Non-CLRS4 Candidate Questions for the AD3 Exam

Prepared by Pierre Flener

Version of 19th June 2025

1 Dubious Alliance Detection

Your cybersecurity consultancy has been contacted by a firm that records all the internet accesses by the users of its computer system. Assuming that each user accesses at most one IP (internet protocol) address per second, a log file records, for each user u and second s, the value $IP(u, s)$ that is the IP address ($\perp$ if none) accessed by u during s.

The firm's computer system was used to launch a complex attack on some remote sites. The attack was carried out by accessing t distinct IP addresses over t consecutive seconds: it accessed address a_1 during second 1, and so on, up to address a_t during second t. Inspection of the log file revealed that no single user accessed all the involved IP addresses at the appropriate times: there is no user u such that $IP(u, s) = a_s$ for each second s in $1\,..\,t$.

The firm now requests your help to detect if a small alliance of its users might have jointly carried out the attack. A subset U of users forms a ***dubious alliance*** if, for each second s in $1\,..\,t$, there is at least one user u in U for which $IP(u, s) = a_s$. The ***dubious alliance detection*** problem asks: given a number n and the collection of all values $IP(u, s)$, is there a dubious alliance of size at most n?

Perform the following sequence of tasks:

1. Prove that this problem is NP-complete, by a *single* reduction, *directly* from the decision version of the ***vertex-cover problem***, whose optimisation version asks to find a vertex cover of minimum size in a given undirected graph $G = (V, E)$, that is a minimum-size subset $V' \subseteq V$ such that if $(u, v) \in E$, then either $u \in V'$ or $v \in V'$ (or both).

2. How would you use a decision algorithm for the dubious alliance detection problem in order to help the firm find the smallest size of a dubious alliance?

3. What would you say to the firm before designing and running your algorithm of task b (and how much would you charge for the project)?

2 Project Course Design

Your timetabling consultancy has been contacted by a university to plan a project course, with one meeting per week, namely ℓ lectures by outside speakers followed by p project sessions.

There are n candidate speakers, with a subset L_i thereof available during week i in $1\,..\,\ell$. Each project session j in $1\,..\,p$ requires having seen the background material of at least one of a subset P_j of the candidate speakers. The ***project course design*** problem asks: given these sets, is it possible to select exactly one candidate speaker within each L_i so that the students will have seen at least one of the speakers in each P_j?

For example, consider $\ell = 2$ lectures, $p = 3$ project sessions, and $n = 4$ candidate speakers, called a, b, c, d, with $L_1 = \{a, b, c\}$, $L_2 = \{a, d\}$, $P_1 = \{b, c\}$, $P_2 = \{a, b, d\}$, and $P_3 = \{c, d\}$. Then the answer is 'yes', since one can select speaker b for week 1 and speaker d for week 2.

Perform the following sequence of tasks:

1. Prove that this problem is NP-complete, by a *single* reduction, *directly* from either the decision version of the ***vertex-cover problem***, which asks to find a vertex cover of minimum size in a given undirected graph $G = (V, E)$, that is a minimum-size subset $V' \subseteq V$ such that if $(u, v) \in E$, then either $u \in V'$ or $v \in V'$ (or both), or 3-***CNF satisfiability***, which asks whether a conjunction of clauses, each of exactly three distinct literals, is satisfiable (where a *literal* is an occurrence of a Boolean variable or its negation).

2. How would you use an algorithm for this decision problem in order to fulfil the university's actual request?

3. What would you say to the university before designing and running that planning algorithm (and how much would you charge for the project)?