VU University Amsterdam

Outpatient Scheduling

Outpatient scheduling for a General Practitioner

PEARL STAM

2011

Outpatient scheduling
Outpatient scheduling for a general practitioner

Pearl Stam

BMI paper
VU University Amsterdam
Faculty of Sciences
Master program Business Mathematics & Informatics
De Boelelaan 1081a
1081 HV Amsterdam

January 2011

Preface

This thesis is part of acquiring the Masters degree in Business Mathematics and Informatics. Business Mathematics and Informatics is a multidisciplinary program, aimed at business processes optimization by applying a combination of methods based upon mathematics, computational intelligence and business management. These three disciplines will also play a central role throughout this thesis.

The subject of this study is outpatient scheduling. The objective of outpatient scheduling is to find an appointment system for which a particular measure of performance is optimized in a clinical environment – it is an application of resource scheduling under uncertainty. First, the particular problem on which we are going to perform the outpatient scheduling is explained. Subsequently, some literature on outpatient scheduling is studied. Eventually, the literature studied is applied to our particular problem, and we shall use a genetic algorithm to try to outperform the scheduling rules from the literature. To measure the performance of the different rules, a simulation tool is created.

I would like to thank Sandjai Bhulai and Dennis Roubos for supporting me during the whole process of writing this paper.

Pearl Stam
February, 2011

Table of Contents

1. Introduction

In a modern society like ours, people are very busy and no one likes to wait. Every minute waiting at a bus stop, a traffic light, or an elevator is considered to be a waste of time.

For a lot of people the biggest burden is to visit the general practitioner, because you know beforehand that you will be sitting in the waiting room for at least a half an hour, if you are lucky, before the doctor will see you.

On the other hand, also the general practitioner is a very busy person, who does not like to wait for his patients to arrive, since he will run out of his schedule if one of them fails to show up in time.

The objective of this paper is to find a reasonable balance between the patient's waiting time, the doctor's idle time and the overtime.

1.1 The problem

The initial problem is to find an optimal appointment system for one general practitioner for one day.

1.1.1 Patients

In our setting, we suppose that there are eight different types of patients. They all have a lognormal service time and their actual arrival time is drawn from a triangular distribution. Below all different types of patients are described.

In the second column the average service time for every patient is shown. The next column shows the standard deviation of the service times. Column four and five give the parameters inserted in the lognormal distribution to achieve the values in the first two columns. Since the lognormal distribution has no upper bound it is truncated to a maximum service time of one hour.

The actual arrival time of patients is simulated by a triangular distribution (of which the parameters are given in the last three columns). Patients usually arrive somewhat early. The average arrival time for a normal patients is three minutes before their scheduled time. The lower limit is eight minutes early and the upper limit is two minutes late. For patients that tend to arrive late the triangle is shifted a few minutes later. The average arrival time is to arrive exactly on time, the lower limit is five minutes early and the upper limit is five minutes late, as can be seen in the table on the following page.

	Avg. service time	St. dev. service time	µ	Σ	Arrival time	Upper limit	Lower limit
Type 1	10	1	2,29761	0,099451	-3	2	-8
Type 2	10	1	2,29761	0,099751	0	5	-5
Type 3	10	2	2,282975	0,198042	-3	2	-8
Type 4	10	2	2,282975	0,198042	0	5	-5
Type 5	20	1	2,994484	0,049969	-3	2	-8
Type 6	20	1	2,994484	0,049969	0	5	-5
Type 7	20	2	2,990757	0,099751	-3	2	-8
Type 8	20	2	2,990757	0,099751	3	8	-2

Table 1.1.1.1 Parameters for the arrival distribution

1.1.2 General Practitioner

The general practitioner always starts working exactly on time. He does not take a break and works until all patients are treated that day.

1.1.3 One day

Every day has eight hours. Four patients of every type are treated per day. Their average service times sum up to 480 minutes which is exactly equal to eight hours. On a day all patients show up, and there are no walk-in patients.

1.1.4 Simulation

To make sure our results are reliable, a simulation program simulates one thousand days per schedule. The performance of each rule is measured by adding up the average waiting time of the patients, the average idle time of the doctor and the average overtime. This is called the fitness of the schedule. In our case the three variables are equally important, but it is also possible to put a weight on each of the variables to indicate the importance.

2. Literature study

The objective of outpatient scheduling is to find an appointment system for which a particular measure of performance is optimized in a clinical environment – it is an application of resource scheduling under uncertainty. The underlying problem applies to a wide variety of environments, such as general practice patient scheduling, scheduling patients for an MRI device, surgical scheduling, etcetera.

2.1 Simulation studies

Outpatient scheduling in health care has been researched extensively over the last 50 years. To model outpatient queuing systems a considerable number of those studies uses simulation techniques. One of the advantages of simulation modeling over analytical approaches is the ability to model complex outpatient queuing systems and to represent environmental variables. Simulation experiments are conducted to evaluate the performance of the system and to understand the relationship between various performance measures and various environmental factors.

The most primitive form of outpatient scheduling is single block scheduling. The single block rule assigns all patients to arrive at the same time. The patients are served on a first come first serve basis. Another, nowadays more common, form of appointment scheduling is the individual block rule. Patients are assigned unique appointment times that are spaced throughout the clinical session.

