Management Summary

By: Audrey Tandiono
Supervisors: Dr Michael O'Sullivan
Dr Cameron Walker

22 December 2011

Problem Summary

This project extends the simulation of the patient transit process at Auckland Hospital in Auckland, New Zealand with a model consisting of two buildings with a total of nine floors of the hospital where most of the transits take place. The simulation was developed by Ivo Matthijssen and further extended by Jordi Timmermans.

Figure 1 Map of the Patient Transit Simulation Model

Orderlies transport patients from their wards to treatment rooms and back to their wards once the treatment has been completed. Some transit patients in certain conditions may also require a transit nurse to assist with the transit. Most transits take place in building 32 and 01 and the simulation has been developed to model the nine floors where a major proportion of the transits occur. The orderly and nurse base is located at level 5 of building 32 as shown in Figure 1 where the dispatcher is also located.

The dispatcher uses the Transit Booking System (INFRA) which shows all the booked requests. The task of the dispatcher is to schedule all the orderlies and transit nurses in such a way that the patient arrives at most 15 minutes before their appointment and at least at their appointment time. However, this is not always possible, so we are also interested in ensuring a patient arrives between 15 minutes before and 15 minutes after their appointment time.

The initial goal of the project is to use historical data obtained from the Transit Booking System to calibrate the existing simulation model to represent the Auckland Hospital's transit system as accurately as possible.

Solution Approach

The simulation was developed in Java using Eclipse as the integrated development environment (IDE) with SSJ library to define the simulation model. Matlab was also used to optimise the difference between the historical data and the simulated patient drop off times. Jordi Timmermans used Matlab Builder JA and the Matlab Optimization Toolbox in his initial calibration of the model, but we do not have a license for Matlab Builder JA (Jordi used his TU Eindhoven license).

Instead we developed "wrapper" code to link Java and Matlab directly (i.e., not through Matlab Builder JA). Using our wrapper we have experimented with `fmincon` (from the Optimization Toolbox), `fminsearchbnd` (a bounded version of `fminsearch` which was available from Mathworks' Matlab Central File Exchange) and `annealing` (also available from Mathworks' Matlab Central File Exchange). We have utilised these functions to calibrate the speed of each nurse, orderly and elevator, as well as the delay times at each ward and elevator. The objective function value that we are trying to optimise is a weighted sum of the mean and standard deviation of the differences between the historical data and the simulated patient drop off times. The active set algorithm in the `fmincon` function finds the minimum of a constrained nonlinear multivariable function using Sequential Quadratic Programming (SQP) and an active set method to solve each QP. This is essentially a descent method (follows the derivatives "downhill"). The `fminsearchbnd` function finds the minimum of unconstrained multivariable function while allowing for bounds using the Nelder-Mead Simplex method, a derivative free method. The `annealing` function uses Simulated Annealing which minimises the function by searching through neighbouring solutions and may initially accept worse solutions to obtain a global minimum (not get trapped at a local minimum).

After obtaining results from the optimisation we verify and validate the solution of the calibration by using different starting points for the optimisation to see if our current solutions are biased due to the optimisation process's starting parameters. By starting with realistic and uniform delays and speeds we are able to check the movement tendencies of the parameters.

We are currently resolving the discrepancy in the way we process the data as in some cases the orderly in our data is actually a nurse and vice versa or there are only one or two nurses catering to a request (i.e. no orderlies). Adjustments in the simulation are in progress to accommodate the adjustments in the input data.

Results and Analysis

Jordi has previously used the `fmincon` function in Matlab to optimise the difference between the historical data and the simulated patient drop off time for the time period 1 October 2011 to 5 October 2011, as shown in Table 1, to calibrate:

- a uniform speed for the orderlies with patients
- a uniform speed for the orderlies without patients
- a uniform speed for the nurses without patients
- delay times for the pickup and drop off at Radiology
- delay times for the pickup and drop off at building 01 level 13 and 14
- delay times for the pickup and drop off at other locations

Orderlies Speed No Patient (m/s)	0.78
Orderlies Speed With Patient (m/s)	7.81
Nurses Speed No Patient (m/s)	2.58
Delay lift building 32 (min)	0.12
Delay lift building 01 (min)	0.11
Delay pickup (min)	0.89
Delay drop off (min)	0.54
Delay pickup b01 level 13+14 (min)	1.38
Delay drop off b01 level 13+14 (min)	0.97
Delay pickup radiology (min)	0.05
Delay drop off radiology (min)	1.32
Elevator speed (m/s)	0.70

Table 1 Jordi's optimisation result 04 Nov 2011

Jordi's optimisation for difference between the historical data and the simulation for each patient drop off time resulted in a mean difference of 0.02 minutes and a standard deviation of 12.13 minutes with an objective function value of 182.02 seconds.

