A Simple Multi-Objective Optimization Algorithm for the Urban Transit Routing Problem

Lang Fan, Christine L. Mumford and Dafydd Evans
School of Computer Science, Cardiff University, UK

Abstract— **The urban transit routing problem (UTRP) for public transport systems involves finding a set of efficient transit routes to meet customer demands. The UTRP is an NP-Hard, highly constrained, multi-objective problem, for which the evaluation of candidate route sets can prove both time consuming and challenging, with many potential solutions rejected on the grounds of infeasibility. In this paper we propose a simple evolutionary multi-objective optimization technique to solve the UTRP. First we present a representation of the UTRP and introduce our two key objectives, which are to minimise both passenger costs and operator costs. Following this, we describe a simple multi-objective optimization algorithm for the UTRP then present experimental results obtained using the Mandl's benchmark data and a larger transport network.**

I. Introduction

Public transportation systems play a very important role in the daily lives of city dwellers and commuters throughout the world; and with the increase in population and current concerns about the environment, efficient urban public transport system are urgently required. Thus, research on the urban transit network design problem (UTNDP) is timely. Generally the UTNDP is concerned with finding a set of routes and schedules for an urban public transport system. In practice, solving this highly complex, multiply-constrained NP-Hard problem involves considering many criteria in order to efficiently meet the needs of passengers while at the same time minimizing the cost to the operators. The different perspectives of users and operators, may result in different criteria for measuring efficiency: the passenger will require a cheap, fast and reliable service, while the operator will need to consider the cost of running the service.

The UTNDP can be subdivided into two components, namely the transit routing problem and the transit scheduling problem [4]. Generally, solving the urban transit routing problem (UTRP) involves devising efficient transit routes on an existing road (or rail) network, with known pickup/dropoff points (e.g., bus stops). On the other hand, the urban transit scheduling problem (UTSP) is charged with finding efficient schedules for the transport vehicles. Due to the complexity of the UTNDP, the two phases are usually implemented sequentially, with the routes determined in advance of the schedules.

In this paper we address the UTRP using an evolutionary multi-objective approach, trading off quality of service with operator costs. The work builds on our earlier metaheuristic framework for the UTRP [8], extending it to multiple objectives. Due to its intrinsic difficulty, metaheuristic algorithms are ideally suited to the UTRP, yet the success of such methods depend heavily on: (1) the quality of the chosen *representation*, (2) the effectiveness of the *initialization procedures* and (3) the suitability of the chosen *neighbourhood moves*. In [7], we proposed a simple model of the UTRP to evaluate candidate route sets from a quality of service viewpoint. Given the practical importance of the UTNDP, it is perhaps rather surprising that little work has concentrated on extracting generic features of the problem, formulating simplified models or devising benchmark data sets to facilitate comparative studies.

For multi-objective optimization problems, there is usually no single optimal solution. Instead, such problems tend to be characterized by a set of alternative solutions, each of which must be considered equivalent in the absence of further information regarding the relative importance of each of the objectives in the solution vectors. Such a solution set is called the *Pareto-optimal set*, and the objective values in the set are located at the *Pareto front*. Pareto-optimal solutions are *non-dominated solutions* in the sense that it is not possible to improve the value of any one of the objectives in such a solution, without simultaneously degrading the quality of one or more of the other objectives in the vector [21]. The urban transit routing problem (UTRP) is a multi-objective problem which usually involves several objectives, such as the total travel time and the total transfer time from the passenger's point of view, and the number of routes and the total sum of all the bus route lengths from the operator's point of view. In reality, transport planners have to develop transit routes based on the practical requirements specified by the various stakeholders, for example bus companies and local government.

In the remainder of this paper, we begin with a brief survey of relevant literature in Section II. Following this, in Section III we describe a mathematical representation of the problem, including the key objectives and constraints along with our method for evaluating candidate solutions. Next, in Section IV we present a simple multi-objective evolutionary algorithm, which is followed in Section V by experimental results which demonstrate the effectiveness of our scheme. Finally, in Section VI we present our conclusions and suggest some ideas for future work.

