ETH zürich

Networked Systems
ETH Zürich — seit 2015

D-ITET

Prof. Dr. Laurent Vanbever

February 2024

Exam: Advanced Topics in Communication Networks

12 February 2024, 15:30-18:00, Room HIL G75

- 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 - 40).
- **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	/33
Programmable data planes	/21
Network verification	/32
Network measurements	/21
Network security	/11
Transport	/21
Sustainable networking	/11
Total	/150

leave blank

leave blank

Task 1: Advanced Routing **33 Points**

a) Hierarchical routing **(9 Points)**

Consider the route reflection topology composed of five routers depicted in Fig. 1. `RR1` and `RR2` act as route reflectors, while `RA`, `RB`, and `RC` 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., `RA` is the client of `RR1`, while `RB` is the client of both `RR1` and `RR2`). The network relies upon an IGP whose weights are depicted next to each link. Three routers—`RA`, `RB`, and `RR2`—receive an *equivalent* eBGP route for the same prefix p (`r1`, `r2`, and `r3` in Fig. 1) meaning that all of their BGP attributes are equal to each other.

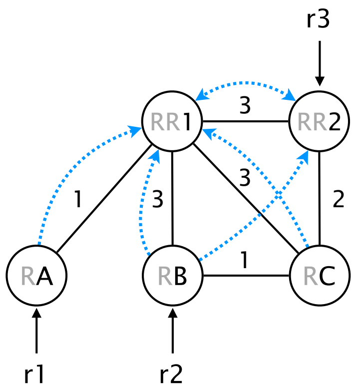

Figure 1: A simple route reflection topology.

(i) Indicate which egress router `RC` choses alongside with the path the traffic it sends to p takes. If something is sub-optimal and/or incorrect about this path, briefly explain it. (2 Points)

Egress router selected by `RC`: ______________________

Path taken by the traffic: ______________________

Something sub-optimal and/or incorrect? ______________________

____

____

(ii) Consider now that the link between `RC` and `RR1` fails. Indicate which egress router `RC` choses after the failure alongside with the path the traffic it sends to p takes. If something is sub-optimal and/or incorrect about this path, briefly explain it. (2 Points)

Egress router selected by `RC`: ______________________

Path taken by the traffic: ______________________

Something sub-optimal and/or incorrect? ______________________

____

____

(iii) In addition of the failure of the link between `RC` and `RR1`, consider now what would happen if `RB` also looses its external route (`r2` stops being advertised). Indicate which egress router `RC` choses alongside with the path the traffic it sends to p takes. If something is sub-optimal and/or incorrect about this path, briefly explain it. (2 Points)

Egress router selected by `RC`: __________

Path taken by the traffic: __________

Something sub-optimal and/or incorrect? __________

(iv) Briefly explain whether adding one extra client session between `RC` and `RR2` would have been enough to prevent sub-optimal routing/forwarding and/or incorrect routing/forwarding in the situation **(iii)** depicted just above. (3 Points)

Would adding a client session (`RC`, `RR2`) prevent sub-optimal routing/forwarding? __________

Would adding a client session (`RC`, `RR2`) prevent incorrect routing/forwarding? __________

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

leave blank

Consider the weighted network topology composed of seven routers depicted in Fig. 2.

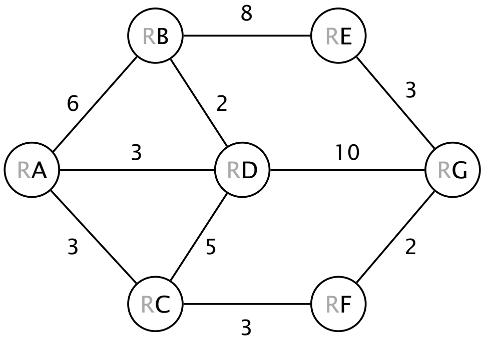

Figure 2: A simple internal topology.

(i) Briefly explain: (i) what a Loop Free Alternate (LFA) is; and (ii) what a link-protecting LFA is. (2 Points)

What is an LFA?

What is a link-protecting LFA?

(ii) Is `RA` a link-protecting LFA for `RD` considering the failure of the link (`RD`, `RC`)? Briefly explain. (1 Point)

(iii) Is `RB` a link-protecting LFA for `RD` considering the failure of the link (`RD`, `RC`)? Briefly explain. (1 Point)

(iv) Is RG a link-protecting LFA for RD considering the failure of the link (RD, RC)? Briefly explain. (1 Point)

(v) Is RB a **per-prefix** LFA for RD for the prefixes originated by RG and when considering the failure of the link (RD, RC)? Briefly explain. (2 Points)

(vi) Is RE a **remote** LFA for RD considering the failure of the link (RD, RC)? Briefly explain. (2 Points)

c) Local prefix filtering and aggregation **(11 Points)**

leave blank

(i) Briefly explain the concepts of strong and weak forwarding consistency and how they differ from each other. (3 Points)

Strong forwarding consistency . . .

Weak forwarding consistency . . .

How do they differ?

(ii) Briefly explain two factors that will influence how "aggregatable" a forwarding table is. (2 Points)

Factor 1:

Factor 2:

In the rest of this question, we ask you to aggregate the forwarding table given below using the Optimal Routing Table Constructor (ORTC) algorithm.

