Modelling and evaluating Obstructive Sleep Apnea management protocols with UPPAAL

Adrian-Catalin Rus
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
a.c.rus@student.utwente.nl

ABSTRACT

Obstructive Sleep Apnea is a sleep disorder which can increase all-cause mortality by threefold in severe cases and its prevalence is estimated to 19%. Patients seeking treatment for this disorder might undergo different procedures before receiving a long-term treatment, but the task of ordering these procedures in an efficient manner is non-trivial. In this research, the modelling tool UPPAAL was used to construct and simulate models of different management protocols for Obstructive Sleep Apnea. We compared and analyzed 14 different protocols by key performance indicators such as time, costs and efficacy in providing a treatment. We concluded that features of a good protocol are: starting with a treatment and including a diagnostic test that can be double-checked with a sleep study and with a successful treatment.

Keywords

obstructive sleep apnea, medical protocols, formal models, simulations

1. INTRODUCTION

Obstructive Sleep Apnea (*OSA* for short) is a condition that causes partial or total blockage of the respiratory tract during sleep. Patients suffering from it have a higher risk of hypertension, insulin sensitivity, "neural injury" and other cardiovascular diseases[2]. More precisely, it has been found that it increases the chances of ischemic stroke by threefold and for atrial fibrillation by fourfold[9]. This is problematic because ischemic stroke and atrial fibrillation are a leading cause of death[11][8]. A review of the epidemiology of *OSA* points to an higher estimated prevalence in men than women (22% respectively 17%) and an up to threefold increase in mortality[5]. Moreover, it has been found that patients suffering from *OSA* have an up to fifteen-fold increase of motor vehicle accidents[6].

Managing a patient suffering from *OSA* is not yet well established because there is a multitude of procedures that work in different manners and address different concerns (e.g. treatments, sleep studies, pre-diagnostic tests). Therefore compiling such procedures into protocols results in a much more varied spectrum of management protocols, which hinders the health care system's duty of choosing the best one. This difficulty mostly emerges from the need to have a large number of patients to test the best candidate protocols, but also from the need of considering varied key performance indicators such as costs, time spent and efficacy in offering a long-term treatment.

We aimed at leveraging the formal and statistical capabilities of UPPAAL[1] to compare multiple protocols without the aforementioned difficulties.

28th Twente Student Conference on IT Febr. 2nd, 2018, Enschede, The Netherlands.

2. APPROACH

UPPAAL allows for hybrid complex systems to be created out of individual models which can be linked together. These systems can then be simulated and specific information can be extracted from the results of the simulations in order to draw conclusions.

2.1 First approach

It is important to note that our initial progress had to be discarded in the later weeks of this research because of its limited scope, high complexity and missing relevant data. Our approach then was to model the patient and the possible treatments for *OSA* in detail and try to answer which ordering of the treatments is the most beneficial to the patient and the health care system[10]. This was problematic because we were trying to model the prevalence of *OSA* severity and complaints (i.e. symptoms and other incommodities caused by *OSA*), but also their evolution in time and how specific treatments can periodically (i.e. yearly) affect these. Moreover, we tried to model the patient's attitude towards a treatment, so he could be able to quit a treatment without reason. All of these aims proved to be unfeasible mainly because of the missing or inconsistent scientific data, which led to many assumptions that had to be made based on our subjectivity.

2.2 Current approach

A core assumption of this approach is that our modelled patients have been referred to the clinic or sleep centre. From there, they will follow our pre-established protocols. We also assume that once a protocol reaches a conclusion about what ought to be done about the patient, the UPPAAL simulation can stop.

The "building blocks" of our current models are:

- *Treatment*: given to a patient, with the possibility of long-term adherence.
- *Sleep study*: used as a "perfect" diagnostic procedure, with no false positives or false negatives.

- *Diagnostic test*: used as a cheaper, but less accurate diagnostic procedure.

