ETH zürich

Networked Systems
ETH Zürich — seit 2015

D-ITET

Prof. Dr. Laurent Vanbever

February 2025

Exam: Advanced Topics in Communication Networks

10 February 2025, 08:30–11:00, Room HIL F 75

- Write **legibly** your ETH student number (legi number) below on this front page.
- **Do not write your name** or use a stamp with your name on it.
- **TRIPLE-check that your legi number is correct!**
  You will not be graded if you make a mistake when writing your number.
- Put your legitimation card (legi) on the most accessible corner of your desk.
  Make sure that the side containing your name and **student number is visible**.

- Verify that you have received all task sheets (Pages 1–32).
- **Do not separate** the task sheets. We will collect the exams after you left the room.
- Write your answers directly on the task sheets.
- All answers fit within the allocated space—often in much less.
- If you need more space, use the **extra sheets** at the end of the exam. **Indicate the task** in the corresponding field, and add a **"see Extra Sheet X"** note in the original task space.
- Read each task completely before you start solving it.

- Answer in **English**.
- **Write clearly** in blue or black ink (not red) using a **pen**, not a pencil.
- **Cancel** invalid parts of your solutions **clearly** (e.g., by crossing them out).
- At the end of the exam, **place the exam face up** on the most accessible corner of your desk. Then collect all your belongings and exit the room according to the given instructions.

- No written material or calculator are allowed.
- It is not required to score all points to get the best mark.

Student legi nr.:

Do not write in the table below (used by corrector only):

Task	Points
Advanced routing	/26
Programmable data planes	/15
Network verification	/21
Network measurements	/16
Network security	/8
Transport	/19
Sustainable networking	/15
Total	/120

leave blank

leave blank

Task 1: Advanced Routing **26 Points**

a) Hierarchical routing **(7 Points)**

Consider the route reflection topology composed of five routers depicted in Fig. 1. `RRA` and `RRB` act as route reflectors, while `R1`, `R2`, and `R3` are route reflector clients. The dashed lines indicate iBGP sessions, with single-headed arrows indicating client sessions (they point towards the route reflector) and double-headed arrows indicating normal iBGP sessions (e.g., `R2` is the client of `RRA`, while `R1` is the client of both `RRA` and `RRB`). The network relies upon an IGP whose weights are depicted next to each link.

The router `RRA` receives an eBGP route `r1` for a certain prefix p. This route has local-pref 50 and an AS-PATH length of 2.

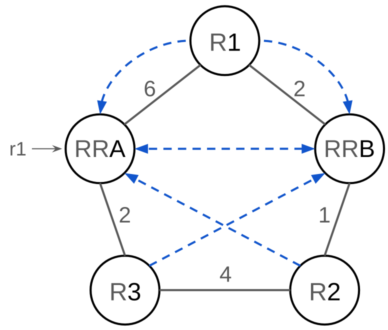

Figure 1: A simple route reflection topology.

(i) Assuming there were no route reflectors, and instead the network was an iBGP full mesh, how many iBGP sessions would there have been? (1 Point)

(ii) Now assume `RRA` and `RRB` act as route reflectors as described. For every router, identify its chosen egress router and the length of the IGP path towards it (write 0 if the router and the egress router are the same). If there is anything suboptimal or incorrect about the paths, explain it. (2 Points)

Router	Egress Router	IGP Weight
R1		
R2		
R3		
RRA		
RRB		

Something sub-optimal?

Something incorrect?

(iii) `R2` has now also received a route (`r2`) towards p. It has local-pref 50 and an AS-PATH length of 5. For every router, identify its chosen egress router and the length of the IGP path towards it (write 0 if the router and the egress router are the same). If there is anything suboptimal or incorrect about the paths, explain it. (2 Points)

Router	Egress Router	IGP Weight
R1		
R2		
R3		
RRA		
RRB		

Something sub-optimal? ____________________________

Something incorrect? ____________________________

(iv) `R1` has now also received a route (`r3`) towards p. It has local-pref 50 and an AS-PATH length of 2. For every router, identify its chosen egress router and the length of the IGP path towards it (write 0 if the router and the egress router are the same). If there is anything suboptimal or incorrect about the paths, explain it. (2 Points)

Router	Egress Router	IGP Weight
R1		
R2		
R3		
RRA		
RRB		

Something sub-optimal? ____________________________

Something incorrect? ____________________________

b) Fast(er) convergence **(9 Points)**

leave blank

Consider the network topology composed of six routers depicted in Fig. 2. **The values assigned to the edges represent bandwidths**.

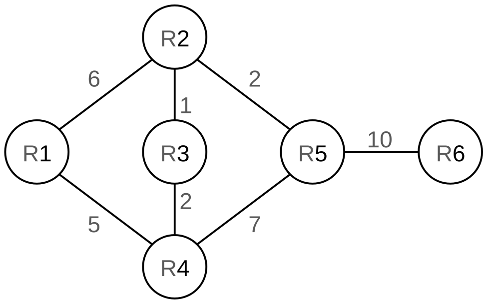

Figure 2: A simple internal topology.