Bailey (1952) was one of the first to analyze an individual block system. At that time in most clinics it was common practice to assign all patients to arrive at the same time.
Bailey combines single block and individual block scheduling. A number of patients is assigned the same arrival time at the beginning of the clinical session. The idea behind this is to keep an inventory of patients so that the doctor's risk of becoming idle is minimized if the first patient arrives late or fails to show up. All other patients are assigned unique appointment times spread throughout the clinical session.
Bailey used a Monte-Carlo simulation technique to find the number of patients to assign an appointment at the beginning of the session and the length of the intervals between the remaining appointment times.
From this he concluded that he should schedule two patients at the beginning of the session. The remaining patients are scheduled at intervals equal to the mean consultation time. This leads to a reasonable balance between the patient's waiting time and the doctor's idle time.
Bailey also found that shorter mean consultation times result in lower patient waiting times. Furthermore, he found that high variability of service times deteriorates both the patients' waiting times and the doctor's idle time.

Assigning time blocks to surgeons on a first come first served basis to find a balance between the surgeon's waiting cost, the idle cost of the facilities and operation room personnel is studied by Charnetski (1984) using simulation. The heuristic found distinguishes different types of procedures having different service time distributions and bases procedure times scheduled for a patient on a function of the mean and the standard deviation of the individual service times.

Ho and Lau (1992, 1999) and Ho, Lau and Li (1995) introduce a number of variable-interval rules and test their performance against traditional ones using simulation. Their best performing variable-interval rule increases appointment intervals toward the latter part of the session. They conclude that there is not one rule that performs well under all circumstances and provide a simple heuristic to assist in selecting an appointment rule for a clinic.
Their assessment of three environmental factors (no-show probability, variability of service times, and number of patients per clinical session) reveals that, among the three, the no-show probability is the major one that affects the performance and the choice of an appointment schedule.

Klassen and Rohleder (1996) classify patients based on their expected service time variability. They use simulation techniques to compare various ways of scheduling patients having a relatively high and relatively low service time variability, when appointment intervals are kept standing. They developed a rule that puts patients with lower service time variability before patients with higher service time variability, which performs better than Ho and Lau's best performing rules.
Later on, they consider the possibility that the scheduler can make errors and that not all patients accept every slot they are assigned to. However, they conclude that their rule mentioned above still performs well under these more realistic assumptions.

An appointment system for a multi-server queuing system, where doctors may arrive late, with constant intervals between two successive appointment times and multiple variable blocks is studied by Liu and Liu (1998). They try to minimize the total cost of the patient's flow-time and the doctors' idle time by developing a simulation search procedure to appoint the number of patients to assign to each block. They suggest a simple procedure to find an appointment rule for a given environment using the properties of the best rules, derived after simulating several environmental factors.

Swisher et al. (2001) developed a discrete-event simulation model to be applied for decision making in outpatient planning. By utilizing this model to a family practice clinic they observe that the results are quite sensitive to changes in the patient mix, patient scheduling, and staffing levels. The effect of patient scheduling is only studied in changing the instant of the appointment, rather than examining several appointment rules.

The majority of the studies mentioned above assume patients are homogeneous for scheduling purposes, and use independently and identically distributed service times for all patients. Furthermore, those studies do not take into account structural latecomers. Another disadvantage is that most studies focus on a particular problem so the solution is not suitable for any problem.

2.2 My approach

From previous studies it is clear that there is not one particular solution to the outpatient scheduling problem.

The objective of outpatient scheduling is to find an appointment system for which a particular measure of performance is optimized in a clinical environment. In most studies this performance measure is a trade-off in the interest of physicians and patients. Patients prefer to have short waiting times, while physicians like to have as little idle time and overtime as possible.

Outpatient clinics can be regarded as queuing systems. A unique set of conditions is considered when designing an appointment system. The simplest case is when all patients arrive on time and one single doctor serves them with stochastic service times. Unfortunately in most clinical sessions this never occurs.
The presence of non-punctual patients, no-shows, walk-ins, and emergencies may intervene to upset the schedule. Furthermore, doctors can arrive late or may be interrupted during the clinical session which messes up the system as well.

The objective of this paper is to find an optimal appointment schedule for a General Practice using genetic algorithms.

The performance of formal studies will be measured using simulation techniques. Subsequently, we shall try to find a schedule that performs better than the ones found from previous studies.

In searching for an optimal appointment system some assumptions have to be made.
We shall try to find an optimal appointment system for one general practitioner. This general practitioner is very punctual. He always begins exactly on time and he never gets interrupted during the clinical session. Patients, however, are not that punctual. Some patients have a reputation of being non-punctual, which is taken into account.
There are several types of patients. Every type has a known service time distribution with known parameters. Presence of walk-ins, no-shows and emergency patients will be neglected.

On the basis of previous assumptions we shall try to find an appointment system that balances the patient’s waiting times, the doctor’s overtime and the doctor’s idle time.

3. Approach

3.1 Methods

Initially we shall look into several appointment scheduling rules mentioned in the previous chapter, and measure the performance on our specific problem. Subsequently we will try to find an appointment system that performs better than those existing rules using genetic algorithms.

