Technische Universität Berlin

FAKULTÄT II
MATHEMATIK UND
NATURWISSENSCHAFTEN
Institut für Mathematik

The Periodic Assignment Problem (PAP) May Be Solved Greedily

by

Christian Liebchen

No. 716/2001

The Periodic Assignment Problem (PAP) May Be Solved Greedily

Christian Liebchen*

November 12, 2001

Abstract

Many public transportation companies operate their networks periodically. One major step in their planning process is to construct a periodic timetable for one abstract period, independently from times during the day. In this paper we show that we may evaluate a periodic timetable very quickly with the number of vehicles required to operate it. This is due to the fact that the Periodic Assignment Problem (PAP) can be solved by a greedy approach. It helps us, at least within a genetic algorithm, to cope with the quadratic objective function in the problem of finding a periodic timetable requiring as few vehicles as possible.

Keywords: Periodic Assignments, Periodic Timetabling, Vehicle Scheduling in Public Transport

1 Introduction

During the planning process for public transportation companies, two major tasks are constructing a good timetable and determining a good vehicle schedule. It is clear that the possibilities for the vehicle schedule depend heavily on the timetable that has to be operated.

As many companies serve their network periodically, in a first step, the timetable is only planned for one abstract period of e.g. 60 minutes, independently from concrete times during the day. In general, there is one abstract plan for the peak hours, one for the evenings, one for the week-ends, and possibly some others.

The problem of scheduling vehicles such that for one day they operate a given timetable has been extensively investigated by Löbel[2], even for the multi-depot case. Decisions that have to be made include at which time vehicles should leave and enter the depots, how they should perform the junction from the peak-hours'

*Author's address: Technische Universität Berlin, Institut für Mathematik, Sekr. MA 6-1, Straße des 17. Juni 136, D-10623 Berlin, Germany. Email: liebchen@math.tu-berlin.de

plan to the evening's plan, and, last but not least, how to serve the abstract plan e.g. for the evenings, which may be operated from 7 p.m. until 11 p.m.

The last point is the only one that has the same horizon as the periodic timetable optimization task, because everything repeats periodically and therefore can be considered as one abstract period. Hence, it would be nice to construct a periodic timetable requiring as few vehicles as possible. But this leads to a mixed integer quadratic program with linear constraints, see Liebchen and Peeters[1]. Remains the problem of at least quickly evaluating a periodic timetable – besides passengers' waiting times – by the number of vehicles required to serve the timetable as well.

The purpose of this paper is to show that within every terminus station the Periodic Assignment Problem (PAP) can be solved by assigning every incoming event greedily – in a FIFO manner – to its nearest unmatched outgoing event. This permits a quick calculation of the resources required to perform a periodic timetable. Notice that for the general assignment problem locally optimal decisions may prevent global optimality. Figure 1 shows that choosing

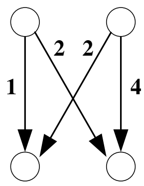

Figure 1: Local vs. global optimality in the general assignment problem

the minimally weighted arc does not yield a minimal assignment.

This paper is structured as follows: First, a point of time will be identified such that for any later point of time within the abstract period there are at least as many arrivals as departures. Starting at this point of time, the optimality of the FIFO strategy will be established by giving a dual certificate that has the same objective value. Finally, it will be proven that we may perform the greedy assignment in any order.

2 A Nice Point to Start From

The problem is as follows. We are given a period length T in which n arrival events as well as n departure events take place. Denote the points of time of the arrivals by $a_1, \ldots, a_n$, of the departures by $d_1, \ldots, d_n$. Assigning event a_i to event d_j involves a cost of $(d_j - a_i) \bmod T$. We want to find an assignment, or equivalently bipartite perfect matching, of minimal cost. For simplicity, we assume that all events take place at distinct points of time.

We count the arrivals that occur before time $t \in [0, T)$ by the function $a(t) := \#\{i \,|\, a_i < t\}$. For the departures, we analogously define $d(t) := \#\{j \,|\, d_j < t\}$. The function $f(t) := a(t) - d(t)$ counts how much more arrivals than departures at the terminus station did take place before time t.

