Online Companion

Appendix C: Additive Model

One can consider an alternative class of additive demand models, in the sense that the price and vehicle effects are additively separable. Such models can be expressed as:

$$d_t(p_t, \{x_{vt}\}_{v\in V}) = h_t^A(p_t) + \sum_{v\in V} B_{vt}x_{vt}. \tag{7}$$

The function $h_t^A(p_t)$ represents the effect of the price vector p_t on demand, and the boost parameter $B_{vt} \geq 0$ corresponds to the absolute increase in demand at time t when vehicle v is used. For example, if $B_{vt} = 65$, assigning vehicle v at time t yields an additional 65 units in sales, relative to the case where this vehicle is not used.

The objective function of the promotion vehicle scheduling problem for the additive model (7) asks to maximize total profit over the selling season:

$$\sum_{t=1}^{T}(p_t - c_t)\cdot d_t(p_t, \{x_{vt}\}_{v\in V}) = \sum_{t=1}^{T}(p_t - c_t)\cdot h_t^A(p_t) + \sum_{t=1}^{T}\alpha_t^A \sum_{v\in V} B_{vt}x_{vt}, \tag{8}$$

where $\alpha_t^A = p_t - c_t$. In this formulation, the first term on the right hand side does not affect the vehicle optimization problem. In addition, α_t^A corresponds to the effect of the price on profits at time t and is simply equal to the profit margin at time t. Since all prices and costs are assumed to be given, α_t^A is a given quantity as well.

The objective function (8), along with the linear constraints specified in Section 3, can easily be expressed as a bipartite b-matching problem, and therefore, can be solved efficiently by various methods. Consequently, the problem is scalable, with running times in milliseconds even for large instances. Nevertheless, as shown in Section 4.2, the additive demand model does not provide a good fit to the data and as a result, is not an appropriate model in this context. The additive model suffers from a scale independence due to assuming an absolute boost independent of the number of sales. In particular, different data points can have a wide range of sales units, such that the additive parameter for the promotion vehicle might be hard to estimate. Thus, a relative term such as in the multiplicative model seems to be more suitable. We refer the reader to Section 4.2, where we examine the prediction accuracy of both models and conclude that the multiplicative model is more suitable for our purposes.

Appendix D: PTAS for Uniform Boosts

In what follows, we devise a polynomial-time approximation scheme for the setting where all vehicles boosts are uniform, and where the base profits of all time periods are also uniform. In other words, the underlying assumption is that $B_{vt} = B$ for any vehicle v and time period t, and in addition, $\alpha_1 = \cdots = \alpha_T = 1$, after normalization. Specifically, for any accuracy level $\epsilon \in (0, 1/2]$, we show how to efficiently construct $\tilde{O}((T/\epsilon)^{O(1/\epsilon^2)})$ linear assignment problems, such that the optimal solution to at least one of them guarantees a $(1-\epsilon)$-approximation for the original problem. (To avoid cumbersome expressions, we use the notation $\tilde{O}(f(n)) = f(n)\cdot \mathrm{polylog}(n)$.)

It is worth mentioning that our algorithm is also useful in the general case, where the vehicle boosts take arbitrary values, say in $[B_{\min}, B_{\max}]$, and similarly, the base profits take arbitrary values in $[\alpha_{\min}, \alpha_{\max}]$. It is not difficult to verify that, by replacing each of the boosts B_{vt} by $B_{\min}$ and each of the base profits α_t by $\alpha_{\min}$, we scale down the optimal objective value by a factor of at most

$$\left(\frac{B_{\max}}{B_{\min}} \right)^{\Delta} \cdot \frac{\alpha_{\max}}{\alpha_{\min}} \ ,$$

where $\Delta = \max_t L_t$ is the maximum capacity of any time period. We can then apply our approximation scheme, losing an additional factor of $1-\epsilon$, and obtaining a solution whose objective value with respect to the original boosts and base profits can only improve.