We used certain permutations of these blocks to build the protocols that we want to examine:

- *P1*: Sleep study → Treatment
- *P2*: Treatment → Sleep study
- *P3a*: Diagnostic test → Treatment or Sleep study, if the test concludes positive for *OSA*
- *P3b*: Diagnostic test → Treatment if the test concludes positive or Sleep study otherwise, with the posibility to take the Treatment afterwards
- *P4*: Treatment → Diagnostic test, if the patient is not compatible with the Treatment → Sleep study, if the test concludes positive for *OSA*

Each protocol reaches a final state that reflects the decision of the health care system, or a failure of the protocol, where applicable. We see the failures as being hidden from the health care system unless they are discovered because of a sleep study. The set of possible final states is the following:

- *Continue treatment*: The Treatment works and the patient will continue to take it in long-term.
- *Change the treatment*: The Treatment does not work, but we know that the patient has *OSA*.
- *No OSA*: We find that the patient does not have *OSA*.
- *False negative*: The Diagnostic test concludes negative, but the patient has *OSA*. This is a failure that can only occur in *P3a* and *P4*.
- *False positive*: The Diagnostic test concludes positive, but the patient does not have *OSA*. This is a failure that can only occur in *P3b*.

Finally, in order to draw appropriate conclusions, appropriate questions must be asked. We focused on asking the following, for each protocol accordingly, based on the final states and variables of the models:

1. **How many weeks does a patient spend using the protocol?**
2. **What are the costs of using the protocol?**
3. **What is the probability that the patient continues using the treatment?**
4. **What is the probability that a patient is refused treatment, although he has *OSA*?** (false negative)
5. **What is the probability that a patient is given treatment, although he does not have *OSA*?** (false positive)

2.3 Implementation

In our models, the aforementioned *Treatment* has been taken to be Continuous Positive Airway Pressure (CPAP for short) because it is considered the most effective *OSA* treatment[12]. We assume that CPAP, if compatible with the patient, treats *OSA* after a certain period. So after trying CPAP, if the patient's symptoms do not disappear, it is because the patient does not suffer from *OSA*, but another sleep disorder. For the *Sleep study* we chose polysomnography (PSG for short) because it is the most accurate sleep study method currently known[7]. Our *Diagnostic test* was chosen to be four tests instead of just one, because there are significant differences in terms of their purpose and performance[3][4].

Implementation decisions were made with the consultation of *OSA* experts (see Section 6), who agreed with or shifted our initial estimates of the values to be used in our models. The implementation characteristics of our building blocks were finally chosen to be:

- CPAP
  - time spent: 13 weeks if the patient is incompatible or 3 weeks otherwise.
  - money spent: 345 euros
- PSG
  - time spent: 24 weeks
  - money spent: 1500 euros
- Diagnostic tests
  1. BerlinQ questionnaire
     - sensitivity: 46%
     - specificity: 70%
     - time spent: 1 week
     - money spent: 50 euros
  2. StopBang questionnaire
     - sensitivity: 80%
     - specificity: 55%
     - time spent: 1 week
     - money spent: 50 euros
  3. Philips questionnaire + ODI (Oxygen desaturation index) measurement
     - sensitivity: 100%
     - specificity: 35%
     - time spent: 1 week
     - money spent: 120 euros
  4. Philips questionnaire + Nasal flow measurement
     - sensitivity: 63%
     - specificity: 90%
     - time spent: 1 week
     - money spent: 120 euros

Above, by *sensitivity* we refer to the confidence of a test to conclude in a true positive, while by *specificity* we refer to the confidence of a test to conclude in a true negative. This means that we can easily model the probability of a test to conclude in a false positive and false negative by computing $1-specificity$ and $1-sensitivity$, respectively. Using these definitions, our assumption of PSG being the

perfect diagnostic procedure translates into having a *sensitivity* and *specificity* of 100%.