In this task, unless otherwise stated, paths are selected according to the **widest-path** algebra, as seen in the lecture:

$$(\mathbb{R}_0^+ \cup \{\infty\}, \geq, \min, 0).$$

In case of ties, prefer the path with the smallest number of hops.

(i) Is `R4` a per-link Loop-Free Alternate (LFA) for `R1` considering the failure of the link (`R1`, `R2`)? Briefly explain. (4 Points)

Can you modify the bandwidth of a single link to change whether or not `R4` is an LFA? That is:

- If you answered that it **is** an LFA, can you change a single link's bandwidth such that it no longer is an LFA?
- If you answered that it **is not** an LFA, can you change a single link's bandwidth such that it becomes an LFA?

If a change is possible, identify such a link and the new bandwidth it should have; if no such change is possible, explain why.

(ii) In the original topology, is `R5` a **remote** LFA for `R1` considering the failure of the link (`R1`, `R2`)? Briefly explain. (2 Points)

(iii) Assume that you are a router in a **random** network. Your knowledge of the graph is limited: you know your neighbors and the bandwidths of the links connecting you to them; however, you know nothing beyond that. Suppose that one of the links connecting you to one of your neighbors fails: how would you choose a router from your remaining set of neighbors in order to *maximize* the likelihood of it being an LFA? Briefly explain. (2 Points)

(iv) Briefly explain how your previous answer would change if the network used **shortest path** routing instead of widest path routing. (1 Point)

c) Distributed prefix filtering and aggregation **(10 Points)**

leave blank

Consider the DRAGON algorithm as explained in class, applied to commercial relationships: it guarantees that if a route towards a child prefix q is filtered and aggregated into the route towards the parent prefix p, the new route towards q will have local-pref greater or equal than the previous one, where the local-pref is defined as customer over peer over provider.

(i) Assume that a router filters q, and the new local-pref is the same as the old one. Is the AS-PATH length of the resulting route guaranteed to be smaller or equal to the length of the original q route? If so, explain why. If not, draw a counter-example in the box below. (2 Points)

(ii) We want to design a variation of DRAGON such that the route attribute is formed by a tuple (local-pref, AS-path length), sorted first by local-pref and then AS-path length. That is, we want to guarantee that the new route will have either: (i) a local-pref strictly greater than the original one; or (ii) a local-pref equal to the original one and an AS-path length smaller than or equal to the original one.

Either prove or disprove the following: We can extend DRAGON to preserve route consistency and guarantee optimality according to this new attribute.

Hint: try thinking algebraically. (3 Points)

Now consider the network topology composed of nine nodes depicted in Fig. 3. A weight is assigned to each edge. Node `R5` advertises prefix p and node `R8` advertises prefix q. q is a child prefix of p.

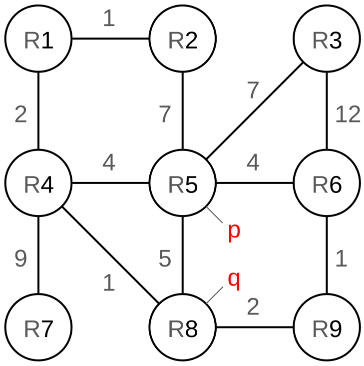

Figure 3: A simple inter-domain topology.

DRAGON can work on different algebras. In the rest of the exercise, consider that the network chooses paths using the **shortest path** algebra.

(iii) List every router that filters prefix q in Fig. 3. (2 Points)

(iv) A router `A` filtering a child prefix q will be able to send packets towards q over a path (through the parent prefix p) that `A` sees as equally cheap or cheaper. However, assuming the weights represent latencies, this does not imply that the packet actually traverses an equally short or shorter path. Draw an example network where the packet ends up taking a longer path after filtering.

Hint: A 3-node network is enough to show this. (3 Points)

leave blank

Task 2: Programmable data planes 15 Points

You are provided with a P4 program designed to detect and mitigate potential SYN flood attacks. In this task, you will analyze this program, explain its behavior, and propose enhancements.

