Minos : A Lightweight and Dynamic Defense against Traffic Analysis in Programmable Data Planes

Zihao Wang, Qing Li, Guorui Xie, Dan Zhao, Kejun Li,
Zhuochen Fan, Lianbo Ma, Yong Jiang

1 Background

- **Encrypted traffic analysis attacks capture user traffic by its 5-tuple (Source IP, Destination IP, Source port, Destination port, Protocol) to infer users' intentions.**
  - Website Fingerprinting
  - Encrypted DNS Fingerprinting
  - IoT Fingerprinting
  - Video Stream Fingerprinting
  - etc.

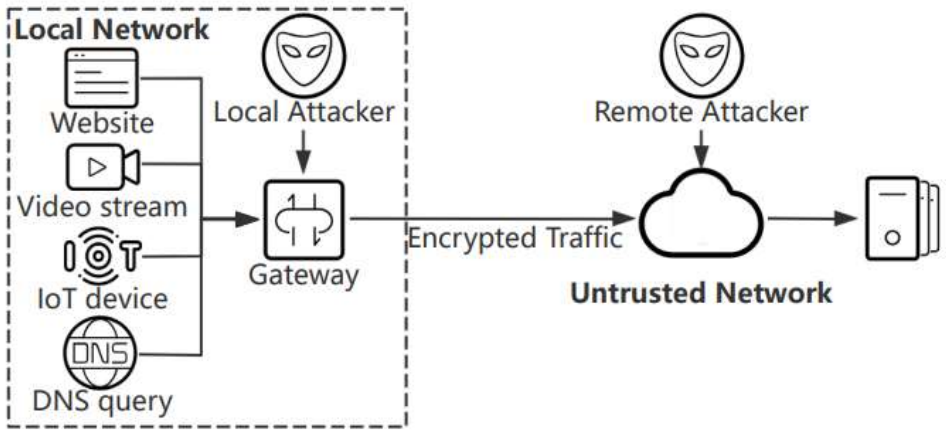

Figure 1: An example of encrypted traffic analysis.

Posing significant threats to user privacy!

1 Background

- **Proxy-based defenses: obfuscating 5-tuples**
  - ➢ Socks-based proxies: Shadowsocks, V2ray.
  - ➢ Secure gateways: IPsec gateways, MACsec gateways.
- **Pros and cons:**
  - ✓ Lightweight and Scalable.
  - × Weak against advanced attacks, lacks Traffic Anonymity.

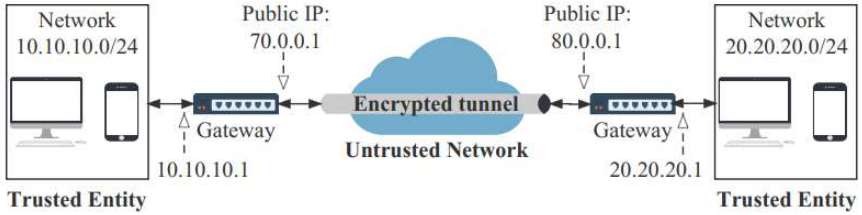

Figure 2: An example of IPsec gateway.

- **Traffic Morphing-based defenses: obfuscating traffic traces**
  - ➢ Insert, delay, combine, or split packets.
- **Pros and cons:**
  - ✓ Provides Traffic Anonymity.
  - × Low throughput, lacks Identity Anonymity.

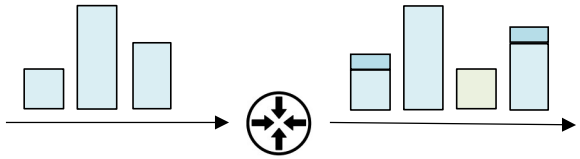

An example of Traffic Morphing-based defenses

1 Background

- **Programmable Switches provides the opportunity to enhance the defense mechanisms**

✓ Lightweight and Scalable. High throughput (100Gbps line rate).

✓ Data plane programmability to support customized functions.

× Limited stages and RAM space.

× Limited computations.

- **Previous programmable switches-based defenses**

➢ Cannot reach line rate due to bandwidth overhead.

➢ Relies on IPsec or MACsec gateways to provide Identity Anonymity.

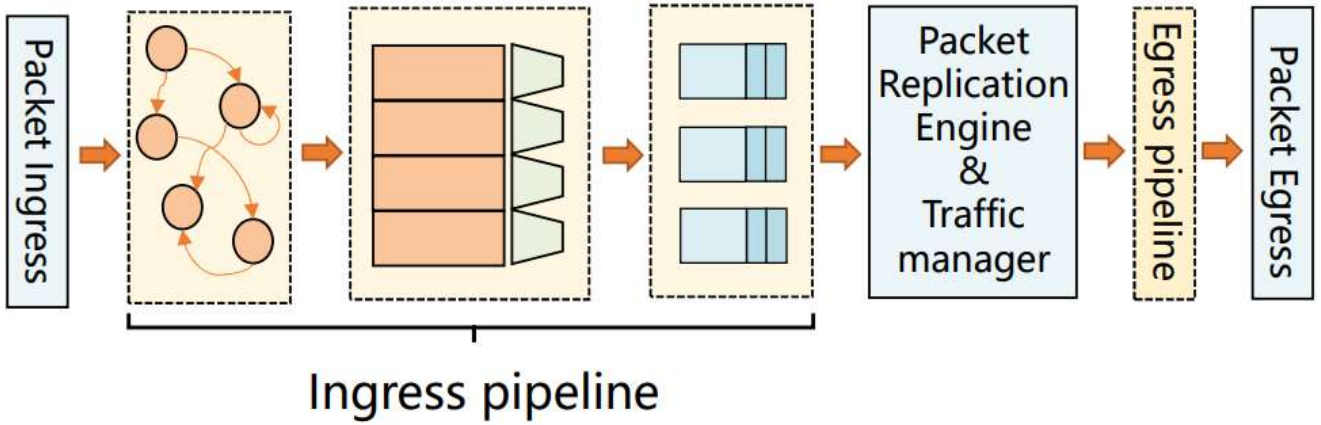

1 Background

- **Previous defenses: high bandwidth overhead, lacks scalability or fails to provide Anonymity**

- **Minos: a line rate traffic analysis defense scheme based on programmable switches.**
  - Work as gateway switches to establish end to end tunnels.
  - Lightweight.
  - Scalable.
  - Provide both Identity Anonymity and Traffic Anonymity.