* ⇒ 1
0* ⇒ 2
1* ⇒ 2
01* ⇒ 3
001 ⇒ 3
101 ⇒ 1

(iii) Fill in the empty binary tree below with the tree state obtained at the end of pass 1 (normalization) when running ORTC on the table above. (2 Points)

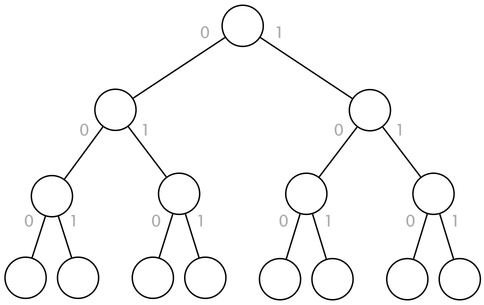

Figure 3: State reached after pass 1.

(iv) Fill in the empty binary tree below with the tree state obtained at the end of pass 2 (next-hop ranking) when running ORTC on the table above. (2 Points)

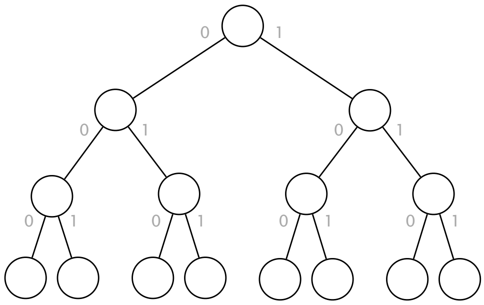

Figure 4: State reached after pass 2.

(v) Fill in the empty binary tree below with the tree state obtained at the end of pass 3 (prefix filtering) when running ORTC on the table above **alongside** with the final forwarding rules obtained (write the rules in the table below). (2 Points)

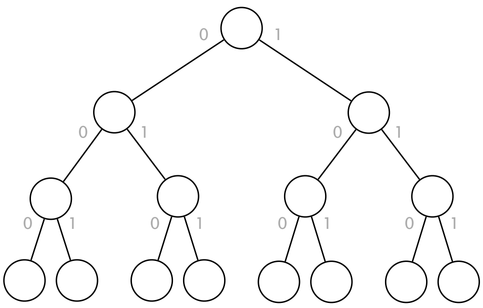

Figure 5: State reached after pass 3.

Figure 6: Final forwarding table obtained.

d) Distributed prefix filtering and aggregation **(4 Points)**

leave blank

(i) Draw an inter-domain network topology composed of 6 routers in which one router advertises a parent prefix p; another router advertises a child prefix q (contained in p); and **all** the other 4 routers can filter the prefix q following the DRAGON filtering rules. Draw provider-customer links using straight lines, with the providers above customers, and peer-to-peer links using dashed lines (as we have done in the exercises). (2 Points)

Briefly explain why is your topology correct: ________

(ii) Same question as above except that this time your topology should be such that **none** of other 4 routers should be able to filter out the prefix q following the DRAGON filtering rules. Use the same drawing convention (see above). (2 Points)

Briefly explain why is your topology correct: ________

leave blank

Task 2: Programmable data planes — 21 Points

a) General Questions (4 Points)

leave blank

For each of the following statements, indicate whether they are *true* or *false*. There is always one correct answer. Grading is as such: 4 points for four correct answers, 2 points for three correct answers, and 0 points otherwise.

true ☐ false ☐ During packet recirculation in a P4 program, a packet can retain and transport computation state information by appending it as an additional header.

true ☐ false ☐ Detecting heavy-hitter flows in P4 requires the use of stateful data structures.

true ☐ false ☐ In a P4 switch, packet drops will not occur unless explicitly specified in the P4 code.

true ☐ false ☐ The P4 programming language supports the use of a division operator.

leave blank

b) Probabilistic data structures (17 Points)

The objective of this task is to implement a Bloom filter designed to block network packets targeting specific ports. This functionality is essential in firewall configurations to prevent unauthorized packet entry into a network.

The Bloom filter you will work on comprises two distinct arrays, referred to as registers. Each register is associated with its unique hash function. For a given input x, each hash function $h_i(x)$ corresponds to a specific register R_i.

The hash functions are defined as:

$$h_1(x) = sum(x) \mod 10$$
$$h_2(x) = mult(x) \mod 10$$

In this context, sum(x) calculates the sum, and mult(x) the product, of x's digits.

For instance, for x = 72, $h_1(x)$ yields $sum(72) = 9$, and $h_2(x)$ gives $mult(72) = 14$. Post modulo 10, hash values are 9 and 4, setting bits at these positions in registers $\mathbf{R_1}$ and $\mathbf{R_2}$.

(i) Bloom filter configuration (3 Points)

Insert port numbers 731, 42, and 1410 sequentially into the Bloom filter as shown below. For each port, set the corresponding bits in the registers. Begin with the filter in its initial state (empty arrays) and proceed through each insertion step (first insertion being 731).

	$\mathbf{R_1}$	$\mathbf{R_2}$		$\mathbf{R_1}$	$\mathbf{R_2}$		$\mathbf{R_1}$	$\mathbf{R_2}$		$\mathbf{R_1}$	$\mathbf{R_2}$
0	0	0									
1	0	0									
2	0	0									
3	0	0									
4	0	0	→			→			→		
5	0	0									
6	0	0									
7	0	0									
8	0	0									
9	0	0									

Note: The table above demonstrates the configuration of the Bloom filter for blocking ports 731, 42, and 1410. Each '1' in the table represents a bit set by the hash functions for these ports, indicating that packets destined for these port numbers will be blocked.

(ii) Analysis of blocked web traffic (2 Points)

Web traffic, typically using port 443, is unexpectedly blocked. Diagnose this issue by identifying the cause and provide a rationale for your conclusion.