3.1.1 Individual block

The first appointment scheduling rule that will be used is the individual block rule. Every patient is assigned a unique appointment time spread out over the clinical session.

3.1.2 Bailey & Welch

The second rule to be studied is the Bailey Welch rule. This rule is almost similar to the individual block rule except now two patients are scheduled at the beginning of the clinical session.
The Bailey Welch rule assumes patients have independent identically distributed Gamma service times. Furthermore, it is assumed that all patients arrive exactly on time.

3.1.3 Charnetski

The final rule to be measured is Charnetski's rule. He derived a heuristic that distinguishes different types of procedures having different service time distributions and bases the amount of time scheduled for a patient on the mean and the standard deviation of the individual service times.
Charnetski assumes that all patients arrive exactly on time, and does not take any overtime into account. The service times are normally distributed and truncated from below at 0. The primary focus of the study is to determine a relationship between the scheduled time for a procedure and the average idle time of the specialist and the waiting time of the patients. A standardized prediction heuristic was used to schedule procedure times, given by $d_i(h) = \mu_i + h\sigma_i$ where h is a scalar value held constant across procedures, μ_i is the empirically determined mean for procedure i and σ_i is the procedure's standard deviation. The value $d_i(h)$ represents the amount of time scheduled for procedure i as a function of the factor h.

3.2 Genetic algorithm

We shall approach this problem with a genetic algorithm. The genetic algorithm finds its roots in artificial intelligence. This heuristic is used to generate useful solutions to optimization and search problems. Genetic algorithms belong to the larger class of evolutionary algorithms, which generate solutions to optimization problems using techniques inspired by natural evolution.

3.2.1 Initialization

From an initial population a large number of individual solutions is randomly generated. Occasionally, the solutions may be directed to areas where optimal solutions are likely to be found.

3.2.2 Selection

The performance of every solution is determined by a function called the fitness function. This fitness function indicates the performance of the specific solution compared to the other solutions. A lower score on the fitness function indicates a higher performance. Based on these fitness values the population is diminished, where solutions that are less fit are more likely to be selected for deletion.

3.2.3 Reproduction

The next step is to generate a second generation population of solutions from those selected. This can be realized through genetic operators like recombination and mutation. In recombination two existing solutions are combined to find two new solutions. Mutation slightly changes several randomly selected existing solutions.
Occasionally, the solutions with the highest fitness are excluded from the evaluation process mentioned above to prevent the algorithm to get stuck at a local optimum. The process is repeated until the population has reached the same size as the previous generation. These processes ultimately result in the next generation population that is different from the initial generation. Generally, the average fitness will have increased by this procedure for the population, since only the best solutions from the first generation are selected to find new solutions.

3.2.4 Termination

The selection and reproduction steps are repeated until a termination condition has been reached.
Common terminating conditions are:

- There might be a minimum criterion to the problem. The process can be terminated if the minimum criterion is satisfied or a fixed number of generations is reached;
- The highest ranking solution's fitness is reaching or has reached a plateau such that successive iterations no longer produce better results.

3.3 Simulation

To model the outpatient queuing system we shall use a simulation tool. As mentioned before, one of the advantages of simulation modeling over analytical approaches is the ability to model complex outpatient queuing systems and represent environmental variables. With this simulation tool we will be able to measure the performance of our different scheduling systems.

3.3.1 Assumptions

For our simulation tool some assumptions have to be made in order to determine the performance of our scheduling systems.

Both the arrival process and the service time of a patient are modeled as a probability process. There are four types of treatments, with short and long service times, and large and small standard deviations. Some patients tend to arrive late.

Patients are served according to the following rules.

- When there are no patients in the waiting room a newly arriving patient is served immediately;
- If there is one patient in the waiting room this patient is served;
- If there is more than one patient in the waiting room, the patient with the earliest scheduled time is served;
- All patients are served, regardless of the scheduled finishing time of the doctor.

3.3.2 Parameters

Patients almost never arrive exactly on time for their appointment. Although most patients arrive somewhat early, several patients have a tendency to arrive late. Both early and late patients are modeled with a triangular arrival distribution, with a different mean.

There are four types of treatments with either a relatively long or a short service time and a larger or a smaller standard deviation. Service times are drawn from a lognormal distribution.

3.3.3 Heuristic

The diagram below shows the steps of the simulation tool while simulating one day. This will be repeated for a large number of days.

Preliminary, the following steps have to be taken.

1) The schedule to be measured has to be entered in the simulation tool;
2) For every patient in the schedule the actual arrival time has to be determined;
3) All patients must be sorted in order of their actual arrival time.

