Solutions

Marking: **M1** a correct method, **A1** the correct answer, **B1** an independent mark.

Each answer shows the steps, then the **result**.

2.1

`10.1.2.200` → C

`10.1.3.1` → B

`11.0.0.1` → D

2.2

via Y $1 + 6 = 7$; via Z $3 + 2 = 5$

5, next hop Z

2.3

distance vector: only its neighbours' distances to each destination

link state: the whole map of links and costs

2.4

B lost C, but A still advertises "C at 2" (via B)

B believes it, sets C at 3 via A; A then sets 4 via B, and so on

with split horizon A does not tell B about routes it learned from B, so B sees C as unreachable at once

3.1

(a) B 1; C 3 (via B); D 6 (via C); E 7 (via D)

(b) B for every destination

(c) D: A–B–D = **7**; E: $7 + 1 = \mathbf{8}$

3.2

(a) each network decides by its own business and security policy; it may not want to carry others' traffic for free, even on a short path

(b) the customer route: traffic sent that way earns money, while the peer link is free and the ISP would pay for transit

(c) a network can announce someone else's addresses (a hijack), and traffic is pulled to it

3.3

(a) both `0.0.0.0/0` routes are in the table, ISP 1 with the lower cost; when ISP 1 is declared dead its route is removed, and the ISP 2 route is used

(b) up to **40 s**; $3 \times 0.3 = \mathbf{0.9\ s}$

exercises.lol exercises.com.cn

(c) per-packet splitting sends one connection's packets over paths of different delays; they arrive out of order and TCP slows down; hashing each connection to one link keeps its packets in order

(d) fast failure detection (0.9 s) with ISP 1 preferred, and per-connection balancing only if ISP 2's capacity is needed: a 40 s outage would drop calls, while sub-second failover is barely noticed