Scheme	Lightweight	Scalable	Anonymity	
			Identity Anonymity	Traffic Anonymity
Proxy-based Defenses	✓	✓	✓	×
Traffic morphing-based Defenses	×	×	×	✓
Ditto [28]	×	✓	×	✓
Minos	✓	✓	✓	✓

Comparison of the existing solutions.

2 Overview

- **Minos contains three modules:**
  - The **Proxy Module** acts as a gateway switch that replaces the source IP of each packet with the gateway IP and sends packets at line rate.

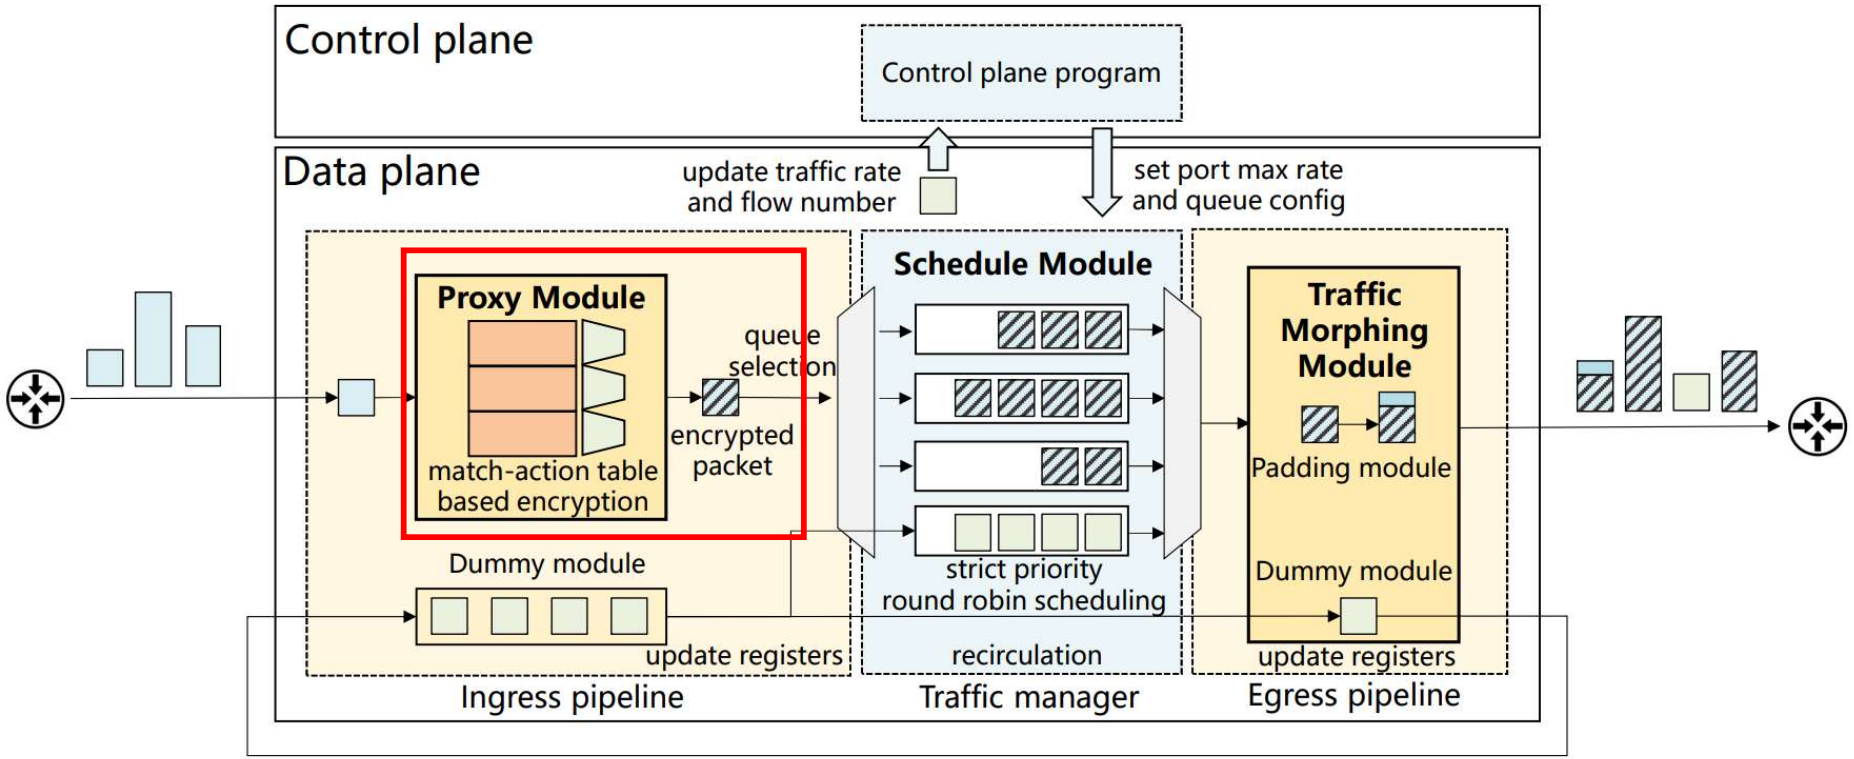

2 Overview

- **Minos contains three modules:**
  - The **Schedule Module** aggregates different flows to the same destination to obfuscate the original traffic with minimal dummy packets.