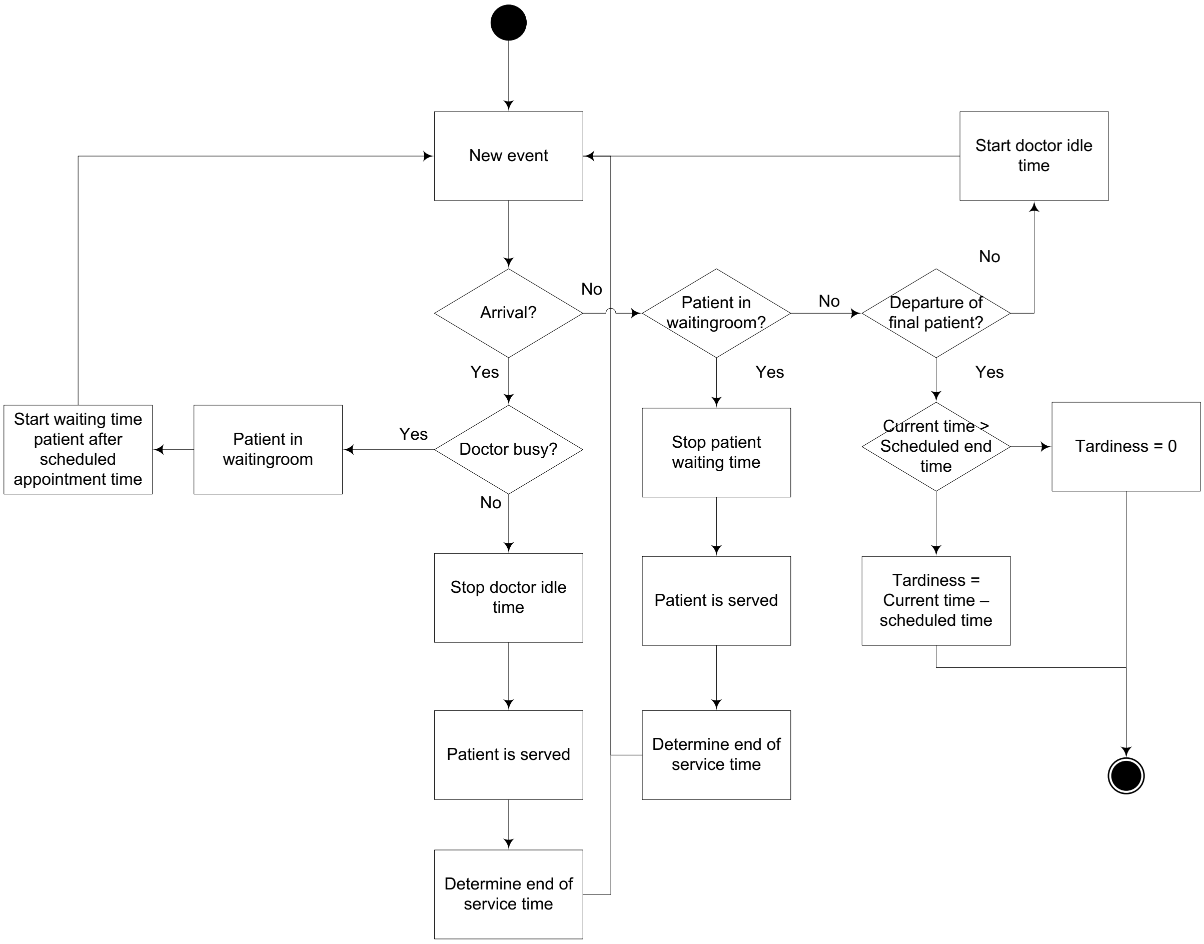

Image 3.3.3.1 Steps of the simulation tool while simulating one day

The doctor idle time is retained, as is the waiting time for every patient. At the end of the day, the tardiness is determined.
The doctor idle time is the time in between two services. The waiting time of the patient is the difference between their appointment time and the time the patient is actually served. This excludes any waiting prior to appointment time, because additional waiting due to early arrival is voluntary and is not a consequence of the appointment schedule. If the patient is served before the appointment time the waiting time is zero. Late patients may

consider some additional waiting as normal, being partly their own fault. Their waiting time is set to zero as well.

3.3.4 Output

The following output is generated from the simulation tool

- The mean and standard deviation of the waiting time of the patient;
- The mean and standard deviation of the idle time of the doctor;
- The mean and standard deviation of the tardiness.

4. Results

The different algorithms we picked to compare to the genetic algorithm do not say that much about the order we have to schedule the patients in, by means of average service times, standard deviations, or arrival times. We have tried several scenarios that seem to be obvious using some knowledge from other literature.

According to Bailey, shorter mean consultation times result in lower patient waiting times. Klassen and Rohleder concluded that patients should be scheduled in order of increasing standard deviation. Furthermore, it seems logical to schedule latecomers after punctual patients, because you do not want your system to get messed up already at the beginning of the day.

Now the only thing we have to investigate is the order of importance of the rules mentioned above.

4.1 Individual Block

The time reserved for every patient is the average service time for that patient. For the individual block system we found the following results:

The patients are scheduled in order of:
1. Increasing service time;
2. Increasing standard deviation;
3. Increasing possibility of arriving late.

Mean overtime	6.6037	Variance overtime	30.9676
Mean waiting time	3.1305	Variance waiting time	15.9294
Mean idle time	6.4110	Variance idle time	18.3622
Fitness	16.1476		

Table 4.1.1 Results for the individual block system with patients scheduled in order of increasing service time, increasing standard deviation and increasing possibility of arriving late.

Subsequently, patients are scheduled in order of:
1. Increasing standard deviation;
2. Increasing service time;
3. Increasing possibility of arriving late.