We assume that CPAP always costs money to try, but if the patient is compatible, he will always spend 13 weeks trying it, no matter if he suffers from *OSA* or not. However, the total time spent will only count as 3 weeks if the patient has *OSA*; the logic being that after 3 weeks it will be clear to the protocol that CPAP must be taken in long-term. For the patient, this assumption translates into 13 weeks spent, but not for the protocol, which counts only the 3 weeks relevant to it. It also serves as a logic base for further improvements, which was an important aspect we kept in mind throughout the research.

2.3.1 Global variables

UPPAAL uses global variables in the same fashion as any commonly-used programming language (e.g. *C++*). We defined ours in the following way:

- *time* as an integer, used for tracking and incrementing the time spent using a protocol.
- *costs* as an integer, used for tracking and incrementing the money spent using a protocol.
- *OSA* as a boolean, used as a perfect indicator of whether the patient has the condition or not.
- *cpapable* as a boolean, used as an indicator of whether the patient is compatible with CPAP or not.
- *positive* as a boolean, used as an indicator of whether the *Diagnostic test* concludes positive for *OSA* or not, no matter if the test is mistaken or not.

2.3.2 Patient model

We use the patient model (see Figure 1) mostly for initialisation purposes, specifically for determining if he has *OSA* and if he is compatible with CPAP. These insights are only available to the model for analysis purposes, as in practice it is unrealistic to determine if the patient has the condition without doing a diagnostic.

The process of initialising the patient begins in the left-most state "At_clinic", from where a probabilistic choice is made to set the global variable *OSA* to *true* with a probability of 70% or to *false*, with a probability of 30%. The same type of choice is again made to determine the value of *cpapable*. After this, the patient model is locked in either "CPAPable" or "not_CPAPable" end state.

It is important to note that the states marked with a "C" inside them mean that time is not allowed to pass while in those states and that their transitions have priority over the others. By employing this type of states, we make sure that as soon as the simulation starts, our patients are instantly initialised, before the other models can start.

2.3.3 P1 model

In *P1* (see Figure 3) we begin by doing a PSG, which takes 24 weeks and 1500 euros to complete. We use a local clock variable w to measure the time spent in the "do_PSG" state by imposing the invariant $w <= 24$, which allows us to remain within that state for 24 time units. The following transition is only allowed to be taken if $w >= 24$. When it happens, we increment *costs* with the appropriate costs and we increment *time* by using the method "count" (see Figure 2). We then reach a decision node, which splits the

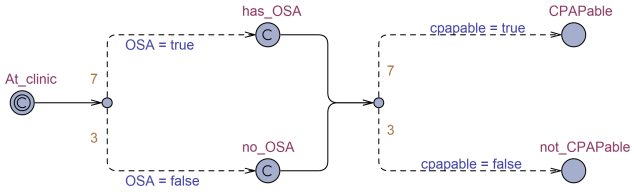

Figure 1. The patient model

trajectory based on the value of *OSA*. If the patient has the condition, he will try CPAP for the cost of 345 euros, then based on whether he is *cpapable* or not, he spends 13 weeks to then adhere to a long-term CPAP or he spends 3 weeks to conclude that he needs some other treatment. Recall that the protocol only increments *time* by 3 weeks for the former option, as discussed in Section 2.3.

Similar to the states in the patient model noted with a "C", the states noted with an "U" mean only that time is not allowed to pass in those states, so they are perceived as being instant. This differentiation is useful especially for decision states, where we do not want to spend time in order to branch.

```
void count(int t){
time = time + t;
}
```

Figure 2. The count method declaration

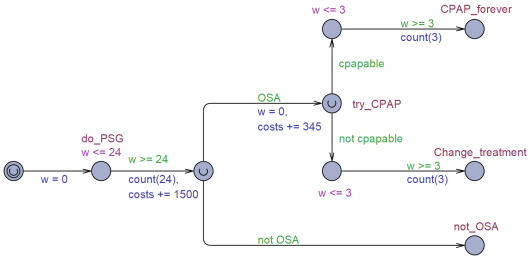

Figure 3. The *P1* protocol model

2.3.4 P2 model

In *P2* (see Figure 4) we begin by directly trying CPAP and paying 345 euros. If the patient is compatible, he will spend 13 weeks and he might adhere to a long-term treatment, but he might not have *OSA*, in which case he will not continue the treatment. If there is no compatibility, as before, 3 weeks will be spent, after which a PSG will be taken to precisely determine if the patient needs another treatment or not.

