Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

A MULTI-OBJECTIVE GENETIC ALGORITHM USING INTERMEDIATE FEATURES OF SIMULATIONS

Hidemasa Muta
Rudy Raymond
Satoshi Hara
Tetsuro Morimura

IBM Research - Tokyo
5-6-52, Toyosu, Koto
Tokyo, 1358511, JAPAN

ABSTRACT

This paper proposes using intermediate features of traffic simulations in a genetic algorithm designed to find the best scenarios in regulating traffic with multiple objectives. A challenge in genetic algorithms for multi-objective optimization is how to find various optimal scenarios within a limited decision time. Typical evolutionary algorithms usually maintain a population of diversified scenarios whose diversity is measured only by the final objectives available at the end of their simulations. We propose measuring the diversity by also the time series of the objectives during the simulations. The intuition is that simulation scenarios with similar final objective values may contain different series of discrete events that, when combined, can result in better scenarios. We provide empirical evidence by experimenting with agent-based traffic simulations showing the superiority of the proposed genetic algorithm over standard approaches in approximating Pareto fronts.

1 INTRODUCTION

Traffic controllers often face daunting tasks in managing their cities' traffic for impending future events, such as major sporting events or natural disasters. They have to find optimal traffic control scenarios with regards to various objectives that reflect the impact of the scenarios against the city's economic and environmental key performance indicators (KPIs). They have to rely on traffic simulations to evaluate various control scenarios, and multi-objective optimization methods to quickly find as many best scenarios as possible for rapid decisions.

Agent-based traffic simulation is one of the popular methods to evaluate traffic scenarios for such work (Gomes, May, and Horowitz 2004; Behrisch et al. 2011; Osogami et al. 2012). It does not require the full knowledge of the macro-dynamics of the traffic. The traffic controllers only need to encode rules for the micro-behavior of the drivers (agents) in their cities, run simulations on their scenario settings, and measure the emerging macro-level results. The agent-based traffic simulations thus allow the operators to evaluate many scenarios that never happened before. Another advantage is that, being a discrete event simulation, it can give the operators detailed pictures about what will happen up to and in the transition from expected events by looking at the intermediate results. However, there are large computational costs that grow at least linearly with the number of cars (agents) in the simulations, and exponentially with the number of types of plausible actions.

978-1-4799-7486-3/14/$31.00 ©2014 IEEE

Evolutionary algorithms are methods of choice to quickly find optimal solutions to these kinds of complex multi-objective problems (Konak et al. 2006). A Genetic Algorithm (GA) is a heuristic search that finds optimal solutions by mimicking natural selections. It starts with a population of candidate solutions (or, individuals) that is evolved by recombination, mutation, and selection operators. The design of such operators is crucial for a GA to find a diverse set of optimal solutions after several evolutions of population (Fonseca and Fleming 1993; Horn et al. 1994; Konak et al. 2006). Traffic controllers can utilize an agent-based traffic simulation and a GA to generate optimal solutions for their tasks as follows. They can encode their traffic-regulation scenarios, such as, road closures, one-way restrictions, and others, in the genomes of individuals in the GA. The recombination and mutation now correspond to combining traffic-regulation scenarios and random changes of regulations, respectively. The traffic simulation is then used to evaluate the values of the simulation scenarios of the individuals. For multi-objective optimization, a selection operator must ensure a diversity of population so that optimal Pareto solutions can be found as many as possible within a limited number of generations. Standard approaches treat simulations as black boxes in the selection operator, namely, scenarios with similar final objectives are often filtered to leave spaces for other individuals surviving into the next generation (Delahaye et al. 2005; Schmöcker et al 2008; Ikeda et al 2007; Brown et al. 2013).

In this paper, we present a novel idea of using intermediate changes of objectives values in an agent-based simulation for a diversity measure to obtain such Pareto-optimal solutions. The idea is based on an observation that even though two simulations have similar final objective values, they are often caused by different underlying physical phenomenon. For example, if the final objectives are the length of traffic jams in a city center, similar values might actually be caused by traffic jams at different parts of the city center. Thus, combining the corresponding two scenarios might turn into distributing traffic to different parts of the city, and hence better final objectives. We propose to refine the diversity measure in a genetic algorithm to incorporate the intermediate changes of objective values that are easily available from an agent-based traffic simulation. Simulations with similar changes of objective values are more likely to have similar underlying phenomenon, and are better candidates to be filtered. On the other hand, although having similar final objectives, those with different changes of objectives when combined might result into better objectives.

We also show by experiments the superiority of the proposed genetic algorithms against standard ones for multi-objective optimization in a large-scale agent-based traffic simulation. We show that the proposed method can produce near Pareto-optimal solutions with less number of generations. This translates into the traffic operators' ability to quickly find optimal scenarios coping with various KPIs in their tasks anticipating future events. As a side contribution, we also propose a new quality measure of multi-objective solutions against Pareto-optimal solutions. We believe the measure might be useful for evaluating other multiple criteria decision makings.