Profit classes. Let x^* be an optimal solution to (P), when the latter is specialized according to our assumptions above, that is,

$$\begin{array}{lll}
(P) \ \max & \displaystyle\sum_{t=1}^{T} B^{\sum_{v \in V} x_{vt}} & \\
(C_1) & \displaystyle\sum_{t=1}^{T} x_{vt} \leq C_v & \forall v \in V \\
(C_2) & \displaystyle\sum_{v \in V} x_{vt} \leq L_t & \forall t \in [T] \\
(C_3) & x_{vt} \in \{0,1\} & \forall v \in V, t \in [T]
\end{array}$$

We assume from this point on that the value of $\mathrm{OPT} = \sum_{t=1}^{T} B^{\sum_{v \in V} x^*_{vt}}$ is known up to an ϵ-factor, meaning that we have an estimate $\widetilde{\mathrm{OPT}} \in [(1-\epsilon) \cdot \mathrm{OPT}, \mathrm{OPT}]$. This assumption can be enforced by employing the greedy algorithm (see Section 5.2), which is guaranteed to obtain an objective value $\mathcal{V}$ satisfying $\mathcal{V} \leq \mathrm{OPT} \leq (\Delta+1)\mathcal{V}$. We can then test all powers of $1+\epsilon$ within the interval $[\mathcal{V}, (\Delta+1)\mathcal{V}]$ as candidate values for $\widetilde{\mathrm{OPT}}$, run our algorithm with each value, and finally return the best solution found over all candidates.

Now, for purposes of analysis, consider a partition of the time periods $1, \ldots, T$ into classes, based on their associated profits with respect to x^*. Specifically, the class $\mathcal{C}_1$ consists of periods t with $B^{\sum_{v \in V} x^*_{vt}} \geq (1-\epsilon) \cdot \widetilde{\mathrm{OPT}}$. Then, for $2 \leq \ell \leq L$, each class $\mathcal{C}_\ell$ corresponds to periods with

$$B^{\sum_{v \in V} x^*_{vt}} \in \left[(1-\epsilon)^{\ell} \cdot \widetilde{\mathrm{OPT}}, (1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}} \right) \ ,$$

where $L = \lceil \log_{1+\epsilon}(T/\epsilon) \rceil$. Finally, we define the last class C_{small} to consist of periods with $B^{\sum_{v \in V} x^*_{vt}} \in [0, (1-\epsilon)^L \cdot \widetilde{\mathrm{OPT}})$. By this definition, the total contribution of class C_{small} to the objective value of x^* is at most $\epsilon \cdot \widetilde{\mathrm{OPT}}$, since

$$\begin{aligned}
\sum_{t \in C_{\mathrm{small}}} B^{\sum_{v \in V} x^*_{vt}} &\leq |C_{\mathrm{small}}| \cdot (1-\epsilon)^L \cdot \widetilde{\mathrm{OPT}} \\
&\leq T \cdot \left(\frac{1}{1+\epsilon} \right)^{L} \cdot \widetilde{\mathrm{OPT}} \\
&\leq T \cdot \left(\frac{1}{1+\epsilon} \right)^{\log_{1+\epsilon}(T/\epsilon)} \cdot \widetilde{\mathrm{OPT}} \\
&= \epsilon \cdot \widetilde{\mathrm{OPT}} \ .
\end{aligned} \tag{9}$$

Constructing the linear assignment problem. For every $1 \leq \ell \leq L$, let R^*_ℓ be the total contribution of the time periods in $\mathcal{C}_\ell$ to the optimal objective function, i.e., $R^*_\ell = \sum_{t \in \mathcal{C}_\ell} B^{\sum_{v \in V} x^*_{vt}}$. Now suppose we have at our possession an integer-valued vector $(k_1, \ldots, k_L)$ that satisfies for every $1 \leq \ell \leq L$

$$k_\ell \cdot \frac{\epsilon \cdot \widetilde{\mathrm{OPT}}}{L} \leq R^*_\ell \leq (k_\ell + 1) \cdot \frac{\epsilon \cdot \widetilde{\mathrm{OPT}}}{L} \ , \tag{10}$$

letting $\tilde{R}_\ell = k_\ell \cdot \epsilon \cdot \widetilde{\mathrm{OPT}}/L$ be the resulting lower bound on R^*_ℓ. Obviously, the vector $(k_1, \ldots, k_L)$ is generally unknown, and we explain in the sequel how to enumerate over only polynomially-many vectors. (Note that a simple enumeration, where each of the coordinates k_ℓ takes one of the values $0, \ldots, T$, makes the total number of options $T^{O(L)} = T^{O((1/\epsilon) \cdot \log(T/\epsilon))}$, which is not polynomial in T.) However, for the time being, with this vector at hand, it follows that each class $\mathcal{C}_\ell$ contains at least $\lceil \tilde{R}_\ell / ((1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}}) \rceil$ time periods, since

$$\tilde{R}_\ell \leq R^*_\ell = \sum_{t \in \mathcal{C}_\ell} B^{\sum_{v \in V} x^*_{vt}} \leq |\mathcal{C}_\ell| \cdot (1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}} \ .$$