2.3.5 Diagnostic test model

This model (see Figure 5) is split into 3 parts: one for picking which test to be done, one for performing the test and one for determining the conclusion of the test. The choice for the test is made based on a synchronization signal received from another model. This is indicated in the model by the cyan-colored text followed by a question mark (e.g. "test_bq?" for choosing BerlinQ). It waits for

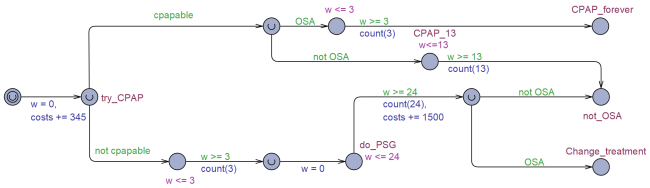

Figure 4. The *P2* protocol model

either of these synchronisation signals to be sent and sets values for the variables *sensitivity* and *specificity* accordingly. It also sets a variable t to the appropriate number of weeks needed to complete the test and increments the *costs* with the right amount. The patient then spends t weeks doing the test and this time is added to the total amount of time spent.

Determining the result of a diagnostic test begins by branching based on the value of *OSA*: if it is *true*, we will use the *sensitivity* of the test for the conclusion and if it is *false*, we will use its *specificity*. Recall that the former is the confidence of a test to conclude in a positive result, while the later is the confidence of a test to conclude in a negative result. So, if the patient has *OSA*, with a probability of *sens*%, he is correctly diagnosed positive, while the probability that he is falsely diagnosed as negative is $(100 - sens)\%$. If the patient does not have the condition, with a probability of *spec*% he is correctly diagnosed as negative, but he is falsely diagnosed as positive $(100 - spec)\%$ of the times. This decision is memorised by using the global variable *positive*.

The final step in this model is to send a synchronization signal to the other models in order to communicate that the *Diagnostic test* has reached a conclusion. This transition creates a loop, arriving at the "Start" state, where this process can be repeated if needed. We currently do not make use of this capability, but it can be used in further improvements of this work.

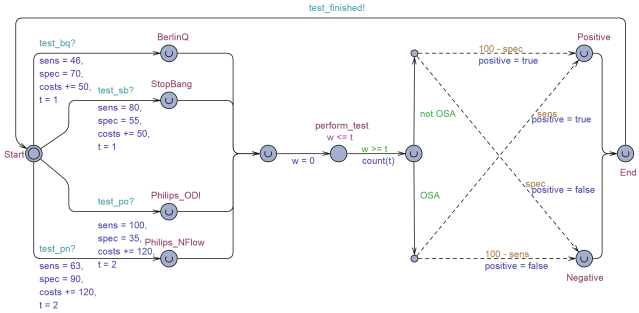

Figure 5. The *Diagnostic test* model

2.3.6 *P3a model*

In *P3a* (see Figure 6) we begin by doing a test of our choice: we send a synchronization signal and wait in the "do_test" state until we receive the synchronization signal "test_finished". We then make a decision based on the result of the *Diagnostic test*. If the result is negative and the patient does not have *OSA*, we stop in the "not_OSA" state, but if the patient has the condition, then we can say that the test has given a false negative result and we stop the protocol. If the result is *positive*, we follow the same steps as in the *P2* protocol.