The ingress pipeline code of the P4 program is provided below.

```
control MyIngress(inout headers hdr,
                  inout metadata meta,
                  inout standard_metadata_t standard_metadata) {

    table ipv4_lpm {
        key = {
            hdr.ipv4.dstAddr: lpm;
        }
        actions = {
            ipv4_forward;
            drop;
        }
        size = 1024;
        default_action = drop();
    }

    register<bit<16>> syn_count;
    register<bit<16>> other_count;

    action ipv4_forward(bit<48> dstMac, bit<9> egressPort) {
        hdr.ethernet.srcAddr = hdr.ethernet.dstAddr;
        hdr.ethernet.dstAddr = dstMac;
        hdr.ipv4.ttl = hdr.ipv4.ttl - 1;
        standard_metadata.egress_spec = egressPort;
    }

    action drop() {
        mark_to_drop();
    }

    action compute_index() {
        bit<16> base = 0;
        bit<16> max_size = 65536;
        hash(meta.index, HashAlgorithm.crc16, base,
             { hdr.ipv4.srcAddr },
             max_size);
    }

    action increment_syn_count() {
        bit<16> current_syn;
        syn_count.read(current_syn, meta.index);
        syn_count.write(meta.index, current_syn + 1);
    }

```

```
    action increment_other_count() {
        bit<16> current_other;
        other_count.read(current_other, meta.index);
        other_count.write(meta.index, current_other + 1);
    }

    action check_syn_threshold() {
        bit<16> syn_cnt;
        syn_count.read(syn_cnt, meta.index);
        const bit<16> MAX_SYN_THRESHOLD = 1000;
        if (syn_cnt > MAX_SYN_THRESHOLD) {
            meta.suspicious = 1;
        } else {
            meta.suspicious = 0;
        }
    }

    action syn_block() {
        if (meta.suspicious == 1) {
            mark_to_drop();
        }
    }

    apply {
        if (hdr.ipv4.isValid()) {
            compute_index();
            if (hdr.tcp.isValid()) {
                if (hdr.tcp.SYN == 1 && hdr.tcp.ACK == 0) {
                    increment_syn_count();
                    check_syn_threshold();
                    syn_block();
                } else {
                    increment_other_count();
                }
            }
            ipv4_lpm.apply();
        }
    }
}
```

(i) A TCP packet arrives with the flags `SYN=1` and `ACK=0` from a source IP that has not been seen before. Describe, step by step, how this packet is processed by the ingress pipeline. In your explanation, identify each relevant pipeline element, explain the purpose of each step, and discuss how the final decision—to forward or drop—is reached. (6 Points)

(ii) The code maintains a separate `other_count`, but it is not currently used to inform forwarding decisions. Explain how `other_count` could be leveraged to reduce false positives in SYN flood detection. Then, propose a new data plane action that uses `other_count` to implement your strategy. Describe the logic of this function and explain how you would integrate it into the existing pipeline. (5 Points)

How `other_count` could help reduce false positives:

How the data plane action logic would look like:

How the data plane action would be integrated into the pipeline:

(iii) Identify two problematic aspects or potential limitations of the design you proposed in **(ii)**. For each issue, propose a mitigation strategy while noting any potential trade-offs (e.g., increased memory usage, added complexity, or a higher likelihood of false positives).

(4 Points)

Problem 1:

Mitigation and trade-off 1:

Problem 2:

Mitigation and trade-off 2:

leave blank

leave blank

Task 3: Verification and Routing Algebra **21 Points**

a) Inter-AS paths with a special peering agreement **(7 Points)**

Consider a modified commercial relationship routing algebra. In addition to provider and customer links (P resp. C), we now distinguish between two levels of peer links: R and Q.

Routes received from both Q- and R-peers are exported to customers (through P links) and to Q-Peers. In addition, routes from Q-peers are exported to R-peers if there is no R link on the route yet (any route can traverse at most one R link).

As an example, consider AS0 in the following network. AS0 would export a route originating from AS3 to AS1, but it would not export a route originating from AS4 to AS1.

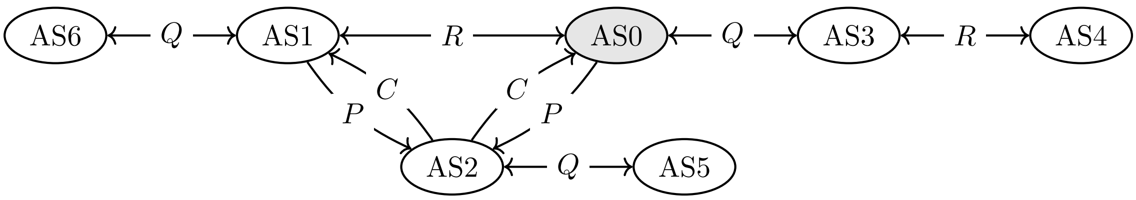

Figure 4: Illustration of attributes added to the route when exporting from arrow tail to arrow head.

(i) Which ASes (excluding AS6) will receive a route originating at AS6? (2 Points)

(ii) Fill out the following table describing the extension operation. (3 Points)

route attribute

new link

$\oplus$	C	R	Q	P	$\bullet$
C					
R					
Q					
P					

(iii) The modified routing algebra uses the following preference relation:

$$C \prec R = Q \prec P \prec \bullet$$

Is this algebra isotone? If so, prove it. Otherwise, provide a counter example.