Mean overtime	6.9242	Variance overtime	35.0269
Mean waiting time	2.9906	Variance waiting time	15.0648
Mean idle time	6.4982	Variance idle time	18.1916
Fitness	16.4131		

Table 4.1.2 Results for the individual block system with patients scheduled in order of increasing standard deviation, increasing service time and increasing possibility of arriving late.

Finally, patients are scheduled in order of:

1. Increasing possibility of arriving late;
2. Increasing service time;
3. Increasing standard deviation.

Mean overtime	6.8204	Variance overtime	38.3545
Mean waiting time	3.5351	Variance waiting time	20.3208
Mean idle time	6.5380	Variance idle time	15.7736
Fitness	16.8935		

Table 4.1.3 Results for the individual block system with patients scheduled in order of increasing possibility of arriving late, increasing service time and increasing standard deviation.

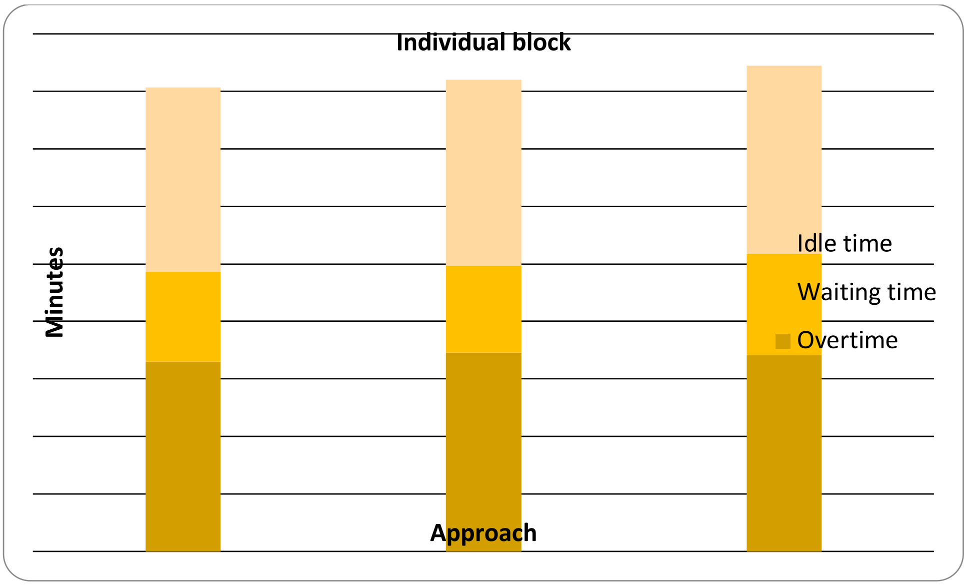

Image 4.1.1 Results for the individual block system.

4.2 Bailey & Welch

The Bailey & Welch rule is very similar to the individual block system. The only difference now is that we schedule the final patient at the same time as the first patient. From the above results we find that there is not that much difference in the performance of the different schedules, but the first scheme works slightly better than the following two, so we will use this scheme to apply Bailey's rule on.

First, we give the final patient from the schedule above an appointment at time zero.

Mean overtime	3.3455	Variance overtime	24.7531
Mean waiting time	13.3588	Variance waiting time	106.8742
Mean idle time	0.0537	Variance idle time	0.1463
Fitness	16.7580		

Table 4.2.1 Results for the Bailey & Welch system, final patient at time zero.

Now we put the last patient with a short service time at time zero and shift all patients with a long service time ten minutes back.

Mean overtime	3.4227	Variance overtime	25.0010
Mean waiting time	6.9368	Variance waiting time	40.5728
Mean idle time	0.8513	Variance idle time	5.3766
Fitness	11.2109		

Table 4.2.2 Results for the Bailey & Welch system, final patient with a short service time at time zero.

Finally, we chop the day in half and pretend the morning is reserved for patients with a short service time, and patients with a long service time are served in the afternoon. Both in the morning and the afternoon, the final patient is put at time zero of the part of day.

Mean overtime	3.8972	Variance overtime	28.7913
Mean waiting time	10.4235	Variance waiting time	92.7438
Mean idle time	1.0056	Variance idle time	4.8501
Fitness	15.3262		

Table 4.2.3 Results for the Bailey & Welch system, final patient with a short service time at time zero and final and first patient with a long service time at the same time.