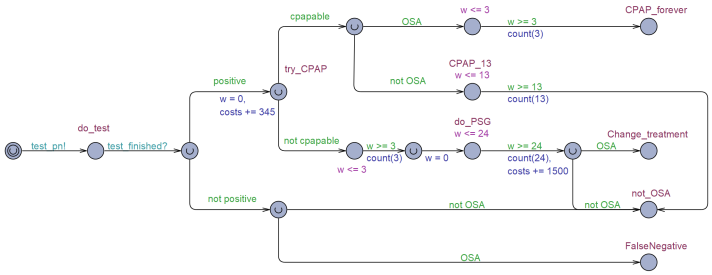

Figure 6. The *P3a* protocol model

2.3.7 *P3b model*

This protocol (see Figure 7) mainly differs from *P3a* by doing a PSG in the case of a negative result of the *Diagnostic test*. This allows the protocol to "save" the patient in case of a false negative of the test, by redirecting him to "try_CPAP". From there, he will arrive either in "CPAP_forever" or in "Change_treatment", depending on the compatibility with CPAP. Another important difference is the possibility to end in the state "FalsePositive", which can happen if the patient is not compatible with CPAP.

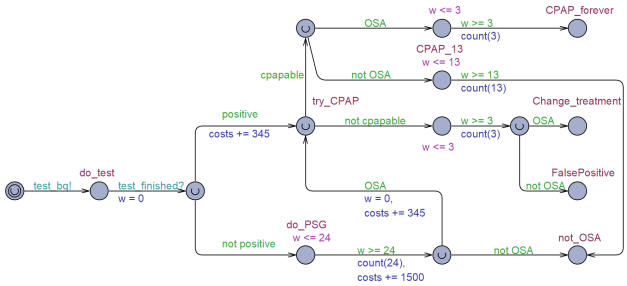

Figure 7. The *P3b* protocol model

2.3.8 *P4 model*

This protocol (see Figure 8) is very similar to *P2* because both start by offering a CPAP, but in *P4* if the patient is not *cpapable*, the patient does a *Diagnostic test*. Here if the result is *positive*, the patient undergoes a PSG to determine whether he should change the treatment or if it is a "false alarm".

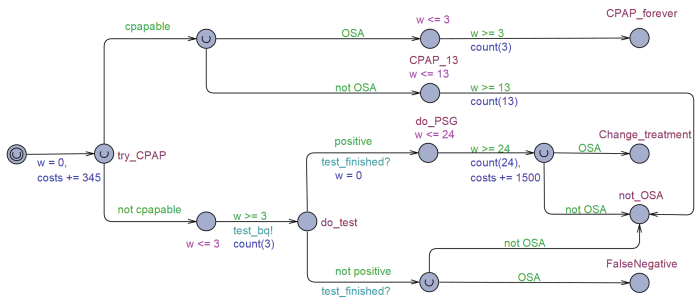

Figure 8. The *P4* protocol model

2.3.9 *Retrieving the results*

Using UPPAAL's Statistical Model Checker tool, we can query the simulations for results. The queries used directly correlate to the five questions mentioned before this subsection:

1. **How many weeks does a patient spend using the protocol?**
   - E[<=100; 1000000](max: time)
   - Reads: "What is the expectancy of *time* if we run one million simulations that stop after 100 time units?"
2. **What are the costs of using the protocol?**
   - E[<=100; 1000000](max: costs)
   - Reads: "What is the expectancy of *costs* if we run one million simulations that stop after 100 time units?"
3. **What is the probability that the patient continues using the treatment?**
   - Pr[<=100](<>P1.CPAP_forever)
   - Pr[<=100](<>P2.CPAP_forever)
   - Pr[<=100](<>P3a.CPAP_forever)
   - Pr[<=100](<>P3b.CPAP_forever)
   - Pr[<=100](<>P4.CPAP_forever)
   - Reads: "What is the probability that within 100 time units the given model arrives at the state "CPAP_forever"?"
4. **What is the probability that a patient is refused treatment, although he has *OSA*?**
   - Pr[<=100](<>P3a.FalseNegative)
   - Pr[<=100](<>P4.FalseNegative)
5. **What is the probability that a patient is given treatment, although he does not have *OSA*?**
   - Pr[<=100](<>P3b.FalsePositive)

In order to ensure accurate results, we changed the default settings of the UPPAAL SMC tool. More precisely, we changed the α, β and ϵ to 0.001 (see Figure 9). This allows, for example, our probability queries to be answered with a confidence interval of 99.9%.