With this lower bound on $|\mathcal{C}_\ell|$, we create a (feasibility) linear assignment problem, with vehicles on the left and with time periods on the right. Each vehicle v has a supply of C_v units. On the other hand, for each class $\mathcal{C}_\ell$, we take a separate subset of $\lceil \tilde{R}_\ell / ((1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}}) \rceil$ time periods, and associate each of them with a demand of m_ℓ units. Here, m_ℓ stands for the minimal number of vehicles needed in order to obtain a per-period profit within the interval corresponding to class $\mathcal{C}_\ell$. That is, m_ℓ is the minimal integer for which

$$B^{m_\ell} \in \left[(1-\epsilon)^\ell \cdot \widetilde{\mathrm{OPT}}, (1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}} \right) \ . \tag{11}$$

Approximation guarantee. By the preceding discussion, the optimal solution x^* is in particular feasible for this assignment problem. As a result, by computing any feasible solution, we obtain an overall profit of at least

$$\begin{aligned}
\sum_{\ell=1}^{L} \left\lceil \frac{\tilde{R}_\ell}{(1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}}} \right\rceil \cdot B^{m_\ell} &\geq \sum_{\ell=1}^{L} \frac{\tilde{R}_\ell}{(1-\epsilon)^{\ell-1} \cdot \widetilde{\mathrm{OPT}}} \cdot (1-\epsilon)^\ell \cdot \widetilde{\mathrm{OPT}} \\
&= (1-\epsilon) \cdot \sum_{\ell=1}^{L} \tilde{R}_\ell \\
&\geq (1-\epsilon) \cdot \sum_{\ell=1}^{L} \left(R^*_\ell - \frac{\epsilon \cdot \widetilde{\mathrm{OPT}}}{L} \right) \\
&= (1-\epsilon) \cdot \left(\sum_{\ell=1}^{L} R^*_\ell - \epsilon \cdot \widetilde{\mathrm{OPT}} \right) \\
&\geq (1-\epsilon) \cdot (1-2\epsilon) \cdot \widetilde{\mathrm{OPT}} \\
&\geq (1-2\epsilon)^3 \cdot \mathrm{OPT} \ .
\end{aligned}$$

The first inequality follows from the bounds on B^{m_ℓ} given in (11). The second inequality follows from the relation between $\tilde{R}_\ell$ and R^*_ℓ in (10). The third inequality holds since $\sum_{\ell=1}^{L} R^*_\ell \geq (1-\epsilon) \cdot \widetilde{\mathrm{OPT}}$, by the upper bound on the contribution of C_{small} in (9). The last inequality holds since $\widetilde{\mathrm{OPT}} \geq (1-\epsilon) \cdot \mathrm{OPT}$.

Efficient enumeration. It remains to show that the number of vectors $(k_1, \ldots, k_L)$ to be considered is polynomial in the input size, for any fixed $\epsilon \in (0, 1/2]$. For this purpose, note that by definition of k_ℓ,

$$\frac{\widetilde{\mathrm{OPT}}}{1-\epsilon} \geq \mathrm{OPT} \geq \sum_{\ell=1}^{L} R^*_\ell \geq \frac{\epsilon \cdot \widetilde{\mathrm{OPT}}}{L} \cdot \sum_{\ell=1}^{L} k_\ell \ ,$$

implying that $\sum_{\ell=1}^{L} k_\ell \leq L/(\epsilon(1-\epsilon)) \leq 2L/\epsilon$. Basic counting arguments show that the number of integer solutions to this inequality is $O(e^{O(L/\epsilon)}) = \tilde{O}((T/\epsilon)^{O(1/\epsilon^2)})$.