The cause of the error: ____________________

Explanation: ____________________

(iii) Expansion of the Bloom Filter (2 Points)

Consider adding an extra array to the Bloom filter. Discuss the criteria for choosing a hash function for this new array, especially in the context of the ongoing web traffic issue, assuming ports 731, 42, and 1410 remain in the filter.

The hash function: ____________________

Explanation: ____________________

(iv) Evaluation of element removal from the Bloom filter (2 Points)

Attempt to remove port 42 from the Bloom filter by resetting its associated bits to 0. Observe any side effects that occur and discuss how this reflects a fundamental limitation of Bloom filters.

The side effect: ____________________

The limitation: ____________________

(v) Implementation of a Counting Bloom filter (8 Points)

This section focuses on implementing a Counting Bloom filter to address the limitations of a standard Bloom filter. In this design:

- Each cell in the standard Bloom filter's register is replaced by a counter.
- Inserting an element increments the counter for its associated bits, while removing an element decrements these counters.
- A query is positive if all relevant counters are greater than zero.
- This approach allows for the removal of elements, contrasting with the binary cells of the standard Bloom filter.

Listing 1: P4 code template for a standard Bloom filter.

```
#define N_CELLS 10

struct metadata {
    bit<32> hash_one;
    bit<32> hash_two;
    bit<16> query_result;
}

metadata meta;
register<bit<16>>(N_CELLS) bloom_array_one;
register<bit<16>>(N_CELLS) bloom_array_two;

action compute_hashes() {
    // Compute hashes, abbreviated for simplicity.
    hash(meta.hash_one, ...);
    hash(meta.hash_two, ...);
}

action insert() {
    // Update Bloom filter arrays.
    bloom_array_one.write(meta.hash_one, 1);
    bloom_array_two.write(meta.hash_two, 1);
}

action query() {
    // Read from Bloom filter arrays.
    bit<16> result_one;
    bit<16> result_two;
    bloom_array_one.read(result_one, meta.hash_one);
    bloom_array_two.read(result_two, meta.hash_two);

    // Set query result based on Bloom filter arrays.
    meta.query_result = result_one & result_two;
}

control bloom_filter_control {
    compute_hashes();
    insert();
    query();

    // Logic to handle query_result...
}
```

The P4 code template provided above demonstrates the standard Bloom filter implementation. Your task is to modify this code to develop the following actions for the Counting Bloom filter:

`increment()`: inserts `x` into the Counting Bloom filter
`decrement()`: deletes `x` from the Counting Bloom filter
`query()`: checks if `x` is present in the Counting Bloom filter

Based on the standard Bloom filter code template, write P4 actions for `increment()`, `decrement()`, and `query()` tailored to the Counting Bloom filter.

`increment()`:

`decrement()`:

`query()`:

leave blank

leave blank

Task 3: Verification and Synthesis 32 Points

a) Warm-up (4 Points)

(i) What guarantees can network verification provide to a network operator? (2 Points)

(ii) Describe two drawbacks of using simulation to verify network configurations. (2 Points)

b) Min-Hop Routing (10 Points)

leave blank

Min-Hop Routing is a protocol in which routers select the route with the least number of hops. Like BGP, Min-Hop Routing is a distance-vector protocol. All neighboring routers exchange their number of hops to the destination. Each router selects **any** one of its neighbors with the smallest number of hops. We encode the routing state of a router `r`:

- `r.Available`: Does `r` know a route towards the destination?
- `r.SelectsFrom`: From which neighbor `n` does `r` know its selected route?
- `r.Hops`: How many hops does `r` take to reach the destination?

You want to verify a network running Min-Hop Routing. In the following, you are given three variations of route-selection equations. For each variation, determine whether it correctly encodes the route selection of Min-Hop Routing as described above. You should assume the route propagation has been correctly encoded.

Hint: We want to verify properties of an unknown implementation of Min-Hop Routing, i.e., the equation should model any correct implementation of Min-Hop Routing.

```
A = If(r.SelectsFrom == n,
       And(n.Available, r.Hops == n.Hops + 1),
       Implies(n.Available,
               Or(r.Hops < n.Hops + 1,
                  And(r.Hops == n.Hops + 1, r.SelectsFrom < n))))

B = If(r.SelectsFrom == n,
       And(n.Available, r.Hops == n.Hops + 1),
       Implies(n.Available, r.Hops <= n.Hops + 1))

C = If(r.SelectsFrom == n,
       And(n.Available, r.Hops == n.Hops + 1),
       Implies(n.Available, r.Hops < n.Hops + 1))
```

(ii) Is A a correct encoding? Justify your answer. (2 Points)

(iii) Is B a correct encoding? Justify your answer. (2 Points)

(iv) Is C a correct encoding? Justify your answer. (2 Points)

(v) Rank the three equations from best to worst. Justify your answer. (2 Points)

Ranking: ________ is better than ________ is better than ________

Reason: ________

(vi) Justify why two routers cannot select each other in any equation above. (2 Points)

c) Specifications (10 Points)

leave blank

You are given the BGP network shown above. It shows physical links and their IGP link weights:

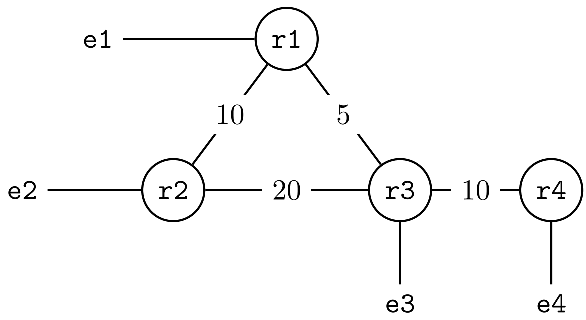

The external networks `e1`, `e2`, `e3`, and `e4` *may* advertise a BGP route towards internal routers `r1`, `r2`, `r3`, and `r4`, respectively. All internal routers are connected in an iBGP full-mesh. The configured BGP Route Maps, however, are unknown.