Lemma 1. Let $t_0 \in \{a_1, \ldots, a_n\}$ minimize the function f, i.e. $f(t_0) = \min\{f(t) \,|\, t \in [0,T)\}$. If we start counting at time t_0, then for any point of time $t \in [t_0, t_0+T)$, we have a non-negative number of vehicles waiting at the terminus station.

Proof. For an easier notation without the modulo operator, we shift all events by an offset $-t_0$, or $T - t_0$ respectively:

$$\tilde{a}_i := \begin{cases} a_i - t_0, & \text{if } a_i \geq t_0 \\ a_i + T - t_0, & \text{if } a_i < t_0 \end{cases} \tag{1}$$

$$\tilde{d}_j := \begin{cases} d_j - t_0, & \text{if } d_j \geq t_0 \\ d_j + T - t_0, & \text{if } d_j < t_0 \end{cases} \tag{2}$$

Count by $\tilde{f}(t)$ the number of vehicles waiting at the terminus station with respect to the events $\tilde{a}_i$ and $\tilde{d}_j$. We show that

$$\begin{array}{ll} \forall\, 0 \leq t < t_0 : & \tilde{f}(t+T-t_0) = f(t) - f(t_0) \geq 0 \quad \text{and} \\ \forall\, t_0 \leq t < T : & \tilde{f}(t-t_0) = f(t) - f(t_0) \geq 0. \end{array}$$

For the case $t_0 \leq t$, we have:

$$\begin{aligned} \tilde{f}(t-t_0) &= \tilde{a}(t-t_0) - \tilde{d}(t-t_0) \\ &= \#\{i \,|\, \tilde{a}_i < t - t_0\} - \#\{j \,|\, \tilde{d}_j < t - t_0\} \\ & (\text{Equation 1}: \ \tilde{a}_i < T - t_0 \Rightarrow \tilde{a}_i = a_i - t_0) \\ &= \#\{i \,|\, t_0 \leq a_i < t\} - \#\{j \,|\, t_0 \leq d_j < t\} \\ &= a(t) - a(t_0) - (d(t) - d(t_0)) \\ &= f(t) - f(t_0). \end{aligned}$$

The case $t < t_0$ is slightly more technical. There, we have:

$$\begin{aligned} \tilde{f}(t+T-t_0) &= \tilde{a}(t+T-t_0) - \tilde{d}(t+T-t_0) \\ &= \#\{i \,|\, \tilde{a}_i < t+T-t_0\} - \#\{j \,|\, \tilde{d}_j < t+T-t_0\} \\ &= \#\{i \,|\, \tilde{a}_i < T-t_0\} + \#\{i \,|\, T-t_0 \leq \tilde{a}_i < t+T-t_0\} - \\ & -\#\{j \,|\, \tilde{d}_j < T-t_0\} - \#\{j \,|\, T-t_0 \leq \tilde{d}_j < t+T-t_0\} \\ & (\text{Equation 1}: \ \tilde{a}_i \geq T-t_0 \Rightarrow \tilde{a}_i = a_i + T - t_0) \\ &= \#\{i \,|\, a_i \geq t_0\} + \#\{i \,|\, a_i < t\} - \\ & -\#\{j \,|\, d_j \geq t_0\} - \#\{j \,|\, d_j < t\} \\ &= (n - \#\{i \,|\, a_i < t_0\}) + a(t) - (n - \#\{j \,|\, d_j < t_0\}) - d(t) \\ &= -a(t_0) + a(t) + d(t_0) - d(t) \\ &= f(t) - f(t_0). \end{aligned}$$

■

By Lemma 1, we know that corresponding to the events $\tilde{a}_i$ and $\tilde{d}_j$, at any point of time, there have been at least as many arrivals as departures. Or, in other words, at any point of time, there can be a non-negative number of vehicles waiting within the terminus station.