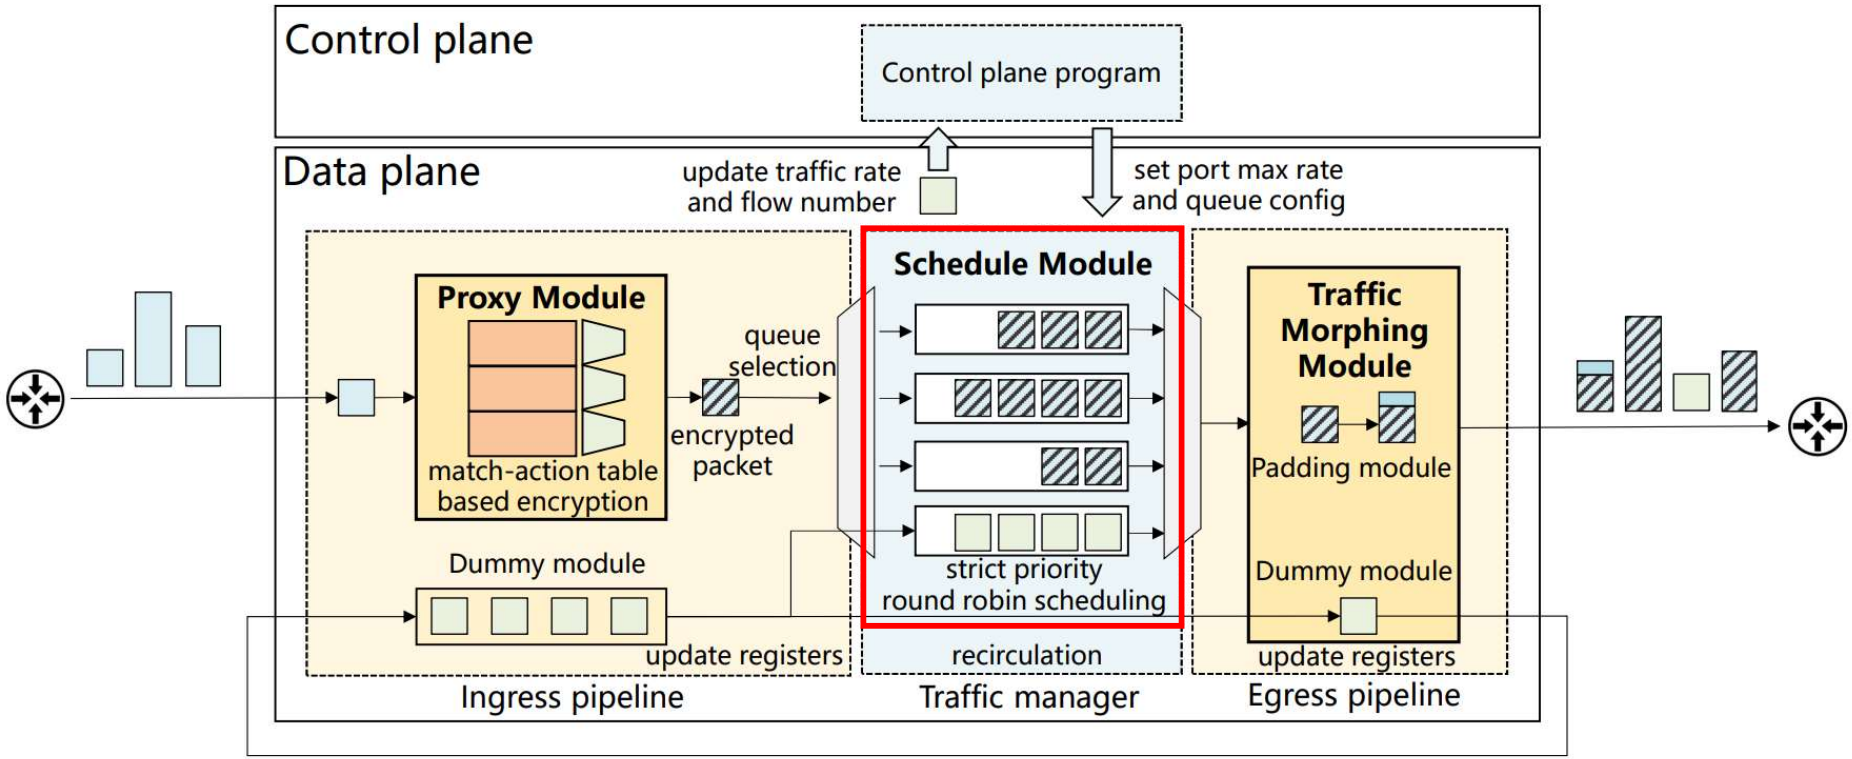

2 Overview

- **Minos contains three modules:**
  - The **Traffic Morphing Module** provides primitives to implement traffic morphing-based defenses, consisting of the *Dummy Module* and the *Padding Module*.

