A one-dimensional cutting stock problem with alternative lengths

Cutting stock problems are integer linear programming problems and a well known class of optimisation problems since its first mention in 1939. They deal with the problem how to cut a given input material in order to fulfil purchase orders so that the wasted input material will be minimised. We can distinguish one-, two- and three-dimensional cutting problems. But even the one-dimensional variant is already NP-hard. The general one-dimensional problem can be characterised as follows: Given an input material with a fixed length and a set of purchase orders. Each order requires a certain number of pieces with the same length but two lengths of different orders may also be different. Then, the input material that is used to handle the orders is to be minimised. Aggravating, the orders may have to be handled until a given period of time, the capacity of cutting machines is limited and a single cut needs a certain period of time. Several algorithms lead to very good but not optimal results like the Delayed Column Generation approach of Gilmore and Gomory published in the early 1960s. The main problem of these algorithms is that they don't handle integrality but result in a solution that contains fractions.

Cutting stock problems arise in many industries. In the steel industry a typical production process consists in rolling tubes out of solid bars of steel. The customer orders a number of tubes with a desired length and diameter. These tubes are the output material of a complex production process. With the length and the diameter as input parameters the needed lengths of input material can be calculated. Due to several degrees of freedom within the production process it is possible to produce the same output with different lengths of input material. Thus, we get alternative cutting lengths that can be used to fulfil customer orders. For example, it is possible to fulfil customer order A with 10 pieces of input length X or with 14 pieces of input length Y.

Now, this variant of the general problem cannot be solved by the use of commonly known algorithms because they always are limited to one length per customer order. In this paper we will present an algorithm that solves the specialised problem. First of all, we will describe the formal model. Because the formal model is NP-hard we cannot optimise a given instance of the problem in acceptable time. Therefore, we introduce a heuristic that leads to good results in an acceptable computational time. First of all, the heuristic is based on the well known Delayed Column Generation approach of Gilmore and Gomory that has to be adapted to the new requirements. Because the Delayed Column Generation approach doesn't handle integrality a post optimisation will be introduced that will lead to satisfying results.

Due to practical relevance, we can ease some side constraints: We can distinguish one main variant and up to five side variants so that the number of alternative cutting lengths can be limited to five. Because of limitations in the production process we are only allowed to use one side variant beneath the main variant. That is because if we are using too much alternatives the production process itself is getting too expensive.
These simplifications do not reduce the complexity of the problem but they lead to a reduction of calculation possibilities so that the latter heuristic will work quicker than without the simplifications.