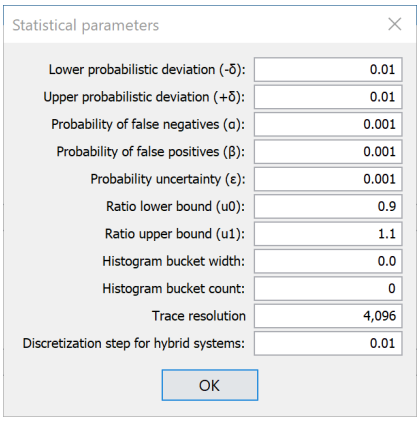

Figure 9. The parameter values used for our simulations

3. RESULTS

Our queries' results were compiled into a spreadsheet (see Figure 10) and color-coded to make them easier to understand. We chose four levels of performance (i.e. best, well, bad and worst performing) and we categorised the results based on comparisons with the other results. For the false negative and false positive chances, we categorised them subjectively, arguing that a probability lower than 5% is the best performing. This subjective categorization is relevant here, especially for the false negatives because of their severe impact on the patient's health span, while the false positives negatively affect the health care system's resource efficiency.

P1				compared amongst each other
Average time spent (weeks)	26			Best performing
Average money spent (euros)	1741			Well perfomring
CPAP_forever chance	0.489			Bad performing
				Worst performing

P2	
Average time spent (weeks)	12.2
Average money spent (euros)	795
CPAP_forever chance	0.49

P3a

Test	BerlinQ	StopBang	Philips + ODI	Philips + NFlow	Averages
Average time spent (weeks)	5.8	9	12.5	7	8.575
Average money spent (euros)	378	602	830	494	576
FalseNegative chance	0.377	0.139	0	0.258	0.1935
CPAP_forever chance	0.225	0.392	0.489	0.308	0.3535

P3b

Test	BerlinQ	StopBang	Philips + ODI	Philips + NFlow	Averages
Average time spent (weeks)	18	11.75	8.5	17	13.8125
Average money spent (euros)	1204	795	586	1166	937.75
FalsePositive chance	0.027	0.04	0.058	0.009	0.0335
CPAP_forever chance	0.489	0.489	0.489	0.49	0.48925

P4

Test	BerlinQ	StopBang	Philips + ODI	Philips + NFlow	Averages
Average time spent (weeks)	8.3	10.4	12.1	9.1	9.975
Average money spent (euros)	544	673	783	593	648.25
FalseNegative chance	0.113	0.042	0	0.077	0.058
CPAP_forever chance	0.489	0.489	0.489	0.489	0.489

Figure 10. The results of our queries

By reading the color-coding scheme, we can quickly indicate that P4 is the best overall performer because of its low number of bad-performing results. The *StopBang*, *Philips+NasalFlow* and *Philips + ODI* variants of this protocol are the overall best choices, although *P2*, *P3b: StopBang*, *P3b: Philips+ODI* and *P3a': Philips + ODI* are close to their performance.

We can also identify that *P1* is overall the worse protocol and that *P3a* (except for the *Philips + ODI* variant) should be avoided because of its problematic high probability of false negatives.

4. RESOURCE EFFICIENCY ANALYSIS

If we want to investigate the resource efficiency of a protocol when arriving at a conclusion, we can do so by backtracking from an end state to the beginning of the protocol and observe the actions and conditions that allow the protocol to reach that state. By resource efficiency we refer to spending at least time and money to reach a given end state. This is important for both the patient and the health care system, precisely: time is more important to the patient and money is more important to the health care system.