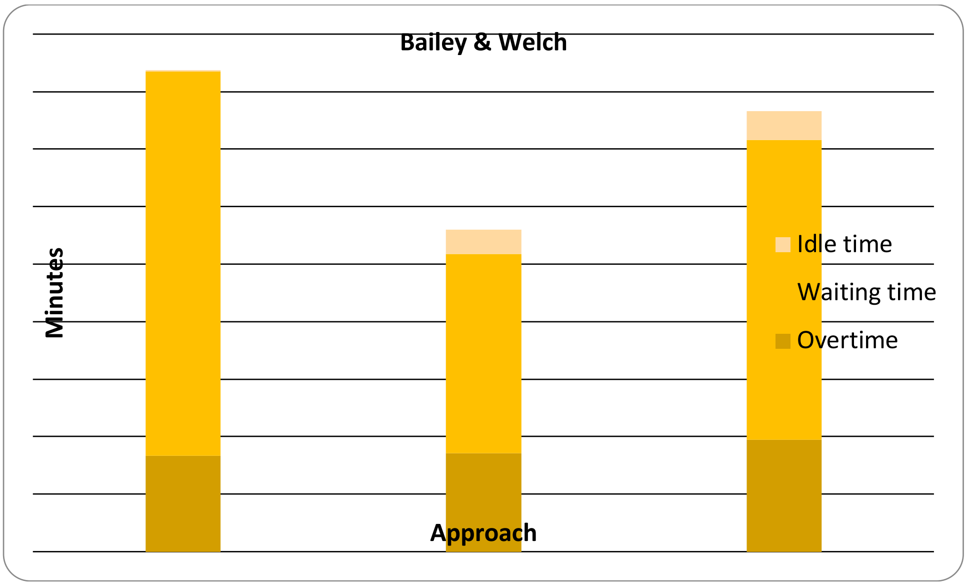

Image 4.2.1 Results for the Bailey & Welch system.

4.3 Charnetski

As mentioned before, Charnetski determines the procedure time for a patient with the following formula $d_i(h) = \mu_i + h\sigma_i$. In this formula, h is a scalar value held constant across procedures, μ_i is the empirically determined mean for procedure i and σ_i is the procedure's standard deviation. The value $d_i(h)$ represents the amount of time scheduled for procedure i as a function of the factor h.

Again, we use the first schedule found for the individual block system, but now the procedure times are calculated using the formula for $d_i(h)$. Since it is a bit silly to schedule a patient at 14 minutes and 23 seconds past nine, we shall round the times in the schedule we find to the closest whole minute.

For h = 0 we find the same results as for the individual block system.

For h = 0.1 we find:

Mean overtime	8.8859	Variance overtime	27.3369
Mean waiting time	2.3253	Variance waiting time	10.6687
Mean idle time	9.0343	Variance idle time	29.4476
Fitness	20.2455		

Table 4.3.1 Results for Charnetski's system, h = 0.1.

Apparently, increasing h results in a higher fitness. We shall try some negative values for h.

For h = - 0.1 we find:

Mean overtime	5.6628	Variance overtime	36.0550
Mean waiting time	4.0786	Variance waiting time	24.1332
Mean idle time	4.8147	Variance idle time	12.2727
Fitness	14.5561		

Table 4.3.2 Results for Charnetski's system, h = -0.1.

For h = - 0.2:

Mean overtime	5.2898	Variance overtime	36.3351
Mean waiting time	5.2630	Variance waiting time	37.4420
Mean idle time	3.6482	Variance idle time	8.0002
Fitness	14.2009		

Table 4.3.3 Results for Charnetski's system, h = -0.2.

For h = - 0.3

Mean overtime	4.6722	Variance overtime	33.3282
Mean waiting time	6.3862	Variance waiting time	49.8470
Mean idle time	2.8881	Variance idle time	4.4373
Fitness	13.9465		

Table 4.3.4 Results for Charnetski's system, h = -0.3.

For h = - 0.4

Mean overtime	4.6302	Variance overtime	34.1016
Mean waiting time	7.8117	Variance waiting time	69.0342
Mean idle time	2.4140	Variance idle time	2.7721
Fitness	14.8559		

Table 4.3.5 Results for Charnetski's system, h = -0.4.

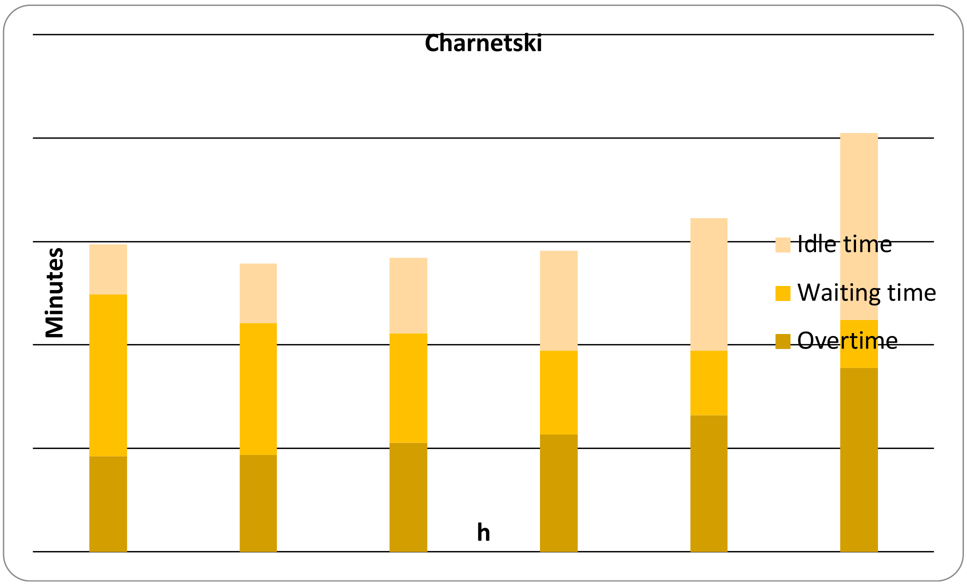

Image 4.3.1 Results for Charnetski's system.