II. The Urban Transit Network Design Problem (UTNDP)

Historically, transport planners have developed reasonably efficient bus route networks and schedules by relying on past experience, simple guidelines and local knowledge. This approach may be adequate for small problems, but for large urban areas where the number of bus routes may be over a hundred and the number of bus stops in the thousands, these methods may not be sufficient to devise good route sets and schedules. For this reason, researchers have shown increasing interest in developing scientific methods for solving the UTNDP in recent decades.

Prior to 1979, the few papers published on the UTNDP considered only specific problem instances [13], [19]. In 1979, Christoph Mandl [14], [15], [16] considered tackling the problem in a rather more generic form, and his common-sense account of the UTNDP in [14] makes remarkably contemporary reading. Mandl concentrated on the UTRP, and developed a two-stage solution; first a feasible set of routes was generated, then heuristics were applied to improve the quality of the initial route set. The route generation phase involved first computing the shortest paths between all pairs of nodes, then seeding the route set with those shortest paths containing the largest number of nodes, respecting the position of any nodes designated as terminals. Unserved nodes were then either incorporated into existing routes in the most favourable way, or new routes created with unserved nodes as terminals. In this first phase, Mandl considered only in-vehicle travel costs when assessing route quality. The route improvement phase involved applying a number of heuristic methods to improve the initial route set: (i) obtaining new routes by exchanging parts of routes at an intersection node; (ii) including a node that is close to a route if travel demand between this node and the nodes on the route is high; and (iii) excluding a node from a route that is already served by another route if the travel demand between this node and the other nodes on the route is small. In this second phase, waiting times were also considered; these were determined by the associated vehicle frequencies.

Following Mandl's pioneering work, heuristic methods have been widely used to solve the UTNDP. In 1986, Ceder and Wilson [3] published a model for simultaneously solving the routing and scheduling problems. The model focused on two routines for generating and testing candidate route sets: *Level I* considered only the passenger's viewpoint, and was aimed at minimizing the total travel time, while *Level II* considered both the passenger's and the operator's viewpoint, balancing travel time and waiting time with the number of vehicles required. Vehicle frequencies and timetables were also set at *Level II*. More recently, Baaj and Mahmassani [2] proposed an heuristic route generation algorithm. Based on travel demand, the algorithm first determines an initial set of skeleton routes which are then expanded to form transit routes. The designer's knowledge and experience were also used to reduce the search space.

With the advancement of the computing technology over the last two decades, metaheuristic techniques have become increasingly popular for solving combinatorial problems. Approaches such as genetic algorithms (GAs), tabu search and simulated annealing have all played important roles in recent research on the UTNDP. Genetic algorithms are particularly popular, and several researchers have used them to determine a route network and the associated vehicle frequencies simultaneously. Pattnaik, Mohan and Tom [18], Tom and Mohan [20], and Agrawal and Mathew [1] all used a binary encoding scheme to represent candidate routes. Candidate routes are pre-determined and stored in a list, then a GA selects routes from this list to form a route set. In general, candidate route sets are initially produced using heuristic procedures, based on shortest path calculations moderated by user-defined guidelines. The genetic operators (mutation and crossover) produce new route set variations for selection, giving the population scope to improve over time, provided selection is biased towards good solutions over poor ones. With this approach, it is important that similar routes should be identified by similar binary codes, so that a simple mutation to a binary code for a particular route, for example, will tend to produce a mutated route having many nodes in common with its parent. In [1] and [20], vehicle frequencies are also encoded as part of the chromosome. On the other hand, Chakroborty and Dwivedi [5] used a different representation by listing the nodes explicitly, rather than binary coding a route as an entity. This work was taken further by Chakroborty [4] to cover both routing and scheduling. For other metaheuristic methods, examples can be seen in Fan and Machemehl's recent papers [9], [10], who used tabu search and simulated annealing to solve a particular class of UTNDP.

Very few research papers have been published on the use of *multi-objective* optimization techniques for the UTRP. Among those that exist, Ceder and Israeli [12] introduced a complex seven-stage approach, which includes several steps to create routes, identify transfers and calculate vehicle frequencies. A number of objectives such as travel time, waiting time, empty space and fleet size were then identified as a set of multi-objective tradeoff solutions to be presented to a human decision maker. Fan and Machemehl [9] proposed a multi-objective decision-making approach. The basic idea is to experiment with different values for the weights of three objective functions, in order to obtain a range of non-dominated results from which a human decision maker can select a suitable compromise solution. Although this multi-objective optimization method is able to find some good solutions to the UTNDP, determining suitable values for the weights requires a large number of experiments, which can be very time consuming. Thus a generic and computationally efficient multi-objective optimization method to solve the UTNDP is desirable.