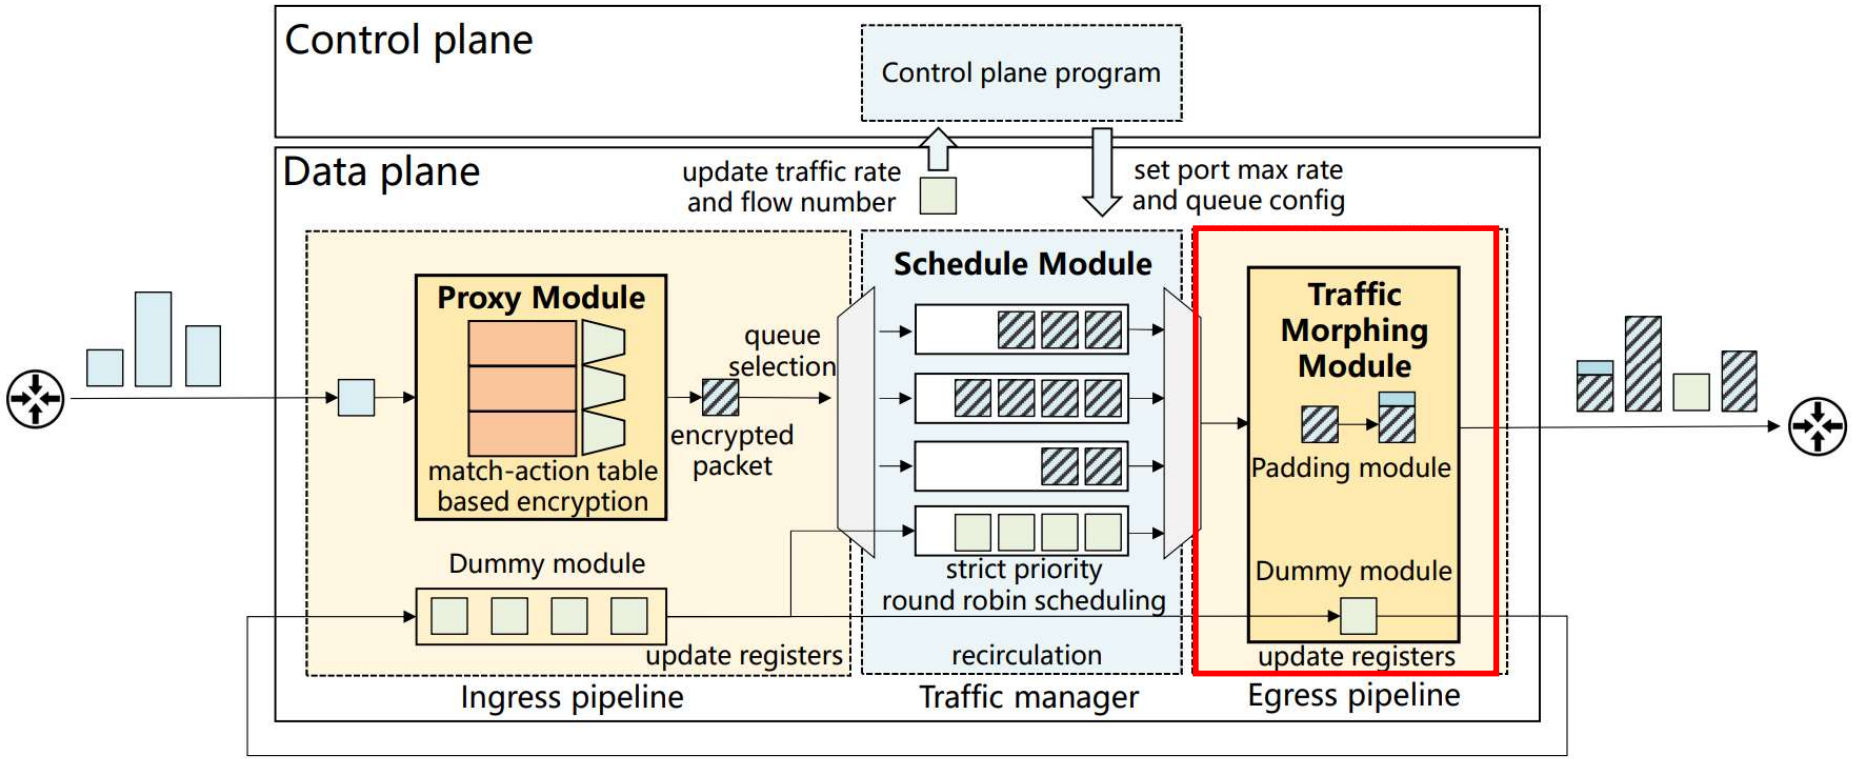

2 System Design-Proxy Module

- **Challenges:**
  - Limited computations: Modulo Operation and division are not supported.
  - Limited stages: A standard Tofino pipeline has only 12 stages.
  - Limited tables: The storage of programmable switches are limited.
- **Minos Proxy Module design:**
  - **Encryption round compression:** Compress four steps in each encryption round into one stage.
  - **Memory Consumption Reduction:** Apply the same set of keys and tables for both encryption and decryption decryption.

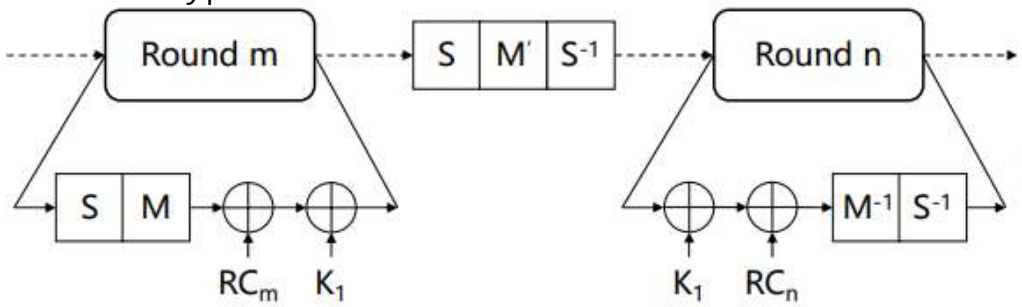

Simplified encryption process of PRINCE

$$D(k_0||k_0'||k_1)(message) = E(k_0'||k_0||k_1 \oplus \alpha)(message)$$

α-reflection property

[1] Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, et al. Prince–a low-latency block cipher for pervasive computing applications. In Advances in Cryptology–ASIACRYPT 2012: 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings 18, pages 208–225. Springer, 2012

2 System Design-Schedule Module

- **Drawbacks of IPsec gateways:**
  - sends packets out in a first-in-first-out manner
- **Minos solution:**
  - Performs flow interleaving to obfuscate the original traffic
- **Challenges:**
  - Prevent packet disorder
  - Support large amount of users
  - Minimize the traffic morphing bandwidth overhead

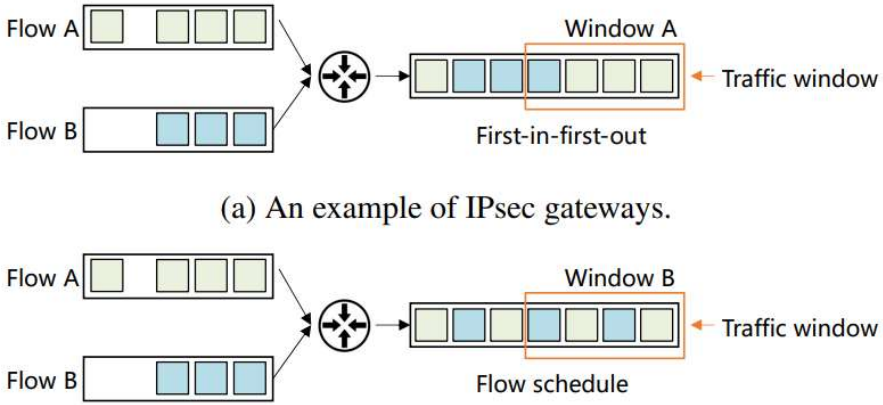

(a) An example of IPsec gateways.