Recall: a routing algebra is isotone if $\forall X, Y, Z \in A$*:*

$$X \preceq Y \Rightarrow Z \oplus X \preceq Z \oplus Y$$

(2 Points)

b) From verification to routing algebra **(14 Points)**

leave blank

You are an operator of a network $G = (V, E)$ that should route traffic according to the **widest path** routing algebra, as seen in the lecture:

$$(\mathbb{R}_0^+ \cup \{\infty\}, \geq, \min, 0).$$

Your goal is to build a verifier comprised of a set of symbolic variables and logical expressions that correctly models widest-path routing. To that end, you come up with the following fixed-point equation for a destination `t`:

$$\forall v, u \in V : \begin{cases} \mathit{Attr}(v) = \min(w(v,u), \mathit{Attr}(u)) & \text{if } \mathit{From}(v) = u \\ \mathit{Attr}(v) \geq \min(w(v,u), \mathit{Attr}(u)) & \text{otherwise} \end{cases} \tag{1}$$

where:

- $\mathit{Attr} : V \to \mathbb{R}_0^+ \cup \{\infty\}$ return the preferred attribute of any node $v \in V$.
- $\mathit{From} : V \to V$ describes the neighbor from which a node selects its route.

Additionally, $w : V \times V \to \mathbb{R}_0^+ \cup \{\infty\}$ maps any pair of nodes to a width. If there is no link between two nodes (*i.e.* $(v, u) \notin E$), then $w(v, u) = 0$, and any node $v \in V$ can reach itself with infinite width, *i.e.*, $w(v, v) = \infty$.

(i) In the symbolic model of BGP presented in the lecture, we used a variable `Available` to encode whether a node actually knows a route or not. The equation above, however, does not include such a term. Explain why this variable is not necessary in our case or for any other routing algebra? (2 Points)

(ii) Observe that Eq. (1) is trivially satisfied if all routers select themselves, i.e., if $\mathit{From}(v) = v$ and $\mathit{Attr}(v) = \infty$ for all $v \in V$. Write (a set of) additional equation(s) that prevent this trivial solution. (2 Points)

(iii) Equation 1 may not yield a *stable routing*. For the following network, find an assignments of the symbolic maps $Attr$ and $From$ that could not be the stable result of any path vector routing protocol, while satisfying Eq. (1) and your solution from task **(ii)**. (3 Points)

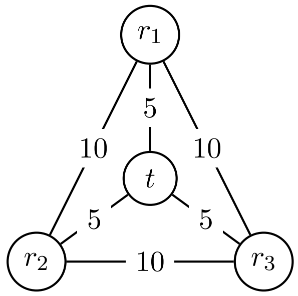

$Attr(r_1) =$ ________ $From(r_1) =$ ________

$Attr(r_2) =$ ________ $From(r_2) =$ ________

$Attr(r_3) =$ ________ $From(r_3) =$ ________

$Attr(t) =$ ________ $From(t) =$ ________

(iv) Let us introduce a symbolic map $PathLen : V \mapsto \{0, 1, \ldots, |V|\}$ to keep track of the length of the learned paths. Use $PathLen$ to write equations that, combined with Eq. (1) and your solution to task **(ii)**, ensure any solution is a *stable routing*. Note that path selection should still be done according to the path widths, not their lengths. In addition, explain why this addition avoids the issue identified in task **(iii)**. (4 Points)

Equations: ________

Explanation: ________

(v) From the lecture, we know that monotonicity implies the existence of stable routing. Further, any stable routing is a solution to the fixed-point equations. However, as you have seen in task **(iii)**, monotonicity does not ensure that *all* solutions to the fixed-point equations are stable routings. Can you describe a property of routing algebra that ensures all solutions to the fixed-point equations are also stable routings? (3 Points)

leave blank

Task 4: Network measurements 16 Points

leave blank

a) Topology inference (6 Points)

In this task, you will analyze a network topology consisting of nine ASes. Your goal is to deduce the business relationships among the ASes – whether they are provider-customer (p2c), customer-provider (c2p), or peer-peer (p2p) – and then answer follow-up questions based on these inferences. The ASes in the topology strictly adhere to the Gao-Rexford principles.

You are given some initial links, as shown in the topology. Note that AS 1, AS 2, and AS 3 form the clique and *do not have any providers*. Additionally, the table below lists a set of BGP paths observed by various vantage points at the end of the BGP convergence process.

Important: In this question you should *not* use ASRank.