III. The Routing Problem

The Urban Transit Routing Problem (UTRP) involves determining a set of efficient transit routes that meet the requirements of both passengers and the operator.

A. Problem Representation

The *transport network* is represented by a (undirected) graph $G(V, E)$ where the nodes $V = \{x_1, \ldots, x_n\}$ represent access points (e.g. bus stops), and the edges $E = \{e_1, \ldots, e_m\}$ represent direct transport links between two access points. A *route* can then be represented by a path in the transport network,

$$r_a = (x_{i_1}, \ldots, x_{i_q}) \quad \text{where} \quad i_k \in \{1, \ldots, n\}$$

and a solution to the UTRP is specified by a *route set*

$$R = \{r_a : 1 \leq a \leq N\}$$

We also define the *route network* associated with a route set to be the subgraph of the transport network containing precisely those edges that appear in at least one route of the route set.

To evaluate a candidate route set, we construct the associated *transit graph* $H = H(\widetilde{V}, \widetilde{E})$ as follows. First, for every node $x_i \in V$ in the transport network, the transit graph contains a node y_{ia} for each route passing through x_i,

$$\widetilde{V} = \cup_{i=1}^{n} \{y_{ia} : x_i \in r_a\}$$

Let $(y_{ia}, y_{jb}) \in \widetilde{E}$ denote the edge from node y_{ia} to node y_{jb} in the transit graph. The edges of the transit graph fall into two categories; *transport edges* correspond to transport links between two nodes on the same route,

$$\widetilde{E}_1 = \cup_{a=1}^{N} \big\{(y_{ia}, y_{ja}) : (x_i, x_j) \in E_a\big\} \qquad (1)$$

where E_a is the set of edges in route $r_a \in R$, while *transfer edges* correspond to transfers from one route to another,

$$\widetilde{E}_2 = \cup_{i=1}^{n} \big\{(y_{ia}, y_{ib}) : x_i \in r_a \cap r_b\big\} \qquad (2)$$

B. Constraints

Some broad criteria regarding what constitutes a good solution to the UTRP are the following [4]:

- The route network is connected.
- The passenger demand is completely satisfied.
- The average number of transfers from one route to another is as small as possible.
- The average journey time is as small as possible.

Because the UTRP is not concerned with scheduling, we will assume that there are sufficient vehicles on each route to ensure that the demand between every pair of nodes is satisfied, provided the route network is connected. (As a result, we must also assume that the transfer time is the same between any two routes in the route set.) Furthermore, each route should have a maximum length, based on considerations such as driver fatigue and the difficulty of maintaining the schedule [22]. Thus we say that a route set is *feasible* provided it satisfies the following constraints,

1. The route network is connected.
2. No route is longer than a predefined maximum length.

C. Evaluation

To evaluate a candidate route set, we consider the *passenger cost* and *operator cost*.

In general, passengers want to travel to their destination in the shortest possible time, but also want to avoid the inconvenience of having to make an excessive number of transfers. Let t_{ij} denote the time required to traverse the