(b) An example of Minos flow schedule.

Figure 6: A comparison of IPsec gateway and Minos.

2 System Design-Schedule Module

- **Minos Schedule Module Design: a Dynamic Flow Scheduling Method.**
  - Stores flow-level information.
  - Performs flow interleaving with round-robin queues.
  - Recording the active flow number in each encryption tunnel to enable Traffic Morphing.

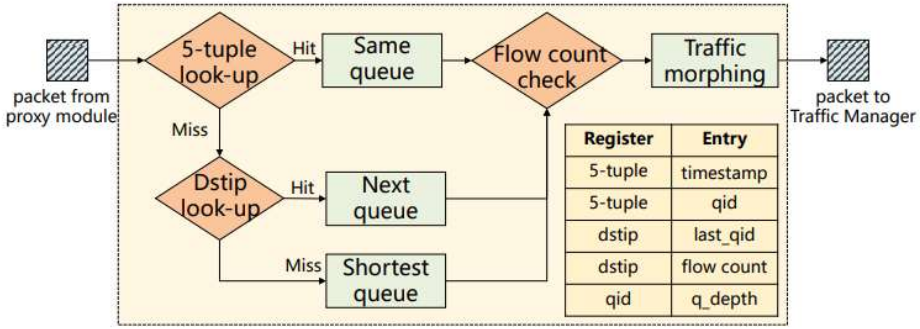

Figure 7: Schedule algorithm of Minos.

- **Key points:**
  - Minos holds multiple encrypted tunnels with different ends, which operates another Minos switch to perform decryption.
  - Attackers can eavesdrop on each encrypted tunnel with the destination IP address.
  - When few flows are transmitting, Minos will enable Traffic Morphing to enhance the defense.

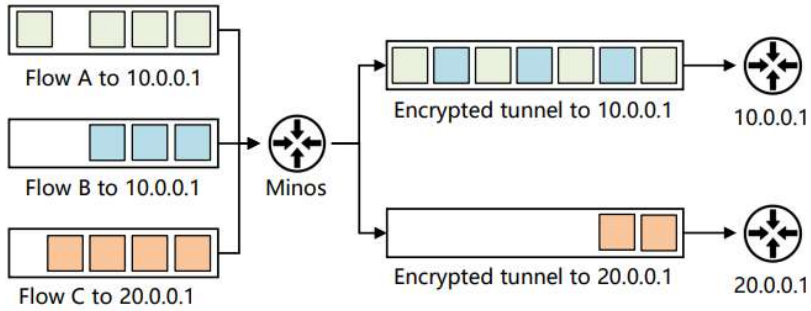

Figure 8: Example of multiple encrypted tunnels.

2 System Design-Traffic Morphing Module

- **Challenges of the Dummy Module:**
  - Dummy packets must be generated in advance and can not be sent in real time.

- **Previous researches fails to provide a lightweight solution:**
  - IMAP[NSDI22]: fill the pipeline entirely with dummy packets
  - Ditto[NDSS22]: buffer the packets in the first round.

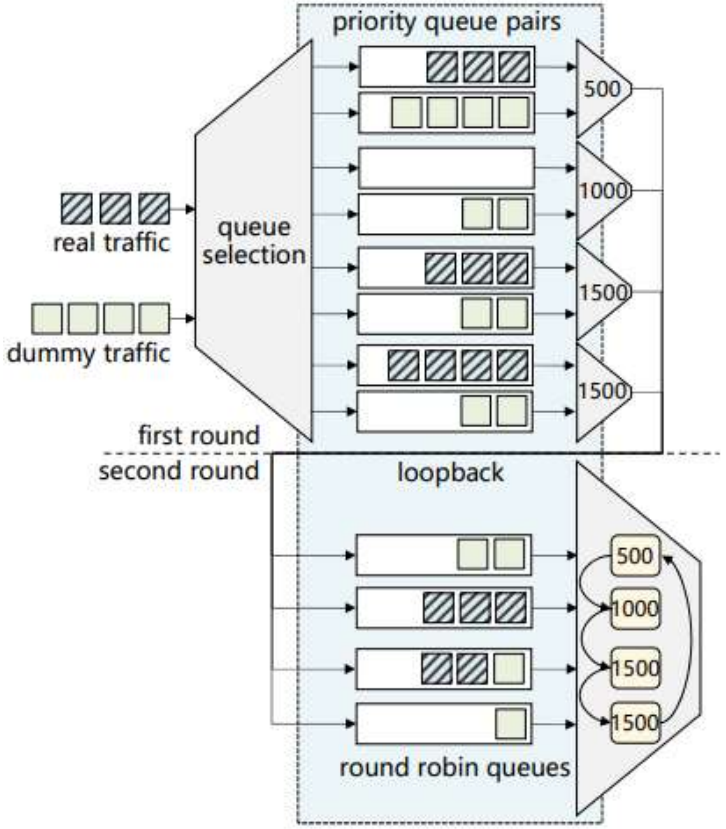

Ditto Traffic Manager

[1] Guanyu Li, Menghao Zhang, Cheng Guo, Han Bao, Mingwei Xu, Hongxin Hu, and Fenghua Li. {IMap}: Fast and scalable {In-Network} scanning with programmable switches. In 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22), pages 667–681, 2022.
[2] Roland Meier, Vincent Lenders, and Laurent Vanbever. Ditto: Wan traffic obfuscation at line rate. In NDSS Symposium, 2022.