You are given the following specification: (We use the same variable as in the lecture. A reference is provided on Page 21)

```
Spec = Or(And(e3.Available,
              r1.SelectsFrom == r3, r2.SelectsFrom == r3,
              r3.SelectsFrom == e3, r4.SelectsFrom == r3),
          And(Not(e3.Available),
              r4.SelectsFrom != e4,
              Implies(And(e1.Available, e2.Available,
                          e1.Route.AsPathLen == e2.Route.AsPathLen),
                      And(r1.SelectsFrom == e1, r2.SelectsFrom == e2,
                          r3.SelectsFrom == r1, r4.SelectsFrom == r1))))
```

(i) According to `Spec`, how should the network handle routes from `e3`? (2 Points)

Property 1: ______________________________

____

(ii) According to `Spec`, how should the network handle routes from `e4`? (2 Points)

Property 2: ______________________________

____

(iii) According to `Spec`, how should the network decide between routes from `e1` and `e2`? (2 Points)

Property 3: ______________________________

____

____

(iv) The following table lists a set of routing states.

- For each external router, we list either `eX.AsPathLen`, or ∞ if `eX.available` is `False`.
- For each internal router, we list `rX.SelectsFrom`.

For each of the routing states, determine whether it can be reached by a valid configuration, i.e., one that satisfies the specification. Briefly justify your answers. (4 Points)

routing state								possible		reason
e1	e2	e3	e4	r1	r2	r3	r4	Yes	No	
1	1	1	∞	e1	e2	e3	r3	☐	☐	
1	1	∞	∞	r2	r1	r1	r1	☐	☐	
3	5	∞	∞	e1	r1	r1	r1	☐	☐	
1	∞	∞	∞	e1	r1	e3	r3	☐	☐	

Variable reference (same as in the lecture)

- An internal router (`r1`, `r2`, `r3`, and `r4`) has the following variables:
  - `rX.Available`: Whether the router selects a route.
  - `rX.SelectsFrom`: The neighbor from which `rX` selects its route.
  - `rX.Route`: The selected route attributes
- An external network (`e1`, `e2`, `e3`, and `e4`) has the following variables:
  - `eX.Available` Whether the network advertises a route
  - `eX.Route`: The advertised route attributes
- A BGP route has the following attributes:
  - `Route.LocalPref`
  - `Route.AsPathLen`
  - `Route.Community`
  - `Route.IgpCost`

d) Synthesis **(8 Points)**

leave blank

(i) Why do we need to parametrize the configuration for network synthesis, and how are configuration parameters represented in SMT? (2 Points)

The following figure shows the inner loop of Counter Example Guided Inductive Synthesis (CEGIS), as presented in the lecture:

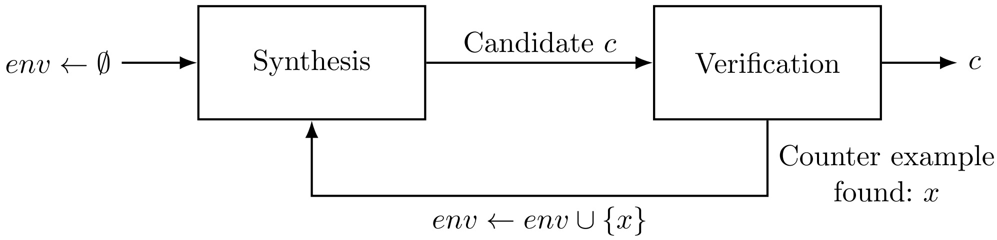

(ii) The **Synthesis** block of CEGIS solves an SMT problem. The SMT solver may return either `sat` or `unsat`. How does CEGIS use that result? (3 Points)

`sat`:

`unsat`:

(iii) The **Verification** block of CEGIS solves an SMT problem. The SMT solver may return either `sat` or `unsat`. How does CEGIS use that result? (3 Points)

`sat`:

`unsat`:

leave blank

Task 4: Network measurements **21 Points**

leave blank

a) Interdomain routing inference (10 Points)

We ask you to infer some of the business relationships in the inter-domain topology composed of 8 ASes depicted in Figure 7. To do so, you will use the **complete** BGP table observations from AS 7 and AS 8 depicted in Figure 8.

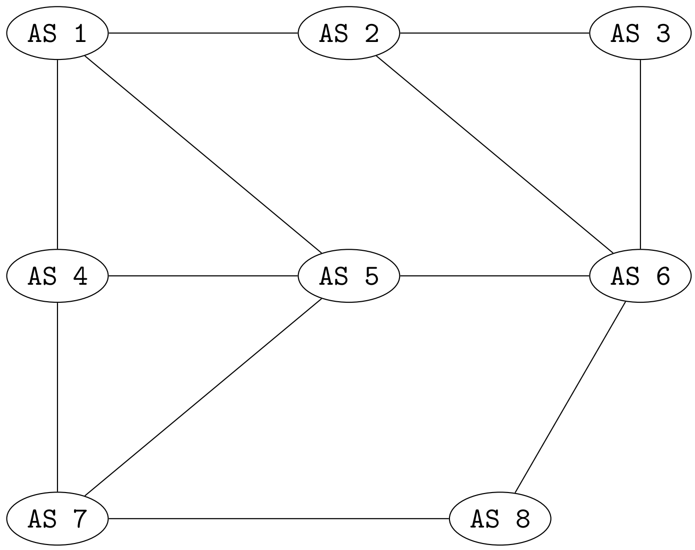

Figure 7: AS-level topology

AS 7 Routing Table

Origin AS	Path
AS 1	[AS 4, AS 1]
	[AS 5, AS 1]
AS 2	[AS 4, AS 1, AS 2]
	[AS 5, AS 1, AS 2]
AS 3	—
AS 4	[AS 4]
	[AS 5, AS 4]
AS 5	[AS 5]
	[AS 4, AS 5]
AS 6	[AS 5, AS 6]
	[AS 4, AS 1, AS 2, AS 6]
AS 8	[AS 8]
	[AS 5, AS 6, AS 8]