We compiled another spreadsheet which we color-coded as

Protocol / End state	CPAP_forever	Change_treatment	not_OSA	FalsePositive	FalseNegative
P1	27 weeks + 1845 euros	27 weeks + 1500 euros	24 weeks + 1500 euros		
P2	3 weeks + 345 euros	27 weeks + 1845 euros	if *cpapable* : 13 weeks + 345 euros if not *cpapable* 27 weeks + 1845 euros		
P3a	3 weeks + 345 euros + test	27 weeks + 1845 euros + test	test + if *positive: 345 euros* + (13 weeks if *cpapable* OR 27 weeks + 1500 euros if not *cpapable*)		test
P3b	3 weeks + 345 euros + test (+ 24 weeks + 1500 euros if not *positive*)	3 weeks + 345 euros + test (+ 24 weeks + 1500 euros if not *positive*)	test + (13 weeks + 345 euros if *cpapable* and if *positive* OR 24 weeks + 1500 euros if not *positive*)	3 weeks + 345 euros + test	
P4	3 weeks + 345 euros	27 weeks + 1854 euros + test	if *cpapable* : 13 weeks + 345 euros if not *cpapable* : 3 weeks + 345 euros + test + (24 weeks + 1500 euros if *positive)*		3 weeks + 345 euros + test

Figure 11. Resources spent for each conclusion, for each protocol

before. In Figure 11 we can observe that *P4* has some undesirable performances, although it is the overall best performer. This is however only problematic when the state "Change_treatment" is reached, which can be avoided by employing *Diagnostic tests* with a high *specificity*. Such a change will also be preferable for minimising the costly impact of "not_OSA" because it would reduce the chances of false positives, so all those costs would not be fruitless.

We see that *P3b* is not the best performer for any state, but that is because of the dependence on the result of the test. We can also identify that when the test concludes negative, the resource efficiency is always worse so it would be desirable for this protocol to employ *Diagnostic tests* with a high *sensitivity*. High *specificity* would also help minimise the costs of arriving in "not_OSA" and "FalsePositive". Another interesting remark about this protocol is that it is the only one that has the possibility to reach "Change_treatment" without needing a PSG, which would be desirable as PSG is the most expensive "building block" in our models.

P3a is overall rather efficient, especially in combination with a test with high *specificity* because it increases the chances of concluding in "not_OSA" and "FalseNegative", but also minimises the costs of concluding in "not_OSA" if the test is *positive* by reducing the possibility of false positives.

Also rather efficient is *P2* which comes close to the efficiency of *P4* because of their very similar structure. Overall, the most inefficient is *P1* because it directly begins with the most costly procedure, PSG.

5. CONCLUSIONS

As we observed in the later two sections, the overall best management protocol is *P4*, followed by *P2* and *P3b*. Improvements to this protocol can be made by coupling it with a higher *sensitivity* (i.e. $> 70\%$) test than our chosen *Diagnostic tests*. Such an improvement will lower the probability of the protocol accepting a false negative to under 5%.

By investigating the performance and efficiency of *P4* compared to the others, we have learnt about two effective factors that improve a management protocol:

1. **Trying a treatment early in the protocol.** This is because it directly provides the patient with the opportunity of a long-term treatment, without having to undergo other costly procedures.
2. **Including diagnostic tests into the protocol.** This helps to reduce the time and money spent by providing an alternative for the costly sleep study. However, it introduces the possibility of false positives and negatives and can lower the probability of offering the patient a treatment. These problems can be properly addressed with adequate diagnostic tests.

We also learnt important lessons from the poor performances of protocols *P1* and *P3a* respectively:

1. **Doing a sleep study first in the protocol greatly increases the time and money spent.**
2. **Blindly trusting a diagnostic test greatly reduces the accuracy of the protocol.** This is especially important if the diagnostic test does not have a high *sensitivity* (i.e. $> 80\%$), but can be remedied by following with a sleep study (as seen in *P3b*), although this would increase the spent resources.

We can conclude that approaches like ours can offer unexpected and important insight into the workings of proposed medical protocols. Considering the advantages of such approaches over conventional testing methods, we believe that the future will bring more importance for mathematical modelling in the context of decision making in the health care systems.