One of the state of the art to search Pareto optimal solutions with simulation is the study by Brown et al. (2013). They used a multi objective genetic algorithm to study bridge retrofit scenarios for a highway network. As in our study, they evaluated the scenarios by agent-based traffic simulations to search for a range of Pareto-optimal solution sets. However, they only used the final objective value of the traffic simulation as the fitness value of each individual solution. Thus, their genetic algorithm is different from our approach because it does not consider historical objectives or intermediate features to preserve diversity. In this paper, we show the advantage of our approach which brings not only the variety of solution sets but also the quick convergence to Pareto-optimal front.

The remainder of this paper is organized as follows. Section 2 describes the concept of what-if traffic simulation for decision making, the role of an agent-based traffic simulation, and the traffic regulation scenarios considered. Section 3 presents key innovations of using intermediate objectives with regards to previously known approaches in multi-objective optimization, as well as the new measure of the quality of a solution set. Section 4 shows experimental results of the proposed method against the state-of-the-art. Section 5 presents related work, while concluding remarks and future work are listed in Section 6.

2 WHAT-IF SIMULATIONS FOR DECISION MAKING

Agent-based simulations are popular for analyzing what-if scenarios in complex processes, such as traffic and social systems. One of the reasons is that because one does not need to have the full knowledge of macro dynamics within such complex systems to test various scenarios that are nevertheless difficult or impossible to test in the real-world. It suffices to define micro rules of agents in the simulations, that define how they interact with the environment and among each other, change the parameters according to the scenario settings, and then run the simulation to see how the interaction of a large number of agents gives rise to macro-level results, such as traffic jams in traffic simulations, or number of products sold in customer-behavior simulations. However, when there are a large number of scenarios to test, agent-based simulations face severe limitation because each run of a large-scale agent-based simulation requires high computational cost. A search algorithm that is capable of efficiently exploring the solution space within possible scenarios must be utilized so that good scenarios can be identified within a limited decision time.

Hereafter, for ease of explanation we describe a typical such case in evaluating scenarios with an agent-based traffic simulation. It is straightforward to see that similar cases exist in other types of agent-based simulations, and with some modifications our approach can be extended to cope with such cases.

2.1 Agent-based Traffic Simulation

We are studying the use of agent-based simulations for the city traffic control. Suppose we have a critical traffic accident on a major road in the city. What kinds of traffic regulation scenario is effective to minimize the traffic congestion? These scenarios require making quick and effective decisions and sending the instructions for traffic control to the local police officers as soon as possible based on the real-time traffic simulation results. To do this, it requires the combination of an agent-based traffic simulator that is capable of evaluating the results of enforcing traffic-regulation scenarios by taking into account typical traffic patterns in the city, and a highly-parallelized population-based evolutionary optimization algorithm to obtain optimal scenarios as quickly as possible.

In our work, we use the IBM Mega-Traffic-Simulator that has often been used to evaluate various what-if scenarios in many cities in the world (Imamichi and Raymond 2013; Suzumura et al. 2012; Osogami et al. 2012; Osogami et al. 2013). Some of the unique features of the Mega-Traffic-Simulator are that it can build a traffic simulation model from traffic patterns of a city by estimating parameters from probe-car data, and it runs on massively parallel computers that enables it to simulate the microscopic traffic flows in the scale of arbitrary city or even the whole Japan (Suzumura et al. 2012). We use the simulator to evaluate the result of traffic-regulation scenarios.

We explain the types, the objectives, and the multi-criteria aspects within the traffic simulation scenarios that give rises to the complexity in finding good scenarios. Hereafter, for explanatory purpose, we will use the case of Nairobi with similar settings as (Imamichi and Raymond 2013) that generates traffic demands, or Origin-Destination (OD) data, based on (Gonzales et al. 2009). It is straightforward to apply our proposed method for traffic simulations in other cities with different settings of scenarios.

2.2 Traffic Regulation Scenarios

There are many possible ways to regulate traffic flows on the roads of a city. For simplicity, we consider scenarios in the form of directing the traffic on a road by one of four actions: *no intervention*, *one-way*, *reverse one-way*, and *close*. If there are r number of roads where those four actions can be executed, then there are 4^r possible scenarios that need to be evaluated by the traffic simulator. It is easy to see that the number of possible scenarios grows exponentially with the locations, or roads, and types of actions, such that it can quickly become impossible to evaluate all scenarios using agent-based simulations.