3 Proving Optimality

From the preceeding section we may assume w.l.o.g. that for any point of time, we have at most as many departures as arrivals. We refer to an arc (i, j) in the bipartite arrival/departure graph as *forward arc*, if $0 \leq a_i < d_j < T$ and as *backward arc*, if $0 < d_j < a_i < T$. The cost of a forward arc (i, j) is $d_j - a_i$, and $d_j + T - a_i$ for a backward arc.

Hence the FIFO strategy, which leads to the assignment

$$\forall i = 1, \ldots, n : \; a_i \rightarrow d_i,$$

uses only forward arcs.

Theorem 1. The FIFO assignment is an optimal periodic assignment, when starting at time t_0.

Proof. To prove optimality of this assignment, we consider the linear programming relaxation of the assignment problem. This is well known to be totally unimodular, i.e. has the same optimal value as the combinatorial problem. Thus we can prove optimality by providing a dual certificate for the optimality of the LP relaxation. Rockafellar[4] gives a combinatorial motivation for the same dual certificate.

The primal linear program to be studied is

$$\begin{array}{ll} \min & cx \\ \text{s.t.} & B^+ x = \mathbf{1} \\ & B^- x = -\mathbf{1} \\ & x \geq 0, \end{array}$$

where B^+ and B^- are the submatrices of the node-arc-incidence matrix

$$B = \begin{bmatrix} B^+ \\ B^- \end{bmatrix}$$

of the assignment graph, carrying the positive resp. negative entries, or in other words, the departure resp. arrival nodes. The vector $\mathbf{1}$ denotes the all-one vector. The corresponding dual linear program can be formulated as

$$\begin{array}{ll} \max & \sum_{j=1}^{n} u_j - \sum_{i=1}^{n} v_i \\ \text{s.t.} & u_j - v_i \leq w_{ij}, \; \forall a = (i, j) \in A. \end{array} \tag{3}$$

In our special case of periodic assignments, the events' points of time fulfill all the requirements, when starting to count at an appropriate point of time, cf. lemma 1. Since the cost w_{ij} of arc $a = (i, j)$ was defined to be at least $d_j - a_i$, the events' points of time are feasible to the dual program (3). And as the FIFO-assignment makes no use of backward arcs, the primal cost of the assignment is

$$\sum_{i=1}^{n} (d_i - a_i) = \sum_{i=1}^{n} d_i - \sum_{i=1}^{n} a_i,$$

which equals the dual objective of the required dual node values. ■

4 Any Point is a Nice Point

The example in figure 2 shows that different orderings of the arrival events may lead to different greedy FIFO assignments. Assume a period length of $T = 4$,

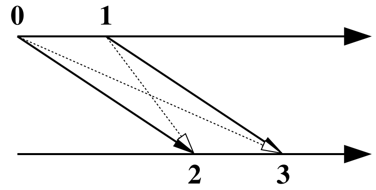

Figure 2: Different orders lead to different periodic assignments

arrivals at times $a_1 = 0$ and $a_2 = 1$, deparuteres at times $d_1 = 2$ and $d_2 = 3$. If we start at time 0, we assign a_1 to d_1. But if we start at time 1, we just get the other possible assignment.

One might ask if the solution quality of the greedy assignment varies under different orderings. But we will show that our shifting has only been necessary for the formal proof. In practice, we may apply the greedy assignment to any order of the arrival events.

Definition. For a given ordering of the arrival events, the *greedy assignment* matches the current arrival event to its nearest unmatched departure event.

Corollar 1. For every ordering, the associated greedy assignment is an optimal periodic assignment.

Proof. From the previous section we know that we can only get worse, if a backward arc has to be involved, because then we get a duality gap of at least T. Assume that an assignment that arises by applying our greedy approach contains a backward arc (a_k, d_l). Let (a_k, d_l) be the first backward arc to be chosen by the algorithm, hence $d_l < a_k$.