5.1 Future work

We have built our models with future work in mind, allowing further detail to be added on top of the existing solution. However, our models arguably lack some crucial details, which could be addressed by future iterations. Therefore we suggest the following additions and improvements:

- Define *OSA* and include it within the models.
- Include the prevalence of *OSA* severity and complaints. This will require relevant and consistent scientific data about these aspects.
- Model what can happen before the patient arrives at the clinic. This could include, but should not be limited to visits to the general practitioner, referral procedures and the evolution of *OSA* severity and complaints over time.
- Do a sensitivity analysis on the values of different parameters in order to appreciate how certain aspects of the models influence the behavior and outcome of the protocols.
- Use a better, more objective and data-driven coloring code for comparing the results.

6. ACKNOWLEDGEMENTS

I would like to thank the *OSA* experts drs. Miranda Wetselaar-Glas (Researcher at LUMC, Leiden) and drs. Piet-Heijn van Mechelen (Chairman of Stichting ApneuResearch) for their vital feedback, information, estimates and faith in our approach. I also want to address special thanks to dr. ir. Rom Langerak for being a well-rounded supervisor.

7. REFERENCES

[1] UPPAAL. https://www.it.uu.se/research/group/darts/uppaal/. Accessed: 21-06-2020.

[2] J. A. Dempsey, S. C. Veasey, B. J. Morgan, and C. P. O'Donnell. Pathophysiology of sleep apnea. *Physiological Reviews*, 90(1):47–112, 2010. PMID: 20086074.

[3] M. M. Eijsvogel, S. Wiegersma, W. Randerath, J. Verbraecken, E. Wegter-Hilbers, and J. V. D. Palen. Obstructive sleep apnea syndrome in company workers: Development of a two-step screening strategy with a new questionnaire. *Journal of Clinical Sleep Medicine*, 12(04):555–564, 2016.

[4] T. M. Fabius, J. R. Benistant, R. G. Pleijhuis, J. V. D. Palen, and M. M. M. Eijsvogel. The use of oximetry and a questionnaire in primary care enables exclusion of a subsequent obstructive sleep apnea diagnosis. *Sleep and Breathing*, 24(1):151–158, 2019.

[5] K. A. Franklin and E. Lindberg. Obstructive sleep apnea is a common disorder in the population—a review on the epidemiology of sleep apnea. *Journal of Thoracic Disease*, 7(8), 2015.

[6] S. Horstmann, C. Hess, C. Bassetti, M. Gugger, and J. Mathis. Sleepiness-related accidents in sleep apnea. *Sleep*, 23:383–9, 06 2000.

[7] V. Ibáñez, J. Silva, and O. Cauli. A survey on sleep assessment methods. *PeerJ*, 6, 2018.

[8] A. Odutayo, C. X. Wong, A. J. Hsiao, S. Hopewell, D. G. Altman, and C. A. Emdin. Atrial fibrillation and risks of cardiovascular disease, renal disease, and death: systematic review and meta-analysis. *BMJ*, 354, 2016.

[9] S. Redline, G. Yenokyan, D. J. Gottlieb, E. Shahar, G. T. O'Connor, H. E. Resnick, M. Diener-West, M. H. Sanders, P. A. Wolf, E. M. Geraghty, T. Ali, M. Lebowitz, and N. M. Punjabi. Obstructive sleep apnea–hypopnea and incident stroke. *American Journal of Respiratory and Critical Care Medicine*, 182(2):269–277, 2010. PMID: 20339144.

[10] C. Rus. Uppaal models (unfinished work), Jun 2020.

[11] A. E. Vidale S. Organizing healthcare for optimal acute ischemic stroke treatment. *J Clin Neurol.*, Apr(2):183–190, 2020.

[12] I.-T. Vlachantoni, E. Dikaiakou, C. N. Antonopoulos, C. Stefanadis, S. S. Daskalopoulou, and E. T. Petridou. Effects of continuous positive airway pressure (cpap) treatment for obstructive sleep apnea in arterial stiffness: A meta-analysis. *Sleep Medicine Reviews*, 17(1):19–28, 2013.