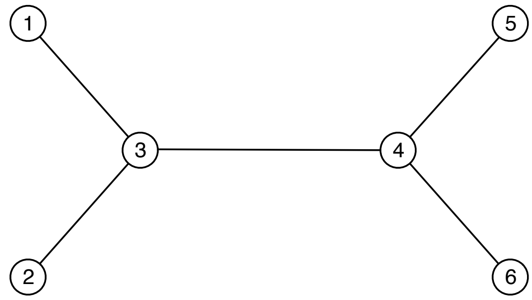

Fig. 1. Transport Network

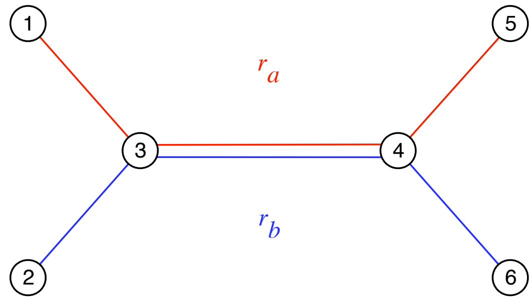

Fig. 2. Route Network

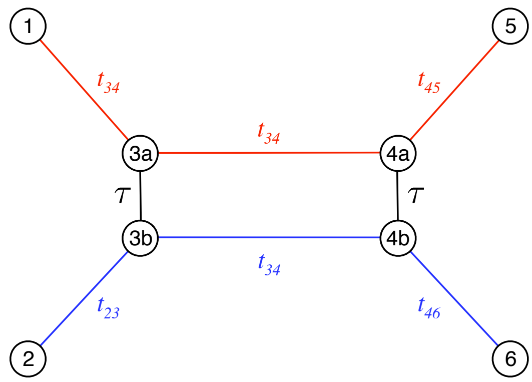

Fig. 3. Transit Network

transport link $(x_i, x_j) \in E$, let τ denote the time required to transfer from route r_a to route r_b, and let θ be an additional time penalty to account for the inconvenience of having to make a transfer (a passenger will only choose to transfer if this reduces the journey time by at least θ time units). For the transit graph $H(\widetilde{V}, \widetilde{E})$ defined above, let the transport edges $(y_{ia}, y_{ja}) \in \widetilde{E}_1$ have length t_{ij}, and let the transit edges $(y_{ia}, y_{ib}) \in \widetilde{E}_2$ have length $\tau + \theta$. Then the minimum journey time from x_i to x_j (including penalties due to transfers) is given by the length of the shortest path from a node in $\{y_{ia} : x_i \in r_a\}$ to a node in $\{y_{jb} : x_j \in r_b\}$ in the transit graph; let $\alpha_{ij} = \alpha_{ij}(R)$ denote the minimum journey time from x_i to x_j.

We assume that passengers will always choose to travel on the shortest-time paths (calculated by Dijkstra's algorithm [6] or Floyd's algorithm [11]) in the transit network. Let

d_{ij} denote the transit demand from node x_i to node x_j (defined in terms of the number of passengers wishing to travel between x_i and x_j). We define the *passenger cost* to be the total journey time over all passengers,

$$C_P(R) = \sum_{i,j=1}^{n} d_{ij}\alpha_{ij}(R) \tag{3}$$

where α_{ij} is the shortest journey time from x_i to x_j using route set R. The passenger cost is proportional to the *mean* journey time. It is beyond the scope of this paper to consider the *variance* of the journey times, which would guard against route sets where a small number of passengers are subject to excessively long journey times.

Certain operating costs, such as the number of vehicles required on each route to guarantee a reasonable quality of service, cannot be addressed without considering vehicle scheduling. However, other operating costs simply depend on the total length of the transport routes (for example, the average fuel cost per vehicle). Thus we define the operator cost to be the *total length* of the route set,

$$C_O(R) = \sum_{a=1}^{N} \sum_{(x,x')\in E_a} \|x - x'\|$$

where $\|x-x'\|$ is the length of the transport link $(x, x') \in E$.

The passenger cost C_P and the operator cost C_O will be traded off as conflicting objectives by our multi-objective optimization algorithm.

IV. A Simple Multi-Objective Optimization Algorithm for the UTRP

To solve the UTRP described above, we present the *Simple Multi-Objective Optimization* (SMO) algorithm, shown as Algorithm 1. This scheme is based on the SEAMO algorithm [21], [17], but without the crossover operator. Algorithm 1 relies on the *Make-Small-Change* procedure, described in [8], which modifies an existing route set to produce a new feasible route set. To apply this procedure, we must impose additional constraints. First of all, we need to set a minimum number of nodes per route to avoid empty or single node routes. Secondly, we ensure that each route is free of cycles and backtracks. The *Make-Small-Change* procedure restricts its attention to the following simple neighbourhood moves in the space of feasible solutions,

1) *a node is added to the end of a route*,
2) *a node is removed from the start of a route*.

The main drawback of adopting such simple neighbourhood moves means that the number of routes in the network has to be chosen prior to optimization.