Consider the last unmatched departure event d_0 in $[0, T)$. We know $d_0 < a_k$, since otherwise a_k would be matched with d_0 instead of d_l. The claim is that then $f(d_0 + \epsilon) < 0$, which would contradict implications of lemma 1. Say that the number of matched departures after d_0 is $M \geq 0$. Since d_0 has been unmatched, hence available, when any of those were matched, they were matched by M arrivals that reside after d_0. By definition, there is no further departure event after d_0, but at least the event a_k as one additional arrival event. Hence, $a(d_0+\epsilon) \leq n-(M+1)$, but $d(d_0+\epsilon) = n-M$. Thus, $f(d_0+\epsilon) \leq -1 < 0$, which contradicts the definition of our value t_0. ■

Hence, no matter which order we apply, every arrival event may be assigned to the earliest unmatched departure. This can be interpreted as follows: By definition, at time t_0 it is possible that there are no vehicles waiting at the terminus station. Since every backward arc would imply a vehicle to wait at time t_0, we know that in an optimal assignment any vehicle must not wait at the terminus station at time t_0.

5 Further Remarks

For the case where lines with different periods share one terminus station, we could ask if we may avoid expanding the lines to their least common multiple period, and assign only lines with the same period. Unfortunately, in general this does not give an optimal solution.

Consider an example with one line being operated every 10, another one every 15 minutes.

arrival	00	02	10	17	20
departure	01	08	16	18	28

Applying the FIFO rule to the single events occuring within the least common multiple of the two lines' periods, the total vehicle waiting time sums up to $1+6+6+1+8=22$ minutes. If the vehicles were required to continue on lines that are operated in the same interval, a total waiting time of $3\cdot 8+2\cdot 14=52$ minutes would have arosen. Calculating on the greatest common divisor of the periods of the unexpanded arcs' endpoints does not help: The plan

arrival	00	02	10	17	20
departure	05	07	15	22	25

implies a total vehicle waiting time of $5\cdot 5=25$ minutes, whereas it seems to be advantagious when avoiding expansion to the least common multiple by calculating on the greatest common divisor of the four original events.

Notice that any two periodic assignments always differ by an integer multiple of the period. This can easily be seen by investigating the few possible cases when switching two pairs' assignments.

Finally, we do not need to restrict ourselves to assignments within one single terminus station. Even connecting trips of vehicles may be considered. We just have to increase the minimal amount of time that has to elapse from the arrival of the vehicle on one line to the time when it is ready for its next departure on a line starting at a different station. The only thing we have to ensure is that within every group we are dealing with the same number of arrival events as departure events, and that within such a group, any arrival may be matched with any departure.

6 Conclusions

We considered the problem of calculating the amount of rolling stock required to operate a periodic timetable. Even if we allow line changes of vehicles within their terminus stations, we are able to quickly calculate the vehicles' total waiting time during one abstract period, which indeed implies the number of vehicles required to serve the timetable. Although for the general assignment problem the greedy approach has to fail, in our special case of assigning periodically incoming events to periodically leaving events, it leads to an optimal periodic assignment.

This observation has a deep impact for the genetic algorithm proposed by Nachtigall and Voget[3] for solving instances of the Periodic Event Scheduling Problem. Although genetic algorithms require to evaluate a huge number of potentially feasible solutions, we may introduce the very important quantity of the exact amount of vehicles required to perform a periodic timetable, because we know how to compute it sufficiently fast.

References

[1] C. Liebchen and L. Peeters. Some Practical Aspects of Periodic Timetabling. In Proceedings of the Annual Conference of the GOR 2001, to appear.

[2] A. Löbel. Optimal Vehicle Scheduling in Public Transit. Ph.D. Thesis, 1997.

[3] K. Nachtigall and S. Voget. A Genetic Algorithm Approach to Periodic Railway Synchronization. Computers and Operations Research **23**, 453–463, 1996

[4] R.T. Rockafellar. Network Flows and Monotropic Optimization. Wiley, 1984.