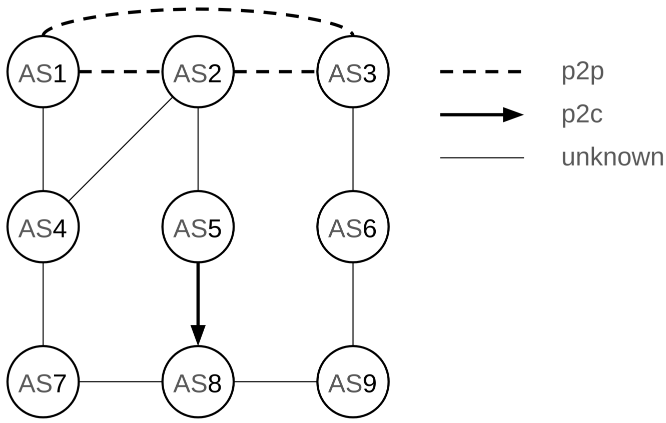

Figure 5: An interdomain topology of 9 ASes.

Vantage Point	Origin AS	Path
AS 4	AS 5	AS 4, AS 2, AS 5
AS 7	AS 5	AS 7, AS 4, AS 2, AS 5
AS 6	AS 5	AS 6, AS 3, AS 2, AS 5
AS 1	AS 8	AS 1, AS 4, AS 7, AS 8
AS 7	AS 9	AS 7, AS 8, AS 9
AS 6	AS 9	AS 6, AS 9
AS 7	AS 3	AS 7, AS 4, AS 1, AS 3

Table 1: BGP paths as observed by various vantage points

(i) Based on the provided paths, classify each edge in the topology in a manner that satisfies the Gao-Rexford conditions for all listed paths. If you cannot "definitively" determine an edge's relationship, *classify it as p2p*. Write your solutions in the table below (you may draw on the graph, but only the answers in the table will be graded). (3 Points)

Link	p2p	p2c	c2p
AS 1 - AS 4			
AS 2 - AS 4			
AS 2 - AS 5			
AS 3 - AS 6			
AS 4 - AS 7			
AS 6 - AS 9			
AS 7 - AS 8			
AS 8 - AS 9			

(ii) We refer to hidden links as links whose presence cannot be directly inferred. However, we can sometimes deduce their absence, since their existence would invalidate certain observed paths. Among the three hidden links provided below, one *cannot* exist. Based on the provided paths, identify this impossible link and briefly explain your reasoning. (3 Points)

- AS 1 – p2c – AS 5
- AS 2 – p2c – AS 8
- AS 3 – p2c – AS 9

b) Network tomography **(10 Points)**

leave blank

In this task, we consider that a link is neutral iff it has the same probability of being good for all paths, otherwise it is non-neutral. We also consider that a network is neutral iff it only has neutral links. Finally, we refer to performance as the logarithm of the probability of (a link, set of links, path, or pathset) being good.

(i) Reverse-engineering a neutral network from tomography equations

Consider a neutral network with an unknown topology of 8 switches (ids 0-7) and no end-hosts. We assume that independence, stationarity, separability, and stability hold. The (correct) equations that describe the performance of three paths of the network are:

- $Y_0 = X_{01} + X_{12} + X_{23} + X_{34}$
- $Y_1 = X_{23} + X_{34} + X_{02}$
- $Y_2 = X_{05} + X_{53} + X_{36} + X_{67}$

where:

- Y_i indicates the performance of path p_i;
- X_{ij} indicates the performance of the **link from switch** i **to switch** j.

▷ Based on the above equations, what is the **minimum** number of distinct links that exist in this network? (1 Point)

▷ Can the network have other links that do not appear in the above equations? Briefly explain. (1 Point)

▷ Indicate the minimal set of switches on which we need to deploy monitoring code to estimate Y_0, Y_1, and Y_2. (1 Point)

(ii) Identifying non-neutral links in a datacenter-like topology

Consider the datacenter topology shown in Fig. 6, where:

- circles indicate servers with ids 0-7;
- squares indicate switches with ids 8-14;
- tomography can only measure performance between servers;
- a server i can send traffic to server j only if $i < j$;
- if server i sends to server j, the traffic follows the **shortest path** (in terms of number of links) from i to j.

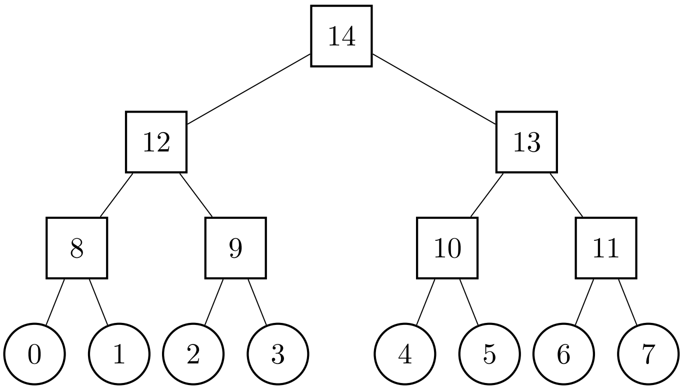

Figure 6: Datacenter-like topology for question (ii).