4.4 Combination

Finally, we combine the best results from the Bailey & Welch rule with the results of Charnetski's rule.

The last patient with a short service time is put at time zero, and the patients with long service time are shifted back the amount of time that is scheduled for this patient. Procedure times are calculated for h equal to - 0.1.

Mean overtime	3.4473	Variance overtime	25.5487
Mean waiting time	8.3176	Variance waiting time	55.0513
Mean idle time	0.3045	Variance idle time	1.2819
Fitness	12.0694		

Table 4.4.1 Results for Bailey & Welch and Charnetski's system, h = -0.1.

The last patient with a short service time is put at time zero, and the patients with long service time are shifted back the amount of time that is scheduled for this patient. Procedure times are calculated for h equal to - 0.2.

Mean overtime	3.3609	Variance overtime	26.0050
Mean waiting time	9.6214	Variance waiting time	69.9474
Mean idle time	0.1364	Variance idle time	0.4456
Fitness	13.1187		

Table 4.4.2 Results for Bailey & Welch and Charnetski's system, h = -0.2.

The last patient with a short service time is put at time zero, and the patients with long service time are shifted back the amount of time that is scheduled for this patient. Procedure times are calculated for h equal to - 0.3.

Mean overtime	3.8402	Variance overtime	28.4364
Mean waiting time	11.2007	Variance waiting time	93.9860
Mean idle time	0.0334	Variance idle time	0.0560
Fitness	15.0743		

Table 4.4.3 Results for Bailey & Welch and Charnetski's system, h = -0.3.

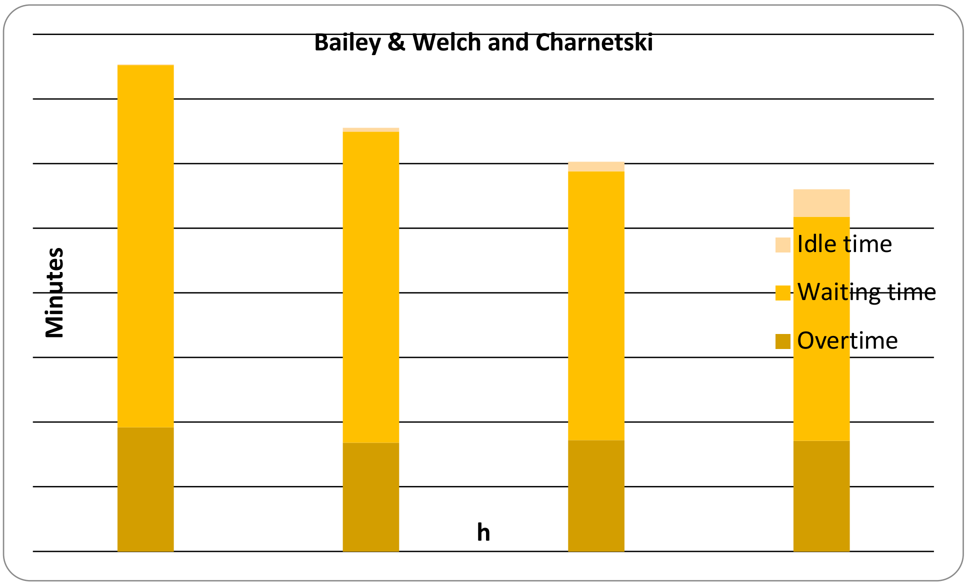

Image 4.4.1 Results for Bailey & Welch combined with Charnetski's system.

4.5 Genetic algorithm

To find a better solution than the previous ones, we tried to optimize the fitness of the schedule using the genetic algorithm solver within MatLab. The variables are the arrival times of the patients, which can vary between 0 and 480.

We started with a population of 20 initial solutions and the number of generations was equal to 100. The initial population contained 20 individuals with 32 random integers between 0 and 480. Unfortunately, this resulted in something worse than everything we found up until now.

Subsequently, we started to increase the population size which improved our results, but we were still not able to find a fitness that exceeded the previous results.

Our next approach was to search in the neighborhood of the best results we found up until now. We created an initial population where every variable from every individual was in a range of 5 minutes of the optimal value up until now. The following results were found applying the genetic algorithm with a population size of 250 and the number of generations equal to 200.

Mean overtime	3.2777	Variance overtime	24.6602
Mean waiting time	5.1208	Variance waiting time	31.3429
Mean idle time	1.2365	Variance idle time	6.5035
Fitness	9.6350		

Table 4.5.1 Results for the Genetic Algorithm near Bailey & Welch.