Algorithm 1 *Simple Multi-Objective Optimization (SMO)*

Generate initial population of feasible route sets.
Calculate passenger and operator costs for each route set.
Record the best-route-set-so-far for both objectives.
repeat
  for each route set in the population
  Apply the *Make-Small-Change* procedure to produce a feasible offspring
  if offspring is a duplicate
  then delete offspring
  elseif offspring improves on either best-so-far
  then offspring replaces parent and best-so-far updated
  elseif offspring dominates parent
  then offspring replaces parent
  elseif offspring and parent are mutually non-dominated
  then find an individual in the population that is dominated by the offspring and replace it with the offspring.
  endif
  endfor
until the stopping condition is satisfied
print all non-dominated solutions in the final population.

V. Experimental Results

First we test the SMO algorithm on Mandl's Swiss road network [14], then compare the results with those previously published in [8]. To the best of our knowledge, Mandl's road network is the only generally available benchmark for the UTRP. Following this, we use the SMO algorithm and the algorithm described in [8] on a larger artificial instance of the UTRP, where we rely on "idealized costs" (see later) to evaluate the performance of our algorithms.

A. Mandl's Network

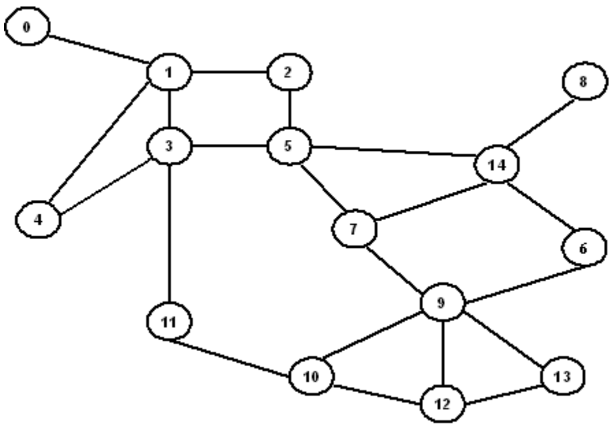

Fig. 4. Mandl's Swiss Road Network

Mandl's Swiss road transport network (see Figure 4) has 15 nodes and 21 links. In our experiments we consider four separate scenarios, namely route sets consisting of 4 routes, 6 routes, 7 routes and 8 routes, with a maximum of 8 nodes in each route. For each scenario, we record the results from 10 replicate runs (each seeded with different random numbers) using a population size of 200. The number of iterations of the multi-objective algorithm used in each experimental run for 4, 6, 7, 8 routes, was 1000, 3000, 4000, 5000 respectively. For each set of experimental runs, we accumulate the results and isolate the non-dominated solutions, trading off passenger cost C_P against the operator cost C_O. Finally, we validate our solutions against our recently published results [8],

Table I shows the best route sets obtained for passengers and operators respectively for each of the 4 scenario. Note that long routes are better for passengers, and short routes for operators.

TABLE I
ROUTES OBTAINED USING THE SMO ALGORITHM

Route No.	Best routes for passengers	Best routes for operators
4	13-12-10-9-7-5-3-4	1-2-5-7-14-6-9-10
	1-3-11-10-9-6-14-8	14-8
	10-9-7-5-2-1-0	13-12-10-11
	4-1-2-5-14-6-9	4-3-1-0
6	12-10-9-7-5-2-1-0	1-2-5-7-14-6-9
	6-14-5-2-1-3-4	8-14
	9-7-5-3-4	1-0
	12-13-9-10-11-3-1-0	13-12-10-9
	9-6-14-8	1-3-4
	11-10-12-13	10-11
7	6-14-7	8-14
	13-12-10-11-3-1-2	10-11
	11-10-9-6-14-8	3-4
	13-9-6-14-5-3-4	13-12-10
	9-7-5-3-4-1-2	1-3
	0-1-2-5-7-9-10-12	0-1-2-5-7-14-6-9
	3-1-0	10-9
8	1-3-11-10-12-13-9	3-1
	11-10-12-13-9-6-14-5	2-1
	4-1-2-5-7-14-8	1-0
	0-1-2-5-7-9-10-12	12-10
	11-10-12-13-9-7-5-3	3-4
	3-5-14-8	14-8
	4-3-5-7-9-10-12	11-10-9-6-14-7-5-2
	11-10-9-6-14-5-2-1	12-13