AS 8 Routing Table

Origin AS	Path
AS 1	[AS 7, AS 4, AS 1]
	[AS 6, AS 2, AS 1]
AS 2	[AS 7, AS 4, AS 1, AS 2]
	[AS 6, AS 2]
AS 3	[AS 6, AS 3]
AS 4	[AS 7, AS 4]
	[AS 6, AS 2, AS 1, AS 4]
AS 5	[AS 6, AS 5]
	[AS 7, AS 5]
AS 6	[AS 6]
	[AS 7, AS 5, AS 6]
AS 7	[AS 7]
	[AS 6, AS 5, AS 7]

Figure 8: Complete routing tables observed at AS 7 and AS 8

You can make the following assumptions:

1. **Import policies:** ASes follow the classical BGP import policies, except for the tie-breaking criteria, i.e.:
   (a) Prefer routes received from customers over those from peers, and routes from peers over those from providers.
   (b) Amongst equally-preferred routes, prefer the routes with shorter AS-path lengths.
   (c) Amongst equally-short routes, prefer the route received from the AS with the smaller AS number (tie-break).
2. **Export policies:** ASes follow the classical BGP export policies, i.e.:
   (a) Routes received from customers are advertised to all: customers, peers, and providers.
   (b) Routes received from peers are advertised only to customers.
   (c) Routes received from providers are advertised only to customers.
3. **There is *no* sibling relationships.** The network topology does not include any sibling relationships.

(i) Consider the path [AS 8, AS 7, AS 4, AS 1, AS 2]. Given that AS 1 and AS 2 have a peering relationship, infer the relationship between AS 7 and AS 8. Provide an explanation and justify why other relationship types are not possible. (1 Point)

Inferred relationship for (AS 7, AS 8): ______________________________

Explanation: ______________________________

____

____

(ii) Given that AS 1 is the provider of AS 5, determine the relationship between AS 5 and AS 7. Briefly explain and justify why other relationships are not possible. (1 Point)

Inferred relationship for (AS 5, AS 7): ______________________________

Explanation: ______________________________

____

____

(iii) Given that AS 3 is the provider of AS 6, infer the relationship between AS 5 and AS 6. Briefly explain and justify why other relationships are not possible. (2 Points)

Inferred relationship for (AS 5, AS 6): ____________________

Explanation: ____________________

(iv) Gao's relationship inference algorithm starts by calculating the degree of each AS. Explain how the selection of vantage points can introduce a bias in this step. Discuss how such bias might degrade the algorithm's accuracy. (2 Points)

How can the vantage point selection bias degrees estimation? ____________________

How can this negatively affect the inference accuracy? ____________________

(v) Cloud and content providers establish *many* peer-peer relationships without offering any transit services. How can this affect Gao's inference algorithm? What are the likely outcomes of the algorithm in this scenario? (2 Points)

Explanation: ____________________

(vi) Suppose a vantage point artificially increasing the degree of an AS by creating fake AS paths crossing it. Describe an approach to identify such inflated ASes in the dataset. (2 Points)

Explanation: ____________________

b) Network tomography (11 Points)

leave blank

This question is about inferring the performance of one or more links of the network N depicted below (Fig. 9). All links ($l_1, \ldots, l_7$) are indicated by solid lines, and all paths ($p_1, \ldots, p_5$) are indicated by dashed lines.

Recall that, during a time interval, a link, set of links, path, or pathset is "non-lossy" if it introduces negligible packet loss during that interval; otherwise, it is "lossy."

Recall that, for a neutral network, performance refers to the logarithm of the probability of (a link, set of links, path, or pathset) being non-lossy. Whenever you indicate performance, **use the following notation:**

- X_i to indicate the performance of link l_i;
- X_{ij} to indicate the performance of the set of links $\{l_i, l_j\}$;
- Y_i to indicate the performance of path p_i;
- Y_{ij} to indicate the performance of path pair p_{ij}.

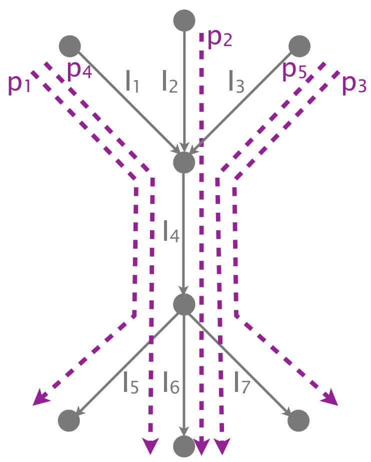

Figure 9: Network N.

(i) Neutral network

Consider that all the links of network N are neutral. Also, independence, stationarity, separability, and stability hold.