▷ Can network tomography prove that the link l_{12-14} between switch 12 and switch 14 is non-neutral? If yes, indicate which of the allowed paths would be used for the tomography equations. If not, explain briefly. (3 Points)

____

____

____

____

▷ Consider now an alternate topology in which switch 14 and its adjacent links (l_{12-14} and l_{13-14}) are replaced by a link l_{12-13} between switch 12 and 13. In this topology, three paths are used:

- p_0: $0 \to 8 \to 12 \to 13 \to 10 \to 4$;
- p_1: $1 \to 8 \to 12 \to 13 \to 11 \to 6$;
- p_2: $3 \to 9 \to 12 \to 13 \to 11 \to 6$.

Using network tomography, write the equations to prove that l_{12-13} is non-neutral by filling in the lines below. Y_i indicates the performance of path p_i, and Y_{ij} indicates the performance of path pair p_{ij}. You must:

- leave empty non-used line(s);
- use $X_{i-j,k-m,\ldots}$ to indicate the performance of links $\{l_{i-j}, l_{k-m}, \ldots\}$, where l_{i-j} (resp. l_{k-m}) is the link between server/switch i and j (resp. k and m). E.g., use $X_{0-8,8-12}$ to indicate the performance of the set of links $\{l_{0-8}, l_{8-12}\}$. (4 Points)

$Y_0 =$ ____

$Y_1 =$ ____

$Y_2 =$ ____

$Y_{01} =$ ____

$Y_{02} =$ ____

$Y_{12} =$ ____

leave blank

leave blank

Task 5: Network Security — 8 Points

a) Defending DDoS attacks at target's network (5 Points)

Figure 7 shows ISPs Swizzcom (controlling router R4) and Moonrise (controlling routers R1, R2, R3, and a public DNS server at `100.64.6.65`). You are Moonrise's operator.

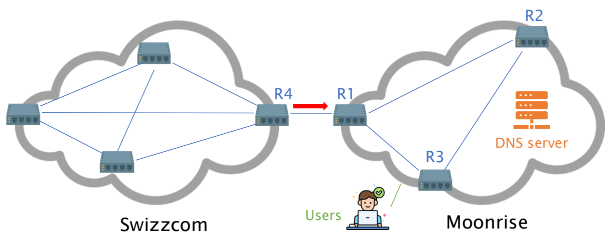

Figure 7: Network setup of two ISPs where Swizzcom is the sole provider of Moonrise.

(i) **Attack.** You observe a DDoS attack toward the DNS server. The traffic's source IPs belong to NTP servers all over the world and destination ports look random. (2 Points)

Affected victim's layer:

true	false	
☐	☐	Network (layer 3)
☐	☐	Transport (layer 4)
☐	☐	Application (layer 7)

Attack strategy:

true	false	
☐	☐	Spoofed Direct-path
☐	☐	Non-spoofed Direct-path
☐	☐	Reflection-amplification

(ii) **Defense.** You want to use Remotely Triggered Black Hole (RTBH) to block the DDoS traffic while maintaining DNS services for your own users. You've already configured a static route for `192.0.2.1/32` that points to interface `Null0` on R1, R2, and R3. You also know that R4 has configured a static route for `192.0.2.2/32` that points to interface `Null0` and that Swizzcom's RTBH community is 666. (3 Points)

The router(s) you send the BGP message from: ____________________

____

The router(s) you send the message to: ____________________

____

The BGP message that triggers RTBH: `ip route` ____________________

____

b) Defending DDoS with `iptables` (3 Points)

leave blank

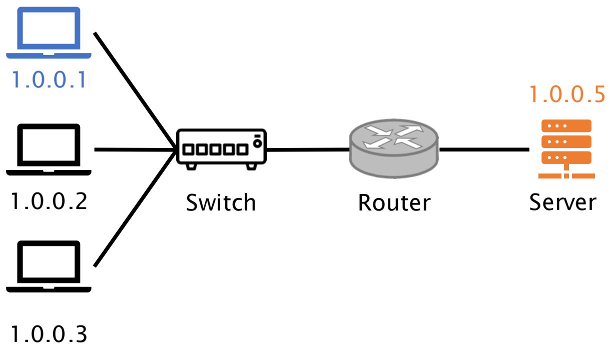

Figure 8: A simple network.