The following parameters [4] are used to evaluate our best route sets found by the multi-objective optimization algorithm:

- d_0 - Percentage demand satisfied without any transfers.
- d_1 - Percentage demand satisfied with one transfer.
- d_2 - Percentage demand satisfied with two transfers.
- ATT - Average travel time (minutes per passenger), including a penalty of 5 minutes per transfer.
- C_O - Operator cost function value (length of the route set)

The values of these parameters for the route sets from Table I are presented in Table II. To validate these results we compare them to our previously published results [8] for the single objective situation, minimizing passenger costs.

From Table II, it is clear that a group of good route sets has been found from the passengers' point of view. The parameter values are very close to our previously published results for a single objective. However, the operator's costs in the "best for passenger" column, are consistently better than the corresponding costs obtained using our previous single objective approach. For the best route sets from the operator's perspective, it is reasonable that the lowest operator cost will correspond with the highest passenger's cost. In general, our multi-objective optimization algorithm can solve Mandl's network problem well. For example, for the four scenarios, assessment parameters of the best route sets for passengers' costs, such as d_0 and ATT, obtained by the SMO are close to our previously published results. We calculated the percentage difference of the SMO results relative to the previously published results, and found these values to lie between 0.83% and 3.40% for d_0, and between 0.19% and 6.33% for ATT. On the other hand, for the best route sets for operator's costs, in the 4 scenarios the total length of these best route sets are the same, namely 63 minutes (evaluated by bus travel time). Non-dominated trade-off solutions from 10 runs for the 4-route scenario can be seen in Figure 5.

TABLE II
THE BEST RESULTS OBTAINED BY OUR SMO ALGORITHM (1) FROM THE PASSENGERS' POINT OF VIEW AND (2) FROM THE OPERATOR'S POINT OF VIEW.

Routes No.	Parameters	Previously Published [8]	Best for passenger	Best for operator
4	d_0	93.26	90.88	61.08
	d_1	6.74	8.35	36.61
	d_2	0.00	0.77	2.31
	ATT	11.37	10.65	13.88
	C_O	147	126	63
6	d_0	91.52	93.19	66.09
	d_1	8.48	6.23	30.38
	d_2	0.00	0.58	3.53
	ATT	10.48	10.46	13.34
	C_O	215	148	63
7	d_0	93.32	92.55	65.64
	d_1	6.36	6.68	26.20
	d_2	0.32	0.77	8.16
	ATT	10.42	10.44	13.54
	C_O	231	166	63
8	d_0	94.54	91.33	59.92
	d_1	5.46	8.67	21.97
	d_2	0.00	0.00	18.11
	ATT	10.36	10.45	13.57
	C_O	283	245	63

B. Experiment on a Larger Network

The larger network we tested has 110 nodes and 275 links, and the total travel demand on the network is 3603360 journeys per day (The data sets with the travel time and demand matrices published on our website: http://users.cs.cf.ac.uk/L.Fan/). Although this network was randomly generated by ourselves, we based its size, connectivity and other features on a bus route map for a major British city. The transit time between each pair of nodes is measured in minutes, and the transfer penalty is set to 5 minutes, as previously. Clearly, compliance with the connectivity and coverage constraints, requires that each node is included in at least one route of a candidate route set, and further that each route has at least one node in common with another route (provided there are two or more routes in a route set). Thus, there is an interrelationship between maximum and average route lengths, on the one hand, with the number of routes in the route set, on the other. For our

TABLE III

COMPARISON RESULTS FOR 110 NODES PROBLEM. MOO DENOTES MULTI-OBJECTIVE OPTIMIZATION

Scenario	Parameters	Best Single objective (passenger)	Best MOO passenger	Best MOO operator
I	d_0	72.91	71.26	48.62
	d_1	20.56	18.88	32.45
	d_2	6.54	9.86	18.93
	ATT	34.60	35.29	37.36
	C_O	2986	2823	1077
II	d_0	71.21	70.01	46.97
	d_1	20.71	19.21	31.84
	d_2	8.08	10.78	21.19
	ATT	35.68	35.89	37.55
	C_O	2378	2257	1265