▷ Assume that we monitor all pathsets of N and we build a system of equations by writing an equation for each pathset/path. State a sufficient condition that allows uniquely inferring the performance of each link using the aforementioned system of equations. Briefly explain why the condition holds for N. (2 Points)

Condition: ______________________________

____

Why the condition holds for N: ______________________________

____

____

▷ For each path/path pair p_1, p_2, p_{12}, write the equation that connects the path/path pair performance to the performance of its links. (2 Points)

Y_1 = __________

Y_2 = __________

Y_{12} = __________

(ii) Non-neutral network

Consider now that all of N's links are neutral *except for* l_4, meaning that the network N is non-neutral. For the rest of this task, **ignore paths** p_4 **and** p_5**.**

▷ To identify l_4 as non-neutral, tomography builds a system of equations, say S. For each of the path/path pair $p_1, p_2, p_3, p_{12}, p_{13}, p_{23}$, write the equation of S that corresponds to it. (3 Points)

Y_1 = __________

Y_2 = __________

Y_3 = __________

Y_{12} = __________

Y_{13} = __________

Y_{23} = __________

▷ Assume that l_4 treats differently p_1 from p_2/p_3 traffic: l_4 is always non-lossy for p_1 traffic, but lossy for p_2/p_3 traffic with probability 0.5. All the other links are always non-lossy. We monitor p_1, p_2, p_3, p_{12}, p_{13}, p_{23} and observe that p_1 is always non-lossy, while each of the rest is lossy with probability 0.5. Show that tomography identifies l_4 as non-neutral. (2 Points)

▷ Assume now that l_4 is lossy only for p_3 traffic, and all other links are always non-lossy. Can tomography conclusively identify l_4 as non-neutral? Briefly explain why. (2 Points)

leave blank

leave blank

Task 5: Network Security 11 Points

a) Applying STRIDE (3 Points)

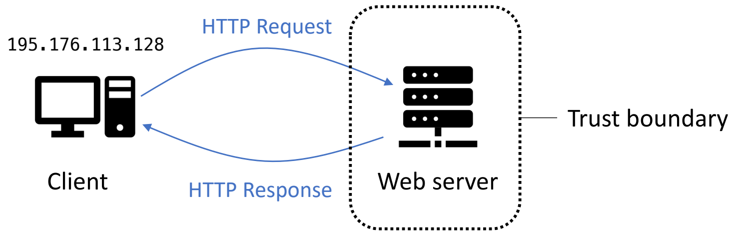

Figure 10: A simple networked system.

The network in Figure 10 consists of a client (with IP 195.176.113.128) and a web server, connected over HTTP. Assume that you are the operator of the server (i.e., the server is within your trust boundary and the client is untrusted). Apply the STRIDE[1] methodology to identify **three** different threats to the server through its communication with the client. Each threat must belong to a different STRIDE category (e.g., only one threat can rely on spoofing). For each identified threat, specify its category and describe a possible mitigation.

Example (exclude it from your answers):

- *Threat:* The client may send an offensive HTTP request to the server and deny sending it during the post-attack analysis.
- *Category:* Repudiation.
- *Mitigation:* The server operator may require the client to sign the HTTP requests and send the signatures along with the requests.

Threat 1: __________

Category 1: __________

Mitigation 1: __________

Threat 2: __________

Category 2: __________

Mitigation 2: __________

[1] Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege

Threat 3: ______

Category 3: ______

Mitigation 3: ______

b) Revisiting *ACC-Turbo* (4 Points)

leave blank

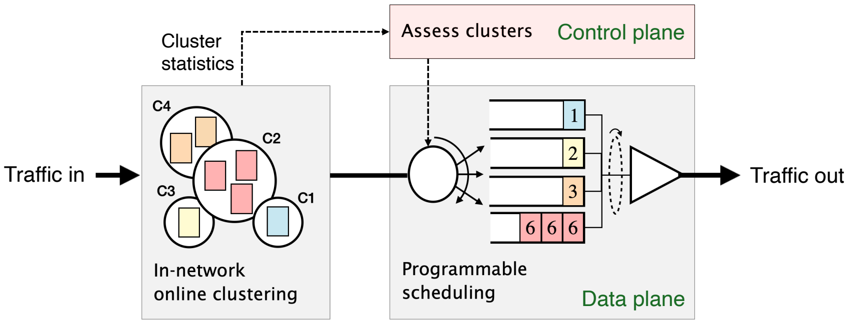

Figure 11: Architecture of *ACC-Turbo*.

ACC-Turbo defends against pulse-wave DDoS attacks by clustering incoming packets based on their similarity and deprioritizing malicious clusters (cf. Figure 11).

(i) Name two reasons why existing DDoS defenses, which rely on a central controller to mitigate attacks, could be vulnerable to pulse-wave DDoS attacks. (1 Point)

Reason 1: ______

Reason 2: ______

(ii) What metric does *ACC-Turbo* use to assess packet similarity? (1 Point)

(iii) Explain one advantage of deprioritizing potentially-malicious traffic as opposed to just dropping it. (1 Point)

(iv) Explain how an attacker could trick *ACC-Turbo* into allocating a high bandwidth share to malicious traffic. (1 Point)

c) Defending DDoS with `iptables` (4 Points)

leave blank

```
    B  1 2.889091491   1.0.0.2 -> 1.0.0.3   TCP 1429 6491 -> 6491 Len=243
    B  2 2.894215096   1.0.0.2 -> 1.0.0.3   TCP 1429 6491 ->  587 Len=652
    M  3 2.899809065   1.0.0.2 -> 1.0.0.3   UDP 1429   22 -> 2890 Len=932
    M  4 2.905407385   1.0.0.4 -> 1.0.0.3   TCP 1429 2191 -> 6490 Len=193
    B  5 2.910978684   1.0.0.3 -> 1.0.0.2   TCP 1429 6490 -> 6491 Len=277
    M  6 2.916530913   1.0.0.6 -> 1.0.0.3   UDP 1429 6490 -> 6491 Len=487
```

Listing 2: A Wireshark trace

You are managing a web server with IP address 1.0.0.3. Listing 2 describes the packets that you observe either incoming or outgoing the server. Note that packets are already classified as benign (B) or malicious (M).