We have extended this optimisation process by allowing for different speeds for each orderly, nurse and elevator. Different delays at each pick-up and drop-off location have also been allowed for. The results from Jordi's optimisation has been used as starting values of the optimisation process to achieve a lower objective function to ensure an improvement in the result while allowing for more realistic bounds.

The result of the optimisation using the active set algorithm in the `fmincon` function is shown on Table 2. Using the parameters below, the mean difference has increased slightly to 0.05 minutes but the standard deviation has reduced to 10.42 minutes with an overall objective function of 158.60 seconds.

Mean Orderlies Speed No Patient (m/s)	1.815
Mean Orderlies Speed With Patient (m/s)	1.778
Mean Nurses Speed No Patient (m/s)	2.567
Delay lift building 32 (min)	1.072
Delay lift building 01 (min)	0.103
Mean Delay pickup (min)	1.111
Mean Delay drop off (min)	0.588
Mean Delay pickup b01 level 13+14 (min)	2.000
Mean Delay drop off b01 level 13+14 (min)	1.250
Mean Delay pickup RAD2 + RAD5 (min)	0.050
Mean Delay drop off RAD2 + RAD5 (min)	4.000
Elevator speed building 32 (m/s)	0.995
Elevator speed building 01 (m/s)	0.995

Table 2 Optimisation result using fmincon function

The result of the optimisation using Simulated Annealing is shown on Table 3. Using the parameters below, the mean difference has decreased to 0 minutes and the standard deviation has reduced to 10.40 minutes with an overall objective function of 155.97 seconds.

Mean Orderlies Speed No Patient (m/s)	1.819
Mean Orderlies Speed With Patient (m/s)	1.783
Mean Nurses Speed No Patient (m/s)	2.568
Delay lift building 32 (min)	1.082
Delay lift building 01 (min)	0.099
Mean Delay pickup (min)	1.105
Mean Delay drop off (min)	0.576
Mean Delay pickup b01 level 13+14 (min)	1.994
Mean Delay drop off b01 level 13+14 (min)	1.250
Mean Delay pickup RAD2 + RAD5 (min)	0.051
Mean Delay drop off RAD2 + RAD5 (min)	3.982
Elevator speed building 32 (m/s)	0.999
Elevator speed building 01 (m/s)	0.998

Table 3 Optimisation result using Simulated Annealing method

The result of the optimisation using the `fminsearchbnd` function is shown on Table 4. Using the parameters below, the mean difference has been reduced to 0 minutes and a standard deviation of 11.14 minutes with an objective function of 167.14 seconds.

Mean Orderlies Speed No Patient (m/s)	1.80
Mean Orderlies Speed With Patient (m/s)	1.77
Mean Nurses Speed No Patient (m/s)	2.57
Delay lift building 32 (min)	0.071
Delay lift building 01 (min)	0.022
Mean Delay pickup (min)	1.11
Mean Delay drop off (min)	0.59
Mean Delay pickup b01 level 13+14 (min)	2.00
Mean Delay drop off b01 level 13+14 (min)	1.25
Mean Delay pickup RAD2 + RAD5 (min)	0.05
Mean Delay drop off RAD2 + RAD5 (min)	4.00
Elevator speed building 32 (m/s)	0.30
Elevator speed building 01 (m/s)	0.30

Table 4 Optimisation Result `fminsearchbnd` function

The average adult walking speed range from 1.27 m/s to 2.53 m/s. Jordi's optimisation resulted in walking speeds of 0.78 – 7.81 m/s - note that 7.81 m/s would give a 100m time of just over 12s which is not realistic for a busy hospital. Our optimisation resulted in walking speeds of 1.77 – 2.57 m/s which are more realistic. The speed of elevators from our optimisations in Table 2 and 3 above are closer to the average speed of a hospital elevator of 1 m/s compared to Jordi's result (1 m/s compared to 0.7 m/s).