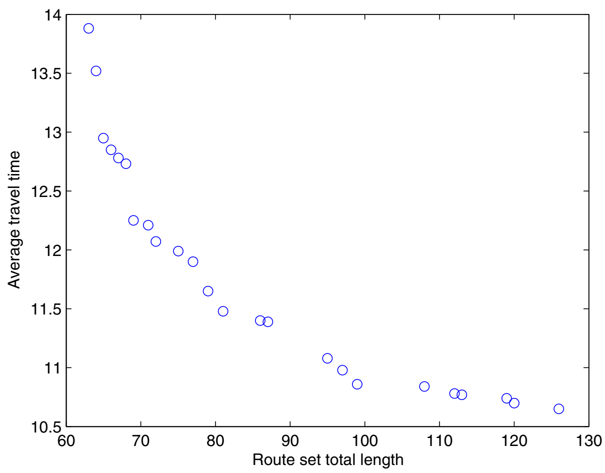

Fig. 5. Non-dominated solutions from 10 runs for 4-route scenario for Mandl's network (route length and travel time both measured in minutes)

experiments we model two scenarios: I) using theoretical considerations, and II) by examining our route map for the British city upon which this data is based. For simplicity, we assigned maximum and minimum route lengths according to the number of nodes in a route, rather than total travel times along the route for our computational experiments.

In Scenario I, we set the minimum route length at 2 nodes, because a route cannot be shorter than this. A maximum route length of 29 was derived by evaluating all the shortest path distances on the transit network, and recording the maximum number of nodes for a path. The total number of routes was set at 55, based on some observations we made on *node frequency* (i.e., the average number of routes in which each node makes an appearance). We examined route maps for various cites of the world, and found remarkable consistency with respect to node frequency. Scenario II, is based on actual route map observations. For this situation, we set the total number of routes to 56, and the minimum and maximum number of nodes in each route to 10 and 22 respectively.

We compare the results obtained using our multi-objective algorithm with those obtained using the single objective approach we used previously in [8]. To do this involved running our single objective software on our new data set for the purpose of the present paper, since we did not publish results for this large data set in our earlier paper. The values of these assessment parameters introduced before for the best route sets obtained by the single objective approach and our multi-objective algorithm can be seen in Table III.

In addition, we also assess the quality of our results by comparing the average travel time (ATT) per passenger with the Lower Bound on the ATT obtained by assuming that every passenger travels on the shortest path on the *transport* network (as opposed to the *route* network) between source and destination without any transfers. The difference between the ATT and the Lower Bound is expressed as a percentage of the Lower Bound. These results for Scenarios I and II can be seen in Table IV.

TABLE IV

COMPARING BEST SOLUTIONS WITH "IDEAL" SOLUTIONS.

Scenario	Lower Bound on ATT	Best ATT obtained by SMO	ATT error %
I	33.84	35.29	4.28
II	33.84	35.89	6.06

On the other hand, from these results, it is clear that our multi-objective optimization algorithm can also find good solutions for these larger network problems. For example, considering assessment parameters (d_0 and ATT) of the best route sets for passenger's costs for the two scenarios, the percentage difference between d_0 obtained by our SMO and d_0 obtained by the single objective method are between 1.69% and 2.26%. At the same time, the percentage difference between ATT obtained by the SMO and ATT obtained by the single objective method are between 0.59% and 1.99%. Note - percentage difference:

$$\|(P_{SMO} - P_{singleobjective})\| \times 100/P_{singleobjective} \quad (4)$$

where P denotes the assessment parameter.

VI. Conclusion

In this paper, we have presented an evolutionary multi-objective optimization algorithm to solve the urban transit routing problem. Our research framework includes a model for the UTRP, a method for representing candidate solutions, a procedure for generating feasible route sets, and a routine for making intelligent neighbourhood moves. These components play an important role in the multi-objective optimization process. Through the experiments on Mandl's benchmark data set, and a much larger data set, we demonstrate that we are able to obtain efficient route sets which balance the requirements of passengers with those of the operator. In future, we intend to refine our model in consultation with planners and other stakeholders. We also intend to develop more sophisticated types of neighbourhood move which provide a mechanism for creating or destroying routes, so that the number of routes will not need to be chosen prior to optimization. At the same time, we also intend to develop improved metaheuristic algorithms to support multi-objective optimization in the UTNDP.