2 System Design-Traffic Morphing Module

- **Minos Dummy Module Design: Priority Queue based Dummy Packet Scheduling.**
  - Buffer dummy packets in the high priority queue.
  - Send dummy packets with Traffic Manager by pausing and resuming queue.
  - Controls sending rate with the following formula: n: the number of rounds in the resume period
    i: the interval between resume and pause instructions
    r & R: output rate of dummy queue and normal queue
    d: the delay of instructions

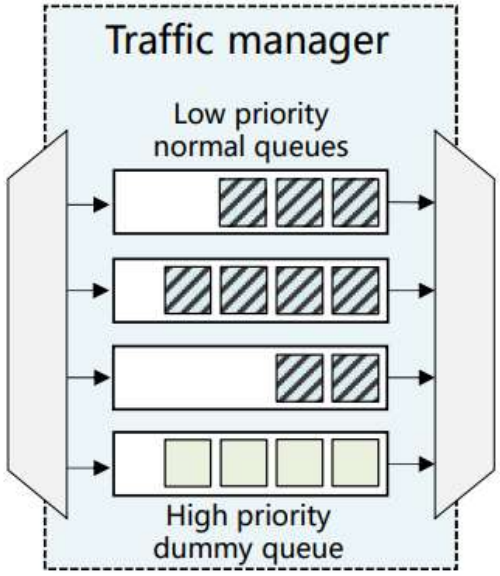

Minos Traffic Manager

$$Dummy\_rate = \left(\frac{1}{n} - d - i\right) * \frac{r}{R}$$

Minos Dummy rate formula

3 Hardware Prototype Evaluation-Implementation

- **Implementation:**
  - We implement a hardware prototype of Minos on two Barefoot Tofino1 switches.

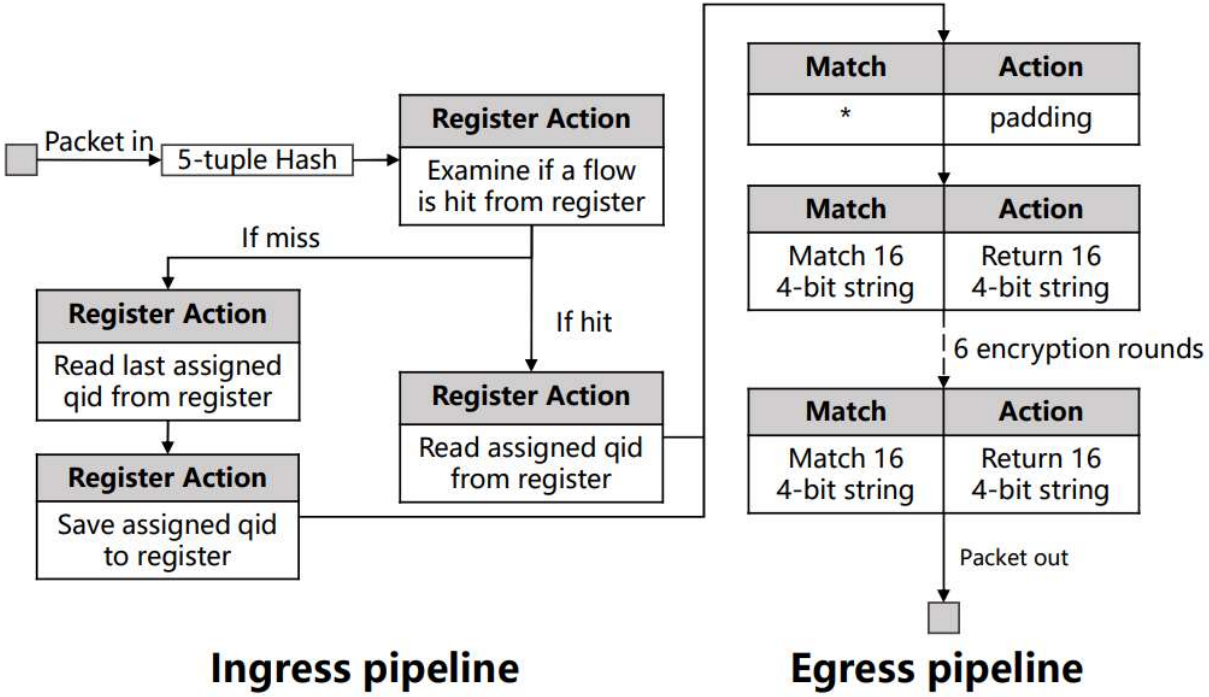

Minos flow table

Module name	Stages	VLIWs	ALUs	TCAM	SRAM
Proxy module	10	6.15%	0%	0%	55.96%
Minos	12	7.81%	0%	0%	59.06%

Hardware Resource utilization.

3 Hardware Prototype Evaluation

- **Evaluation Setup:**
  - Two Minos switches to establish an end-to-end tunnel.
  - A traffic generator to inject packets and record output traffic.
  - Datasets:
    - Website Fingerprinting dataset from [1]
    - IoT Fingerprinting dataset from [2]
    - An even mixture of both dataset
    - TCP packets of 128B
- **Proxy Module Evaluation:**
  - The throughput of Minos Proxy Module with different datasets.
  - All above 94Gbps

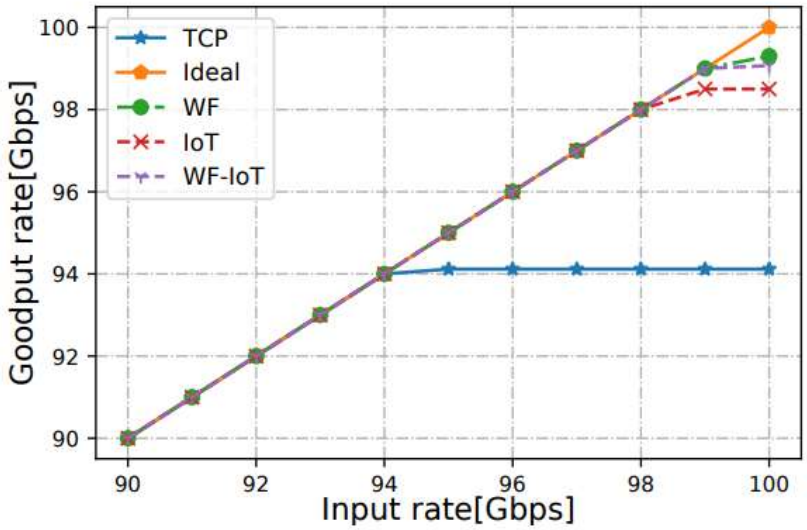

Figure 12: Throughput of Minos proxy module.