You are the operator of the network shown in Figure 8. Your goal is to ensure the data transmitted between the benign host `1.0.0.1` and the server is not affected by three following DDoS attacks. In each attack, you are provided with the summarized conversations from `tshark` traces. Write down the `[MATCH]` field of the `iptables` rule(s) you need on the router to filter out malicious traffic in each case. The relevant `[MATCH]` flags are as follows:

- "`-s <ip address>`": matching the source IP address
- "`-d <ip address>`": matching the destination IP address
- "`-m mac --mac-source <MAC ADDRESS>`": matching the MAC source address

(i)

```
                                          |   <-     |   ->     |   Total
    1.0.0.1 <-> 1.0.0.5                   |   3      |   2      |   5
    1.0.0.2 <-> 1.0.0.5                   |   5      |   90     |   95

    06:bf:13:f4:34:ca <-> 0a:0a:21:3f:74:28 |   3      |   2      |   5
    16:91:75:69:af:12 <-> 0a:0a:21:3f:74:28 |   5      |   90     |   95
```

Listing 1: IPv4 and Ethernet conversations in Attack 1

(1 Point)

(ii)

```
                                            |   <-     |   ->     |   Total
    1.0.0.1 <-> 1.0.0.5                     |   1      |   1      |   2
    1.0.0.2 <-> 1.0.0.3                     |   0      |   43     |   43
    1.0.0.1 <-> 1.0.0.3                     |   55     |   0      |   55

    06:bf:13:f4:34:ca <-> 0a:0a:21:3f:74:28 |   1      |   1      |   2
    16:91:75:69:af:12 <-> ec:61:68:df:df:10 |   0      |   43     |   43
    ec:61:68:df:df:10 <-> 06:bf:13:f4:34:ca |   0      |   55     |   55
```

Listing 2: IPv4 and Ethernet conversations in Attack 2

(1 Point)

(iii)

```
                                            |   <-     |   ->     |   Total
    1.0.0.1 <-> 1.0.0.5                     |   1      |   99     |   100

    06:bf:13:f4:34:ca <-> 0a:0a:21:3f:74:28 |   1      |   1      |   2
    16:91:75:69:af:12 <-> 0a:0a:21:3f:74:28 |   0      |   20     |   20
    ec:61:68:df:df:10 <-> 0a:0a:21:3f:74:28 |   0      |   78     |   78
```

Listing 3: IPv4 and Ethernet conversations in Attack 3

(1 Point)

leave blank

leave blank

Task 6: Transport Protocols 19 Points

a) Warm-up (8 Points)

(i) What is head-of-line blocking? Why is it a problem in TCP, but not in QUIC. (2 Points)

(ii) Why does MPTCP perform a 3-way handshake for each subflow? Provide an example that shows why this is necessary. (2 Points)

(iii) Why is QUIC typically implemented as a userspace library? (2 Points)

(iv) Why does BBR use the bandwidth-delay-product as a target for its sending rate? Name and explain one disadvantage of using a smaller or larger rate. (2 Points)

Disadvantage of smaller rate:

Disadvantage of larger rate:

b) CCA Design Space **(6 Points)**

leave blank

Name three different network signals that are used by existing CCAs to model a network.

Signal 1: ____________________

Signal 2: ____________________

Signal 3: ____________________

Figure 9 shows the trace of an unknown CCA. The figure also shows a packet loss event as a dashed, vertical line. Figure 10 shows the buffer utilization of all routers in the network.

For each signal mentioned above, indicate whether that CCA makes use of it or if it is impossible to say given the data. Justify your answers using evidence from Figures 9 and 10.

You may annotate the plots but **only the answers in the answer fields** on the next page will be graded.