References

[1] J. Agrawal and T. V. Mathew, Transit Route Network Design Using Parallel Genetic Algorithm. Journal of Computing in Civil Engineering, pp. 248-256, 2004.

[2] M. H. Baaj and H. S. Mahmassani, Hybrid Route Generation Heuristic Algorithm for The Design of Transit Networks, Transportation Research, vol. 3(1), pp. 31-50, 1995.

[3] A. Ceder and H. M. Wilson, Bus Network Design, Transportation Res. B, vol. 20B(4), pp. 331-344, 1986.

[4] P. Chakroborty, Genetic Algorithms for Optimal Urban Transit Network Design. Comput. Aided Civil Infrastruct. Eng, vol. 18, pp. 184-200, 2003.

[5] P. Chakroborty and T. Dwivedi, Optimal Route Network Design For Transit Systems Using Genetic Algorithms, Engineering Optimization, vol. 34(1), pp. 83-100, 2002.

[6] E. W. Dijkstra, A note on two problems in connexion with graphs. Numer. Math., vol. 1, pp. 269-271, 1959.

[7] L. Fan and C. L. Mumford, A Simplified Model of the Urban Transit Routing Problem, The 7th Metaheuritstics International Conference, Montreal, Canada, 2007.

[8] L. Fan and C. L. Mumford, A Metaheuristic Approach to the Urban Transit Routing Problem. Journal of Heuristic. Online First, 2008.

[9] W. Fan and R. B. Machemehl, A Tabu Search Based Heuristic Method for the Transit Route Network Design Problem. In: The 9th International Conference on Computer-Aided Scheduling of Public Transport, San Diego, California, 2004.

[10] W. Fan and R. B. Machemehl, Using a Simulated Annealing Algorithm to Solve the Transit Route Network Design Problem. Journal of Transportation Engineering, pp. 122-132, 2006.

[11] R. W. Floyd, Algorithm 97: Shortest Path, Commun. ACM, vol. 5(6), pp. 345, 1962.

[12] Y. Israeli and A. Ceder, Designing Transit Routes at the Network Level. In: IEEE Vehicle Navigation and Information Systems Conference, pp. 310-316, 1989.

[13] W. Lampkin and P. D. Saalmans, The Design of Routes, Service Frequencies and Schedules for a Municipal Bus Undertaking: a case study, OR Quarterly, vol. 18, pp. 375-397, 1967.

[14] C. E. Mandl, Applied Network Optimization. Academic, London, 1979.

[15] C. E. Mandl, Evaluation and Optimization of Urban Public Transport Networks, Third Congress on Operations Research, Amsterdam, Netherlands, 1979.

[16] C. E. Mandl, Evaluation and Optimization of Urban Public Transport Networks, European Journal of Operational Research, vol. 5, pp. 396-404, 1980.

[17] C. L. Mumford, Simple Population Replacement Strategies for a Steady-State Multi-Objective Evolutionary Algorithm, Genetic and Evolutionary Computation Conference (GECCO), Seattle, Washington, USA, 1389-1400, 2004.

[18] S. B. Pattnaik, S. Mohan, and V. M. Tom, Urban Bus Transit Route Network Design Using Genetic Algorithm, Journal of Transportation. Eng., pp. 368-375, 1998.

[19] L. Simman, Z. Barzily, and U. Passy, Planning the Route System for Urban Buses, Comput. Ops. Res., vol 1, pp. 201-211, 1974.

[20] V. M. Tom and S. Mohan, Transit Route Network Design Using Frequency Coded Genetic Algorithm. Journal of Transportation. Eng., pp. 186-195, 2003.

[21] C. L. Valenzuela, A Simple Evolutionary Algorithm for Multi-Objective Optimization (SEAMO). In: Congress on Evolutionary Computation (CEC), Hawaii, 2002.

[22] F. Zhao and A. Gan, Optimization of Transit Network to Minimize Transfers. Final Report, Lehman Center for Transportation Research, Florida International University, 2003.