[1] Vera Rimmer, Davy Preuveneers, Marc Juarez, Tom Van Goethem, and Wouter Joosen. Automated website fingerprinting through deep learning. arXiv preprint arXiv:1708.06376, 2017.
[2] Jingjing Ren, Daniel J Dubois, David Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. Information exposure from consumer iot devices: A multidimensional, network-informed measurement approach. In Proceedings of the Internet Measurement Conference, pages 267–279, 2019.

3 Hardware Prototype Evaluation

- **Schedule Module Evaluation:**
  - Latency brought by the Schedule Module is insignificant.

Table 3: Arrival time of 1000th packet(s).

Origin	Limited output rate	Two flows	Three flows	Four flows
1	1.095861	1.102307	1.106605	1.108754

- **Dummy Module Evaluation:**
  - Examine how can we adjust dummy packet rate.

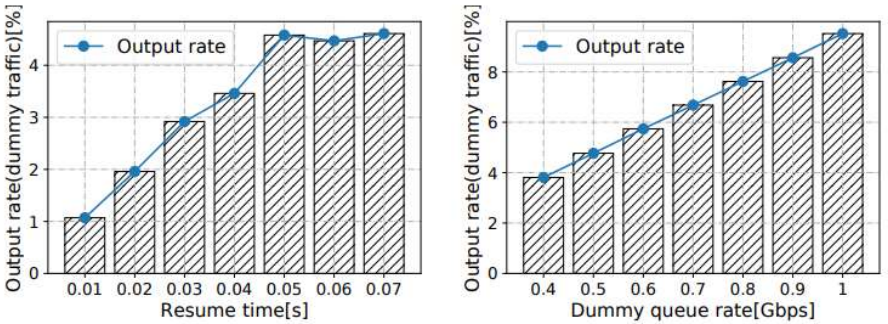

(a) Different intervals (b) Different output rates

Figure 13: Evaluation of the dummy module

- **Overall Evaluation:**
  - Overall goodput rate above 95%
  - Overall latency overhead less than 2.4%

Table 4: Overhead of Minos

Flow type	arrival time of 500,000th packet(s)	Goodput rate(%)
Base	5	100%
One flow	5.02	95.3%
Multiple flows	5.12	99.2%

3 Hardware Prototype Evaluation

- **Comparison with SOTA[1] work:**
  - Better overall throughput with different datasets.
  - Higher TCP throughput.
  - Lower UDP loss.
  - Faster website loading.

- **To conclude, Minos outperforms ditto in all aspects related to throughput and overhead.**

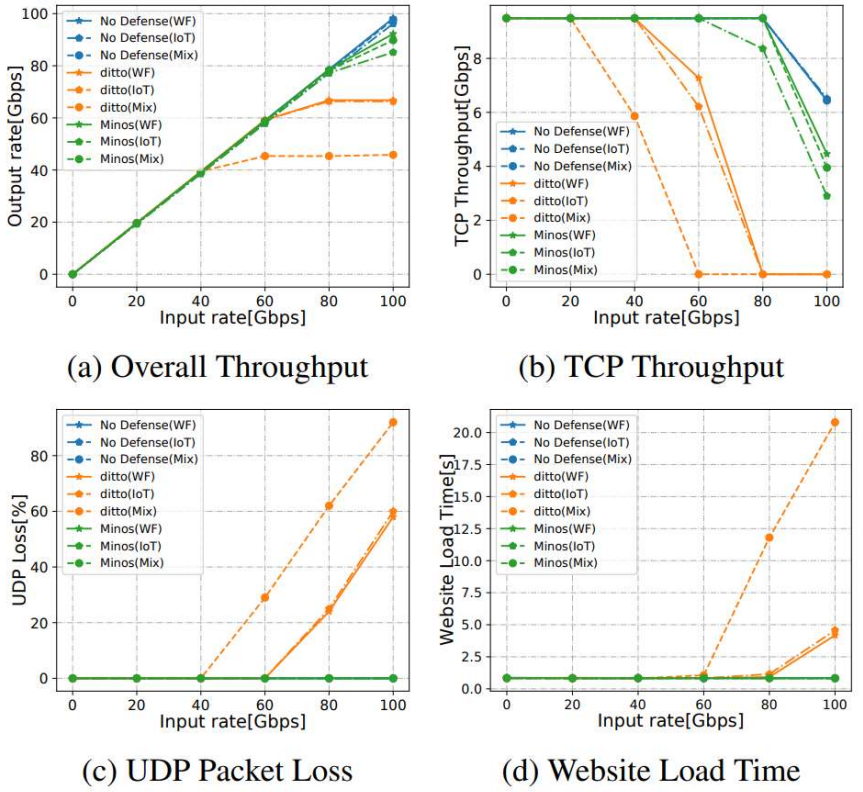

(a) Overall Throughput (b) TCP Throughput

(c) UDP Packet Loss (d) Website Load Time

Figure 14: Evaluation results in comparison with Ditto.

[1] Roland Meier, Vincent Lenders, and Laurent Vanbever. Ditto: Wan traffic obfuscation at line rate. In NDSS Symposium, 2022.

4 Use case: Website Fingerprinting

- **Evaluation setup:**
  - Website Fingerprinting dataset from [1]
  - Attacks from top security conferences: kNN[2], CUMUL[3], kFP[4], DF[5].
  - Mix 1-N flows from different websites.
- **Defense design:**
  - Multiple flows: Flow interleaving without traffic morphing.
  - Few flows: Flow interleaving with traffic morphing which randomly insert dummy packets in first 500 or 1000 packets of each flow.
- **Evaluation Metrics:**
  - Accuracy Metrics: if the ML model returns any right label of the original mixing flows.
  - TopN Metrics: how many right labels are returned, each is counted as 1/N.