For the following **three** `iptables` rule sets:

- Identify the packets that will get dropped if you deploy the rules.
- Describe one way the attacker could bypass the rules, and prevent their malicious packets from being dropped.

(i) `iptables -A INPUT --dport !6491 -j DROP` (1 Point)

Dropped packets: __________

Bypassing strategy: __________

(ii) `iptables -A INPUT -p UDP -j DROP` (1 Point)

Dropped packets: ______________________

Bypassing strategy: ______________________

____

____

(iii)
```
iptables -A INPUT --src 1.0.0.2 -j DROP
iptables -A INPUT --src 1.0.0.3 -j DROP
iptables -A INPUT --src 1.0.0.6 -j DROP
```
(1 Point)

Dropped packets: ______________________

Bypassing strategy: ______________________

____

____

(iv) Now, write a set of up to four `iptables` rules that drop *all and only* malicious packets. Explain your solution briefly. (1 Point)

Rules:

`iptables -A INPUT` ______________________

`iptables -A INPUT` ______________________

`iptables -A INPUT` ______________________

`iptables -A INPUT` ______________________

Explanation: ______________________

____

____

____

leave blank

leave blank

Task 6: Transport Protocols 21 Points

a) Modeling and Design (4 Points)

(i) How do TCP RENO and CUBIC model the network, how do they update this model based on observations, and how do they control their sending rates? (2 Points)

Model:

Observations:

Control:

(ii) In the modern Internet, many flows are small and short: they transmit only a small amount of bytes and are only active for a short time. Explain the impact this has on congestion control with TCP RENO or CUBIC. (1 Point)

(iii) Explain why network address translation (NAT) is a challenge for Multipath TCP, and how this challenge is solved. (1 Point)

b) Buffering **(11 Points)**

leave blank

In this task, you will analyze the impact of congestion control on network queues.

Consider the network shown below with a single sender and receiver connected through a router with link rates $r_{in} = 100\,\frac{packets}{ms}$ and $r_{out} = 24\,\frac{packets}{ms}$. The rates are defined in terms of data packets, and you can assume that ACKs are infinitely small and transmitted instantly. The router has a buffer b for data packets with a capacity of $1200\,packets$.

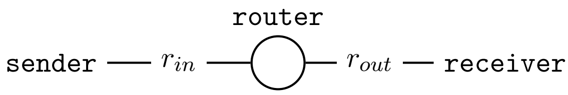

The only traffic in this network is sent from the sender to the receiver. The sender is using a time-based AIMD congestion control algorithm with packet loss as its congestion signal.

(i) Name an example of a real congestion control algorithm that updates its *cwnd* as a function of time since the last loss, and an example of an algorithm that updates its *cwnd* after every RTT. Describe a disadvantage of the RTT-based approach. (1 Point)

Time-based algorithm: ______________________________

RTT-based algorithm: ______________________________

RTT-based disadvantage: ______________________________

(ii) Compute RTT_{max}, i.e., compute the maximum time between sending a (not lost) data packet and receiving its ACK in this network. Explain your steps. (1 Point)

The sender algorithm updates its congestion window (*cwnd*) and sending rate R as follows:

- When loss is detected, the *cwnd* is halved and a new *congestion epoch* begins.
- The sender detects loss **instantly** at time t if the buffer is: *(i)* full ($b(t) = 1200$) *; and (ii)* still filling ($\frac{db}{dt}(t) > 0$). You do not need to consider duplicate ACKs or timeouts.
- During a *congestion epoch*, use the following equation to determine the congestion window at the t since the beginning of the epoch, where $cwnd_0$ is the *cwnd* at epoch start:

$$cwnd(t) = cwnd_0 + t \cdot 2\,\frac{packets}{ms}$$

- During a *congestion epoch*, use the following linear approximation for the sending rate:

$$R(t) = \frac{cwnd(t)}{50\,ms}$$

- Assume the algorithm to be in steady state, oscillating around its operating point. At the beginning of a *congestion epoch*, the *cwnd* is 800, and at the end, it is 1600.

(iii) Analyze the *change* in buffer size over time. In the plot below, draw $\frac{db}{dt}$ between $t = 0ms$ and $t = 500ms$. Assume that at $t = 0s$, a new *congestion epoch* has just started and the buffer is full at $b = 1200\,packets$ (otherwise we would not have detected loss).

Hint: At which rates does the buffer fill and drain, and how long is a congestion epoch?