Our solutions show that speed of nurses is significantly faster than orderlies speed, pick up delays are longer than drop off delays (except for radiology which has a significantly longer drop off delay than pick up delay) and lift delay for building 32 is longer than building 01 (lift delay of building 01 is expected to be longer than building 32). To make sure that our solution is not biased due to the starting point of the optimisation process, we have experimented using different starting values. Starting values and results of the optimisation processes using the simulated annealing method is shown in Table 5 and 6 below.

	Before	After
Mean Orderlies Speed No Patient (m/s)	2	2.034
Mean Orderlies Speed With Patient (m/s)	1.5	1.532
Mean Nurses Speed No Patient (m/s)	2	2.014
Delay lift building 32 (min)	1	1.009
Delay lift building 01 (min)	1.5	1.446
Mean Delay pickup (min)	1.5	1.457
Mean Delay drop off (min)	1	0.980
Mean Delay pickup b01 level 13+14 (min)	1.5	1.471
Mean Delay drop off b01 level 13+14 (min)	1	0.962
Mean Delay pickup RAD2 + RAD5 (min)	1.5	1.479
Mean Delay drop off RAD2 + RAD5 (min)	1	0.936
Elevator speed building 32 (m/s)	1	1.017
Elevator speed building 01 (m/s)	1	1.057
Mean difference (min)		0.302
Standard deviation (min)		12.246
Objective function (sec)		197.286

Table 5 Experimentation starting point and result using uniform starting speeds and delays to start optimisation

	Before	After
Mean Orderlies Speed No Patient (m/s)	1.7	1.697
Mean Orderlies Speed With Patient (m/s)	1.5	1.494
Mean Nurses Speed No Patient (m/s)	1.7	1.704
Delay lift building 32 (min)	0.5	0.500
Delay lift building 01 (min)	1	0.995
Mean Delay pickup (min)	1	1.012
Mean Delay drop off (min)	1	1.012
Mean Delay pickup b01 level 13+14 (min)	1	1.015
Mean Delay drop off b01 level 13+14 (min)	1	0.990
Mean Delay pickup RAD2 + RAD5 (min)	1	1.005
Mean Delay drop off RAD2 + RAD5 (min)	1	0.992
Elevator speed building 32 (m/s)	1	1.030
Elevator speed building 01 (m/s)	1	0.974
Mean difference (min)		0.000
Standard deviation (min)		12.586
Objective function (sec)		188.797

Table 6 Experimentation starting point and result using uniform starting speeds and delays to start optimisation

With a starting speed with no patient of 2 m/s, the optimisation resulted in speed of orderlies to be faster than speed of nurses, while starting with a speed of 1.7 m/s resulted in speed of orderlies to be slower than speed of nurses. This inconsistency does not confirm that the speed of orderlies is slower than the speed of nurses as suggested by our optimisation results.

By starting with pick up delays of 1.5 min and drop off delays of 1 min, the optimisation did not significantly change the initial 0.5 min difference between pick-up and drop off delays as both delays decreased by approximately the same amount. Using a starting pick up and drop off delay of 1 min resulted in longer pick up delays than drop off delays which verified our optimisation result that showed longer pick up delays (except for radiology where our result showed a significantly longer drop off delay).

Starting with a longer lift delay for building 01 resulted in decrease in the initial 0.5 min difference between the delays as seen in both Table 5 and 6 which may suggest that the lift delays in both buildings may not be different.

Due to the numerous adjustable parameters (a total of 292 parameters) in the optimisation process, results from experimentations above are not conclusive as there are many ways to compensate for the increase or decrease of a parameter. For example, an increase in speed may be compensated by an increase in delays. Further investigation is required to verify and validate our results.

A summary of the obtained optimisation results is shown on Table 7. Currently, the mean difference between simulated drop off times and historical data drop off times is very close to 0, but the variation still requires improvement as suggested by the standard deviation of more than 10 minutes.

	Jordi	fmincon	annealing	fminsearchbnd
Mean (min)	0.02	0.05	0.00	0.00
Standard deviation (min)	12.13	10.42	10.40	11.15
Objective function (sec)	182.02	158.60	155.97	167.19

Table 7 Summary of optimisation results

Future Work

- Complete adjustments to resolve the orderly-nurse discrepancy in the data.
- Use distributions instead of averages for speeds and delays to allow for variations
- Create metrics for comparisons between simulated results and historical data
- Implement orderlies and nurses schedules (data required)
- Experiment with different dispatching policies
- Validate solutions by physically measuring speeds and delays in the hospital and checking with the hospital staff to ensure realistic parameters
- Experiment with fast transits.