Leaf sequencing

Intensity Modulated Radio Therapy (IMRT) is a procedure used to kill cancer cells. In classical radiotherapy, a single high-intensity radiation beam is focused on the cancer cells and kills them. The beam, however, also damages all the (sane) tissues it passes through. On the other hand, IMRT uses many different low-intensity beams, which only converge in the exact spot where the cancer cells are living. In this way the collateral damage for the patient is minimised.

The entire IMRT procedure is optimised using Operational Research tools, and is usually sequenced in three phases:

- Beam Angle Optimisation (BAO), used to decide from which angles the various beam will be emitted;
- Fluence Map Optimisation (FMO), used to determine the intensities of the beams;
- Leaf Sequencing (LS), used to schedule the beam-emitting machines so that they overall emit beams at the intensities required by FMO.

Here, we will focus on Leaf Sequencing. We represent the target area as a bidimensional matrix I. The value I_{ij} represents the intensity that we would like to convey to the cells in position (i, j).

9	4
6	1

According to the matrix in this example, the cells in position $(0, 0)$ needs to be hit with the beams for 9 minutes; the cells in position $(0, 1)$ need to be hit for 4 minutes; analogously, the cellz in position $(1, 0)$ need to be hit for 6 minutes, and those in position $(1, 1)$ for 1 minute.

This intensity can be obtained, for example, by hitting each area in sequence. For example, one could perform 9 minutes of radiotherapy in configuration

1	0
0	0

followed by 4 minutes in configuration

0	1
0	0

and by 6 minutes in configuration

0	0
1	0

and finally 1 minute in configuration

0	0
0	1

Such an approach would force the patient to be under the beams for a total of $9 + 4 + 6 + 1 = 20$ minutes. If we also consider the time T (for example $T = 2$) that is needed to switch the beam-emitting machine from one configuration to the next, the total time used for such a therapy session would be $20 + 2 \cdot 3 = 26$ minutes.

The machines, however, are capable of irradiating more than one area at a time. So, another schedule that can be used to perform the same therapy is the following: 1 minute in configuration

1	1
1	1

followed by 3 minutes in configuration

1	1
0	0

and finally by 5 minutes in configuration

1	0
1	0

This approach gives a radiation time of $1 + 3 + 5 = 9$ minutes, and a total time of $9 + 2 \cdot 2 = 13$ minutes.

The objective of the Leaf Sequencing problem, is then to decide a sequence of machine configurations, such that the total makespan of the procedure is minimised, and all areas are hit with exactly the intensity they need.

Let C be the set of configurations available for a machine; for each configuration $c \in C$ and each position (i, j) in the target area, the parameter $\delta_{cij} \in \{0, 1\}$ will have value 1 if position (i, j) is hit when using configuration c, or 0 otherwise. For example, configuration

$c =$

1	0
0	1

will give $\delta_{c00} = \delta_{c11} = 1$ and $\delta_{c01} = \delta_{c10} = 0$. In the following, we assume that matrix I is an $n \times m$ matrix.

Consider then variables $x_c \in \mathbb{N}$ for each configuration $c \in C$; the value of x_c will represent the number of minutes for which configuration c is used. Also consider variables $y_c \in \{0, 1\}$ for each configuration $c \in C$; y_c will have value 1 if configuration c is used, or value 0 otherwise.

It is easy to see that a model for the Leaf Sequencing problem is the following:

$$\min \quad \sum_{c \in C} x_c + T \cdot \left(\sum_{c \in C} y_c - 1 \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{c \in C} \delta_{cij} x_c = I_{ij} \qquad \forall i = 1, \ldots, n \ \forall j = 1, \ldots, m \tag{2}$$

$$x_c \leq M y_c \qquad \forall c \in C \tag{3}$$

$$x_c \in \mathbb{N} \qquad \forall c \in C \tag{4}$$

$$y_c \in \{0, 1\} \qquad \forall c \in C \tag{5}$$

Where M is a sufficiently big number.