[1] Vera Rimmer, Davy Preuveneers, Marc Juarez, Tom Van Goethem, and Wouter Joosen. Automated website fingerprinting through deep learning. arXiv preprint arXiv:1708.06376, 2017.
[2]Tao Wang, Xiang Cai, Rishab Nithyanand, Rob Johnson, and Ian Goldberg. Effective attacks and provable defenses for website fingerprinting. In 23rd USENIX Security Symposium (USENIX Security 14), pages 143– 157, 2014.
[3] Andriy Panchenko, Fabian Lanze, Jan Pennekamp, Thomas Engel, Andreas Zinnen, Martin Henze, and Klaus Wehrle. Website fingerprinting at internet scale. In NDSS, 2016.
[4] Jamie Hayes and George Danezis. k-fingerprinting: A robust scalable website fingerprinting technique. In 25th USENIX Security Symposium (USENIX Security 16), pages 1187–1203, 2016.
[5] Payap Sirinam, Mohsen Imani, Marc Juarez, and Matthew Wright. Deep fingerprinting: Undermining website fingerprinting defenses with deep learning. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, pages 1928– 1943, 2018.

4 Use case: Website Fingerprinting

- **Defense performances:**
  - Mixing 4 flows is enough to decrease the accuracy of each attacks to below 20%.
  - The lightweight defense mechanism can reduce the accuracy of each attacks to around 40%.

Table 5: **Defense performances**

Defense	Accuracy%				Precision%				F1			
	kNN	CUMUL	kFP	DF	kNN	CUMUL	kFP	DF	kNN	CUMUL	kFP	DF
No defense	84.62	91.47	91.68	96.9	83.46	90.72	91.56	96.75	0.831	0.9091	0.9147	0.9676
Tamaraw	2.55	4.42	4.99	1.47	3.57	8.18	6.43	5.32	0.0223	0.049	0.0528	0.0206
WTF-PAD	20.38	44.42	57.28	81.06	18.43	44.14	56.11	78.86	0.1766	0.4305	0.5491	0.7857
Minos-500	5.63	39.82	31.73	7.06	7.96	55.68	38.96	6.7	0.05	0.4	0.28	0.05
Minos-1000	7.38	38.73	32.66	6.67	8.49	52.61	38.31	6.7	0.06	0.38	0.28	0.045

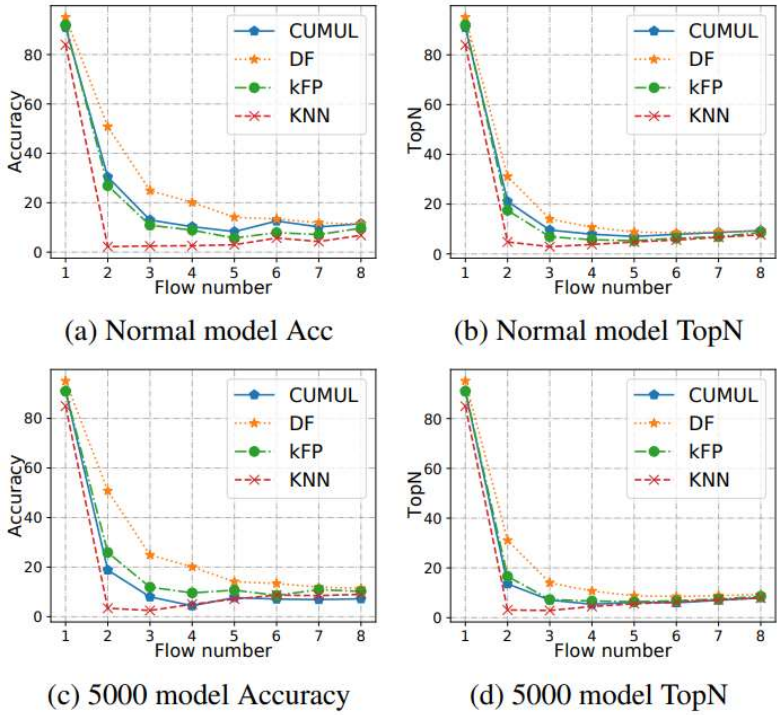

(a) Normal model Acc (b) Normal model TopN

(c) 5000 model Accuracy (d) 5000 model TopN

Figure 15: Evaluation of flow mixture

4 Use case: Website Fingerprinting

- **Defense throughput:**
  - Multiple flow Scenario: insignificant bandwidth and latency overhead.
  - Single flow Scenario: insignificant latency overhead and low bandwidth overhead.

Table 6: Overhead of each defense mechanism

Defense	Parameters	Overhead	
		Latency(%)	Bandwidth(%)
Tamaraw	$\rho_{out} = 0.04, \rho_{in} = 0.012, L = 50$	14.23	143.82
WTF-PAD	Normal rcv	0	60
Minos-500	$\Omega = 0.6$, window=500	0	6
Minos-1000	$\Omega = 0.6$, window=1000	0	12

- **To conclude, Minos reduces the accuracy of each attacks with limited overhead and achieves line rate packet transmitting.**

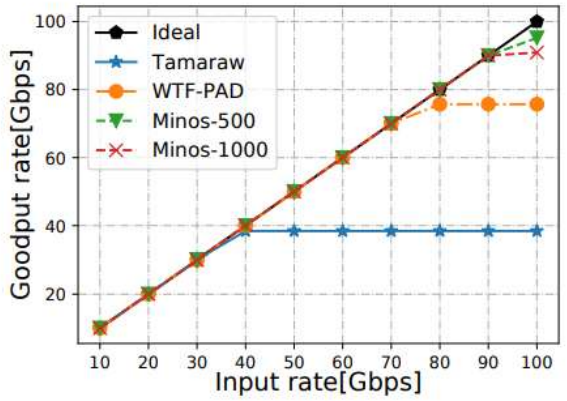

Figure 16: Throughput of each defense.

5 Summary

- **Minos: a line rate traffic analysis defense scheme based on programmable switches.**
- **Three key Modules:**
  - A **Proxy Module** that achieves **line rate encryption** within constrained computations and resources by the *Encryption Round Compression method*.
  - A **Schedule Module** that handles concurrent user traffic in a dynamic and scalable manner with a *Dynamic Flow Scheduling Method*.
  - A **Traffic Morphing Module** that realizes real-time dummy packet insertion with *Priority Queue based Dummy Packet Scheduling*.
- **The Evaluation result proves that Minos is lightweight, scalable and provides both Identity Anonymity and Traffic Anonymity.**

Thanks for your listening!