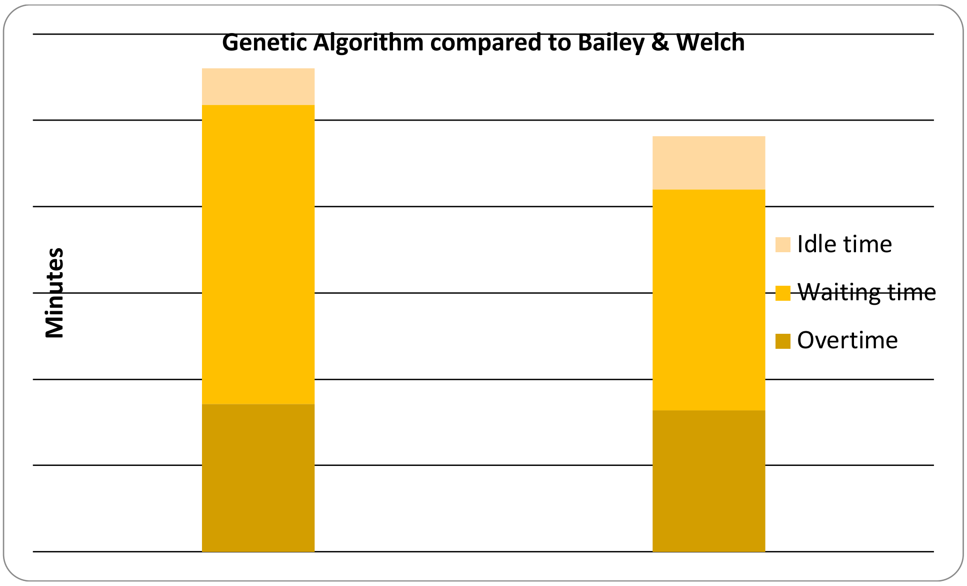

Image 4.5.1 Results for the Genetic Algorithm compared to Bailey & Welch.

In the graph above it can be seen that with the genetic algorithm we are able to improve the Bailey Welch rule. Although the idle time has increased a little bit, we have managed to decrease the overall fitness almost two minutes by taking a lot of time of the average waiting time for the patients.

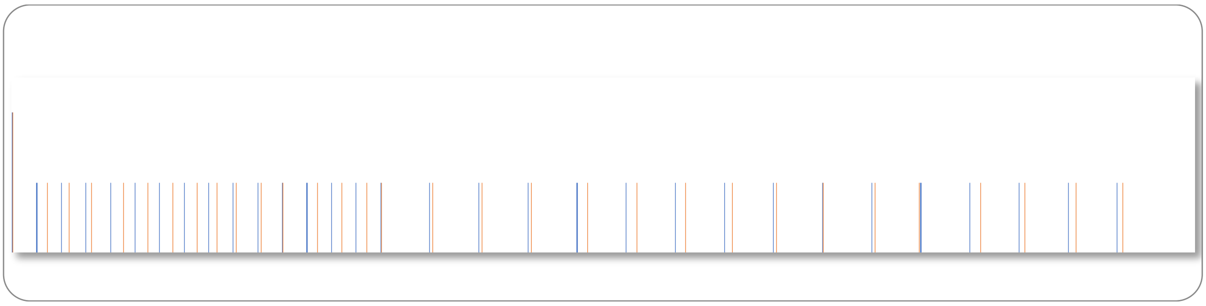

Image 4.5.2 The optimal schedule found with the Genetic Algorithm next to the Bailey & Welch schedule.

The image above shows the optimal solution found with the Bailey & Welch approach, in blue, and the optimal solution found with the genetic algorithm in red. Every bar represents an appointment. The two higher bars at time zero represent two arrivals at the same time.

It is remarkable that the genetic algorithm schedules patients that arrive in time very close to the Bailey & Welch schedule, while patients that have a tendency to arrive late are scheduled several minutes later than Bailey & Welch.
Furthermore, we see that the third patient to arrive is shifted four minutes forward compared to the Bailey & Welch schedule. This makes sense, since this patient has to wait for two patients to be served in front of him, instead of one.

5 Conclusion and discussion

5.1 Conclusion

From the results we can conclude that it is indeed possible to find a schedule with the genetic algorithm that outperforms the other algorithms mentioned.
However it was not that easy as we thought beforehand. We got stuck in a local optimum several times, so we really had to push the algorithm in the right direction to get some satisfying results.

We have seen that, although the Bailey & Welch rule is almost sixty years old, it is still a very good performing rule that could not be matched by Charnetski in this setting. Also we saw that it is wise to consider some scheduling rules instead of simply implementing the individual block system, since every rule we applied outperformed the individual block rule.

Unfortunately, we noticed that the variations of the waiting time, idle time and overtime were pretty high, which means that the results we found may count for the long run, but we will probably not find the same results if we observe only a few days.

5.2 Discussion

For this paper some assumptions had to be made that might not be completely credible in reality. For instance, we assumed there were no no-shows, while in reality it is always possible that patients fail to show up. And it is not very likely that the doctor will always start exactly on time.

Also, not every patient is willing to accept any time of day he or she is assigned to. For instance, people that have a job, often want an appointment at the beginning or the end of the day and will not settle for an appointment at noon.

Furthermore, the general assumption of independence between the arrival and the service patterns may be questionable. In practice, doctors may increase their service rate during peak hours knowing that there are many patients waiting.

6 References

N. Bailey (1952). A study of queues and appointment systems in hospital outpatient departments, with special reference to waiting-times.

J.R. Charnetski (1984). Scheduling operating room surgical procedures with early and late completion penalty costs.

C.-J. Ho and H.-S. Lau (1992). Minimizing total cost in scheduling outpatient appointments.

Tugba Cayirli and Emre Viral (2003). Outpatient scheduling in healthcare: A review of literature.