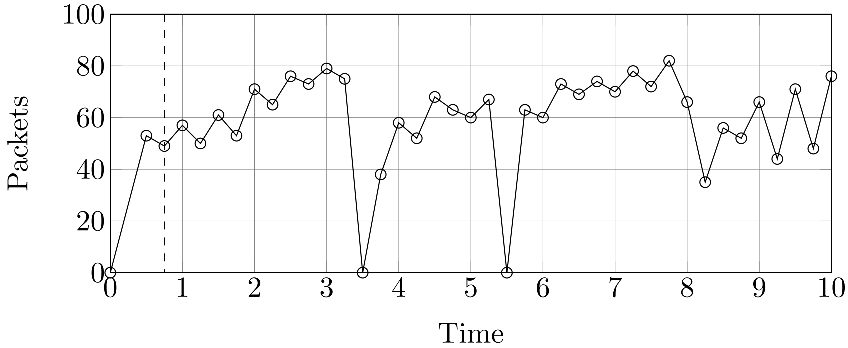

Figure 9: A trace of an unknown CCA

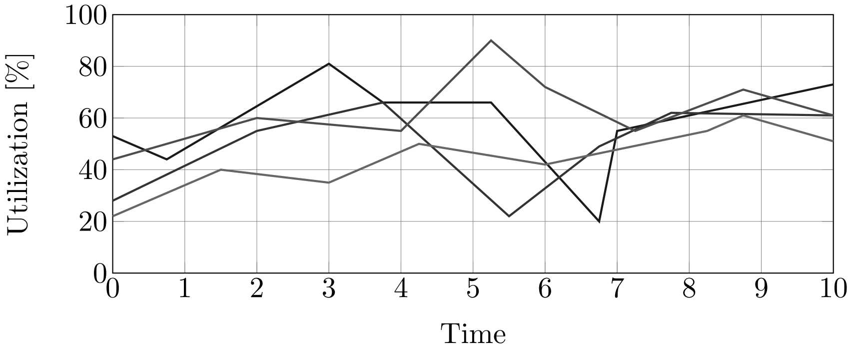

Figure 10: The recorded utilization of 4 buffers in the network

Signal 1: true ☐ false ☐ inconclusive ☐

Justification: ______________________________

Signal 2: true ☐ false ☐ inconclusive ☐

Justification: ______________________________

Signal 3: true ☐ false ☐ inconclusive ☐

Justification: ______________________________

c) Contention in the public internet **(5 Points)**

leave blank

Consider the network in Figure 11. A and B are hosts connected to each other via R1 and R2, respectively. Beside the traffic between A and B, R1 and R2 also forward traffic from the public internet. This traffic has an unpredictable pattern and always enters the network via X and leaves it via Y

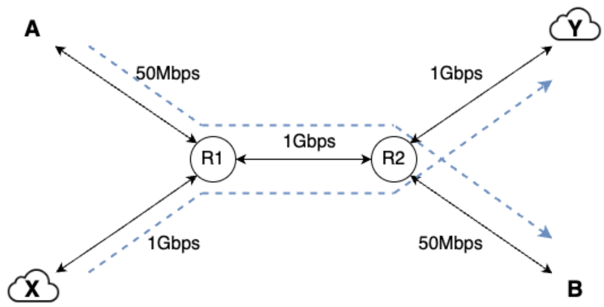

Figure 11: The public internet with A and B connected to it via R1 and R2

(i) Assume host A uploads a large file to host B. With respect to this connection, name an advantage and disadvantage of increasing R1's buffer size for the link to R2. Justify your answer. *Note*: A and B maintain no open connection to the public internet. (2 Points)

(ii) With respect to the same upload, is an increase of R2's buffer size beneficial as well? Justify your answer. (1 Point)

(iii) Name one simple strategy that A can do in order to improve performance without changing anything about the network. Explain your answer. (2 Points)

leave blank

Task 7: Sustainable networking **15 Points**

In a 2024 Nature's Scientific Reports paper, researchers estimated and compared the carbon footprint of writing a page of text for ChatGPT and humans. For ChatGPT, the researchers accounted for:

- 1.84 gCO2e/query for the energy used for training the model once per month;
- 0.38 gCO2e/query for the energy for inference (i.e., answering the query);
- 0.0001 gCO2e/query for the manufacturing and recycling cost of the hardware used for training the model;
- A scaling factor of 0.6 for the fact that one typical ChatGPT query produces more words than one page of text (250 words).

This results in an estimated footprint of 1.32 gCO2e/page of text produced by ChatGPT.

For the humans, the researchers considered the annual carbon footprint estimated for an entire country and divided by the number of inhabitants and the number of hours in the year to get a footprint per person per hour. This is then multipled by the average time a professional writter takes to produce 250 words (0.86h, or 50min). This results in an estimated footprint of 1426 gCO2e/page of text for a writter located in the USA.

The paper concludes that producing a page of text with ChatGPT is about 1000 times more carbon efficient than if written by a human located in the USA.

(i) What type of analysis was done for estimating the footprint of ChatGPT/humans writing a page of text? Which dimensions were considered? Tick all that apply. (4 Points)

ChatGPT

Type of analysis

☐ Attributional ☐ Consequential
☐ Top-down ☐ Bottom-up

Dimensions considered

☐ Operational ☐ Embodied
☐ Scope 1 ☐ Scope 2 ☐ Scope 3

Humans

Type of analysis

☐ Attributional ☐ Consequential
☐ Top-down ☐ Bottom-up

Dimensions considered

☐ Operational ☐ Embodied
☐ Scope 1 ☐ Scope 2 ☐ Scope 3

(ii) The footprint comparison between ChatGPT and humans is not fair. Why? (2 Points)

(iii) How would you change the footprint estimation methodology for ChatGPT, humans, or both in order to make a fair comparison? (4 Points)

(iv) Let us assume a fair comparison concludes that ChatGPT is (still) more efficient at text writing than humans. There exists a risk that transitioning to using ChatGPT only for text writing increases the footprint of text writing overall. (3 Points)

What is the name of this effect?

Explain the theory about how this effect works.

Give one real-life example where this effect happened.

(v) Can you compute an estimate for the footprint of running one *additional* ChatGPT query based on the data provided above? If no, explain why not. If yes, give the numerical value and explain your reasoning. (2 Points)

Extra Sheet

In case you need more space, use the following pages. Make sure to always indicate the task to which the answer belongs (e.g., *3 d) (ii)*).

Task: ____________

Task: ____________

Task: ______

Task: ______