Reports from the group

"Combinatorial Optimization and Graph Algorithms"

of the Department of Mathematics, TU Berlin

716/2001 *Christian Liebchen:* The Periodic Assignment Problem (PAP) May Be Solved Greedily

711/2001 *Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, Joseph S. B. Mitchell, and Martin Skutella:* The Freeze-Tag Problem: How to Wake Up a Swarm of Robots

710/2001 *Esther M. Arkin, Sándor P. Fekete, and Joseph S. B. Mitchell:* Algorithms for Manufacturing Paperclips and Sheet Metal Structures

705/2000 *Ekkehard Köhler:* Recognizing Graphs without Asteroidal Triples

704/2000 *Ekkehard Köhler:* AT-free, coAT-free Graphs and AT-free Posets

702/2000 *Frederik Stork:* Branch-and-Bound Algorithms for Stochastic Resource-Constrained Project Scheduling

700/2000 *Rolf H. Möhring:* Scheduling under uncertainty: Bounding the makespan distribution

698/2000 *Sándor P. Fekete, Ekkehard Köhler, and Jürgen Teich:* More-dimensional packing with order constraints

697/2000 *Sándor P. Fekete, Ekkehard Köhler, and Jürgen Teich:* Extending partial suborders and implication classes

696/2000 *Sándor P. Fekete, Ekkehard Köhler, and Jürgen Teich:* Optimal FPGA module placement with temporal precedence constraints

695/2000 *Sándor P. Fekete, Henk Meijer, André Rohe, and Walter Tietze:* Solving a "hard" problem to approximate an "easy" one: heuristics for maximum matchings and maximum Traveling Salesman Problems

694/2000 *Esther M. Arkin, Sándor P. Fekete, Ferran Hurtado, Joseph S. B. Mitchell, Marc Noy, Vera Sacristán and Saurabh Sethia:* On the reflexivity of point sets

693/2000 *Frederik Stork and Marc Uetz:* On the representation of resource constraints in project scheduling

691/2000 *Martin Skutella and Marc Uetz:* Scheduling precedence constrained jobs with stochastic processing times on parallel machines

689/2000 *Rolf H. Möhring, Martin Skutella, and Frederik Stork:* Scheduling with AND/OR precedence constraints

685/2000 *Martin Skutella:* Approximating the single source unsplittable min-cost flow problem

684/2000 *Han Hoogeveen, Martin Skutella, and Gerhard J. Woeginger:* Preemptive scheduling with rejection

683/2000 *Martin Skutella:* Convex quadratic and semidefinite programming relaxations in Scheduling

682/2000 *Rolf H. Möhring and Marc Uetz:* Scheduling scarce resources in chemical engineering

681/2000 *Rolf H. Möhring:* Scheduling under uncertainty: optimizing against a randomizing adversary

680/2000 *Rolf H. Möhring, Andreas S. Schulz, Frederik Stork, and Marc Uetz:* Solving project scheduling problems by minimum cut computations (Journal version for the previous Reports 620 and 661)

674/2000 *Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Sándor P. Fekete, Joseph S. B. Mitchell, and Saurabh Sethia:* Optimal covering tours with turn costs

669/2000 *Michael Naatz:* A note on a question of C. D. Savage

667/2000 *Sándor P. Fekete and Henk Meijer:* On geometric maximum weight cliques

666/2000 *Sándor P. Fekete, Joseph S. B. Mitchell, and Karin Weinbrecht:* On the continuous Weber and k-median problems

664/2000 *Rolf H. Möhring, Andreas S. Schulz, Frederik Stork, and Marc Uetz:* On project scheduling with irregular starting time costs

661/2000 *Frederik Stork and Marc Uetz:* Resource-constrained project scheduling: from a Lagrangian relaxation to competitive solutions

Reports may be requested from:

Hannelore Vogt-Möller
Fachbereich Mathematik, MA 6–1
TU Berlin
Straße des 17. Juni 136
D-10623 Berlin – Germany

e-mail: moeller@math.TU-Berlin.DE

Reports are also available in various formats from

`http://www.math.tu-berlin.de/coga/publications/techreports/`

and via anonymous ftp as

`ftp://ftp.math.tu-berlin.de/pub/Preprints/combi/Report-`*number*`-`*year*`.ps`