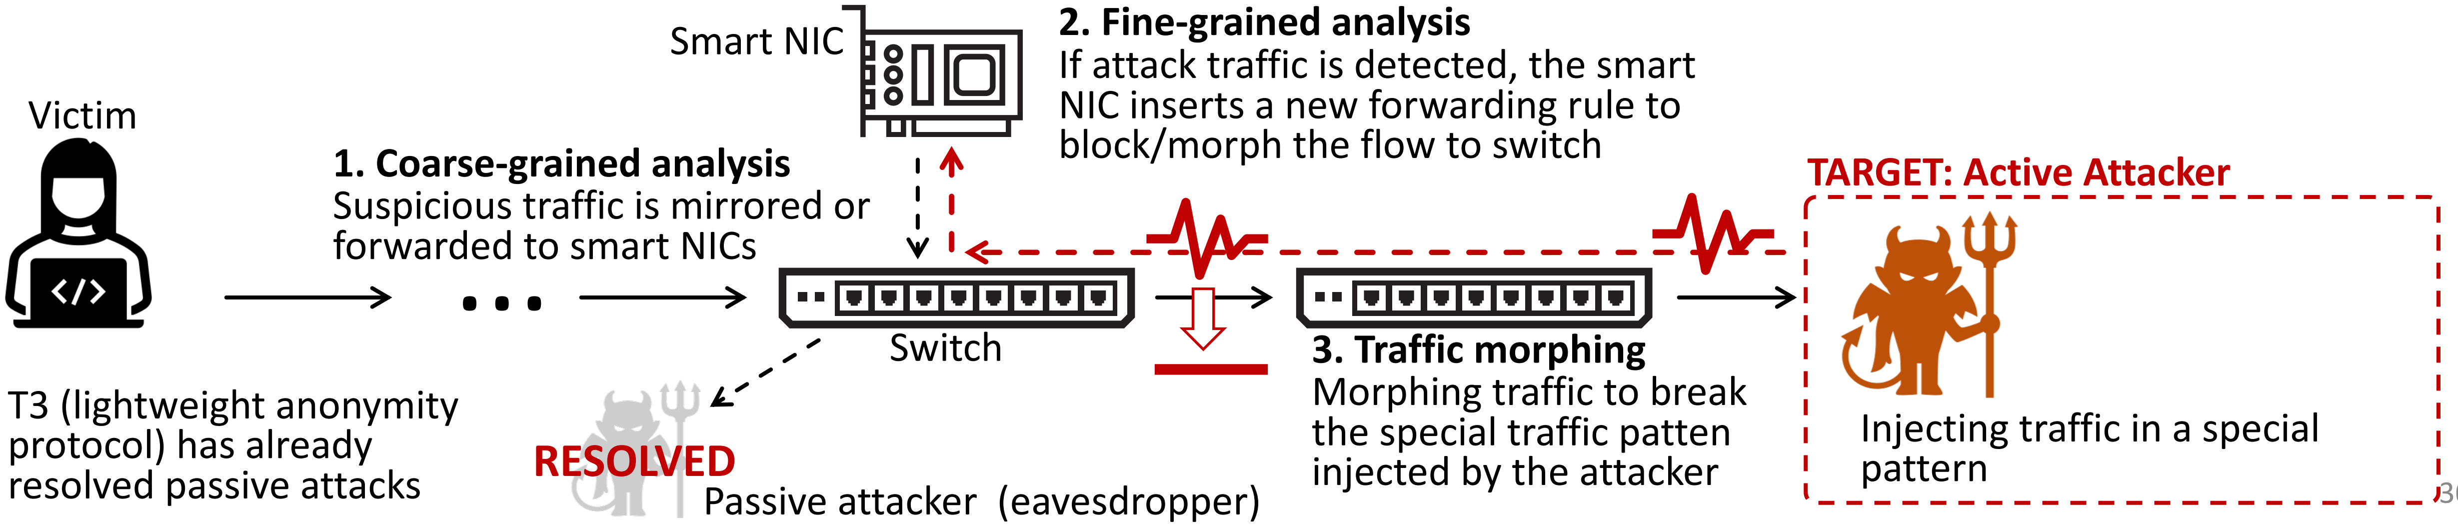

Dissemination of Project Achievements: Summary

- We are designing the architecture to publish a joint paper
- Publications so far:

1. Kentaro Kita, Junji Takemasa, Yuki Koizumi, Toru Hasegawa, “Secure Middlebox Channel over TLS and its Resiliency against Middlebox Compromise,” in Proceedings of IEEE INFOCOM 2023, May 2023.
2. Cuidi Wei, Ahan Kak, Nakjung Choi, and Timothy Wood, “Towards a Scalable 5G RAN Central Unit,” in Proceedings IEEE INFOCOM 2023 Workshop on Next-generation Open and Programmable Radio Access Networks (NG-OPERA), May 2023
3. Ryu Watanabe, Ayumu Kubota, and Jun Kurihara, "Application of Generalized Deduplication Techniques in Edge Computing Environments," to Appear in Proceedings of AINA 2023 (M2EC-2023), Juiz de Fora, Brazil, Mar. 2023.
4. Yutaro Yoshinaka, Junji Takemasa, Yuki Koizumi, Toru Hasegawa, “Design and analysis of lightweight anonymity protocol for host- and AS-level anonymity ,” Computer Networks, Volume 222, 109559-109559, Feb. 2023.
5. S. Panda, K. K. Ramakrishnan and L. N. Bhuyan, "Synergy: A SmartNIC Accelerated 5G Dataplane and Monitor for Mobility Prediction," *2022 IEEE 30th International Conference on Network Protocols (ICNP)*, Lexington, KY, USA, 2022.