(6 Points)

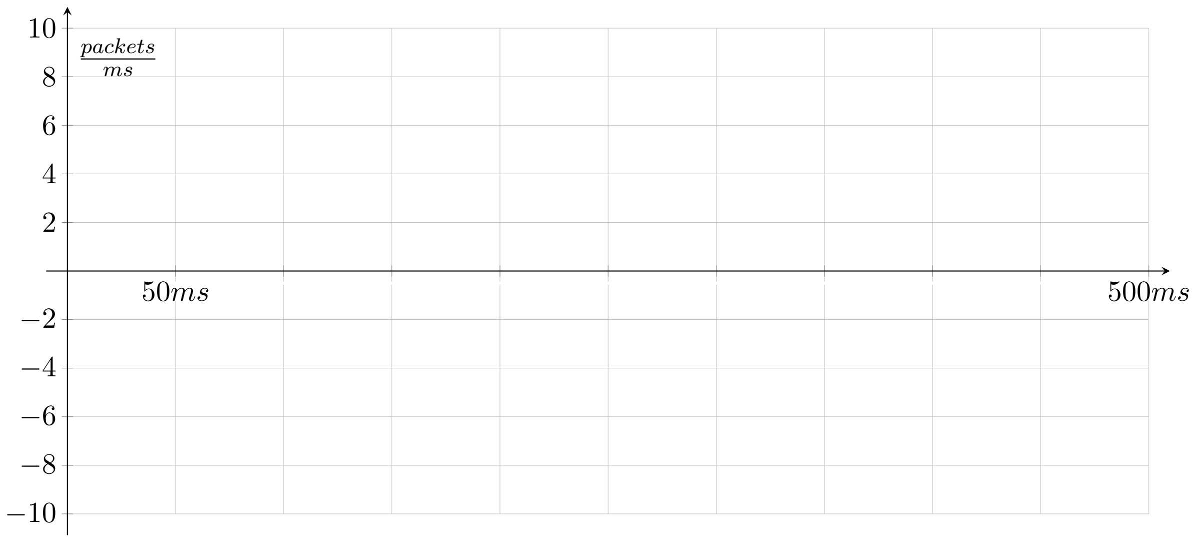

Notes (not graded): ______________________________

(iv) Is the buffer ever empty? If yes, give a time t at which this happens. If not, give the minimum buffer occupancy b_{min}. (2 Points)

Ever empty? ____________________

Computation: ____________________

(v) Consider a second flow on the same sender that uses a *delay-based* congestion control algorithm instead of a loss-based one: a new *congestion epoch* starts if the observed RTT exceeds a threshold RTT_{thres}.

Is it possible to choose a $RTT_{thres} < RTT_{max}$ such this both the loss-based and delay-based algorithm share the bottleneck bandwidth fairly?

Provide an example or explain why it is impossible.

Hint: Which algorithm reacts to congestion first? (1 Point)

c) BBR Control Space (6 Points)

leave blank

In this task, you will discuss why BBR uses *both* a congestion window and a pacing rate. Consider the network shown below, where two flows transmit data using a simplified version of BBR. Router buffers have a size of **200 packets**.

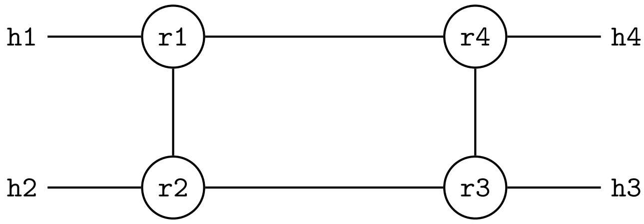

At time t = 0s, the two flows follow disjoined paths.
Flow one follows the path $h1, r1, r4, h4$; flow two follows the path $h2, r2, r3, h3$.

At time t = 5s, flow one is rerouted and now follows the path $h1, r1, r2, r3, r4, h4$; both flows now share the bottleneck link $r2, r3$.

To simplify the state measurements, assume that the bottleneck bandwidth $BtlBw$ and propagation delay $RTTProp$ are known exactly, but changes are detected with $100ms$ delay. The table below shows the state as well as the computed pacing rate and congestion window for both flows. The left side shows $t < 5.1s$ and the right table shows $t \geq 5.1s$, reflecting the $100ms$ delay in detecting the change after the top flow has been rerouted at $t = 5s$.

			$t < 5.1s$	$t \geq 5.1s$
Measured	$BtlBw$	$\frac{packets}{s}$	1000	500
	$RTTProp$	ms	120	120
Computed	$PacingRate$	$\frac{packets}{s}$	1000	500
	$Cwnd$	$packets$	120	60

(i) What if BBR would *only use the provided congestion window* without a pacing rate? Explain a downside and give a concrete numeric example in the scenario above in terms of reduced throughput and/or additional delays and/or packet losses. (3 Points)

(ii) What if BBR would *only use the provided pacing rate* without a congestion window? Explain a downside and give a concrete numeric example in the scenario above in terms of reduced throughput and/or additional delays and/or packet losses. (3 Points)

leave blank

Task 7: Sustainable networking **11 Points**

(i) Among the following, mark as *true* the energy units and as *false* the others. (1 Point)

- true ☐ false ☐ kW/h
- true ☐ false ☐ J
- true ☐ false ☐ gCO2e
- true ☐ false ☐ J/s
- true ☐ false ☐ kWh
- true ☐ false ☐ W

(ii) The Greenhouse Gas (GHG) protocol classes emissions into three scopes. Explain briefly which type of emissions fall into each scope. (3 Points)

Scope 1: __________

Scope 2: __________

Scope 3: __________

(iii) Let us assume that the carbon footprint of streaming Netflix in Switzerland is estimated to be 55gCO2e per hour of streaming. How much would carbon emissions be reduced if you stream only one hour and go to bed instead of streaming two hours? Explain your answer. (2 Points)

(iv) Let us assume a new-generation router consumes half the energy of its older generation. From a sustainability standpoint, why is it not always better to upgrade to the new generation? (2 Points)

(v) Name two energy-saving techniques that can theoretically save tens of % of energy in lightly-loaded networks. (2 Points)

Technique 1: __________

Technique 2: __________

(vi) In practice, what limits the efficiency of these techniques today? (2 Points)