To show how our proposal can help identifying best scenarios, we will present two cases of traffic regulation scenarios upon a traffic accident on a road near the *Central Business District* (CBD) of Nairobi (see the *cone* mark in left parts of Figures 1 and 2): the small case where we know all optimal scenarios by simulating all possible scenarios, and the large case where the optimality of scenarios obtained is not guaranteed. We will compare the performances of our proposal on both the small and large cases. We will show that our proposal outperforms standard approaches in the small case, and show similar trends in the large case.

The small case consists of scenarios on applying one of the four actions on 5 roads nearby the location of the traffic accident, as shown in the left part of Figure 1 with "!" marks. Clearly, there are 4^5 =1024 possible scenarios by the combination of the four actions on the five roads. One of the scenarios is shown in the right part of the figure.

The large case consists of scenarios considering more number of roads regulated with the four actions as shown in Figure 2. There are 16 roads on a larger portion of area surrounding the accident site, and thus 4^{16} possible different scenarios that need to be simulated. There are now more than a million times as many as scenarios of the small case. One of the scenarios is shown in the right part of Figure 2.

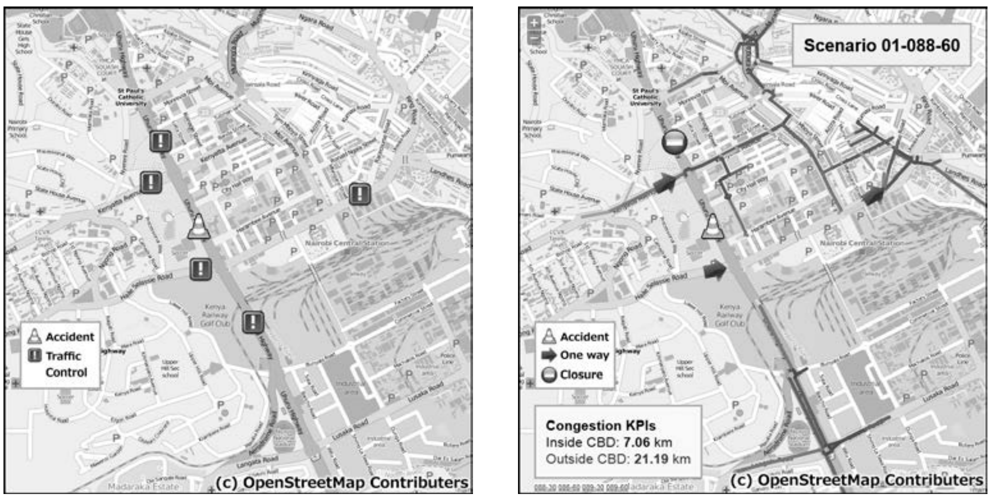

Figure 1: The small case: regulating traffic on 5 roads in Nairobi

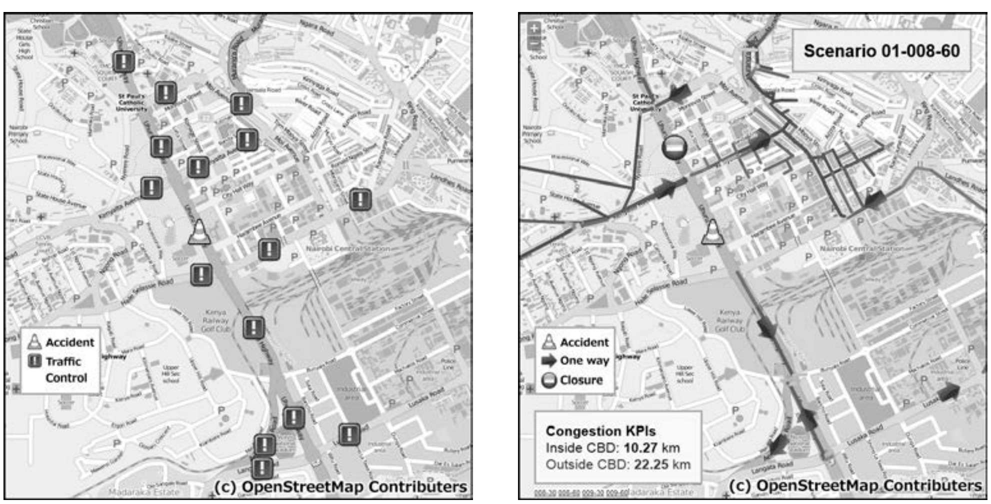

Figure 2: The large case: regulating traffic on 16 roads in Nairobi

2.3 Objectives

IBM Mega Traffic Simulator, like other agent-based traffic simulators, provides detailed information on the states of agents (cars), such as their speed, and macro-level statistics on each road, such as the number of cars passing through it, at every atomic time step of the simulation. This allows us to find roads with slow vehicles (slower than 10km/h) and calculate the total length of those roads. As the multi objectives of the scenarios, we use the total length of such roads for the urban area (inside the CBD) and the suburban area (outside the CBD) as shown in Figure 3. These two objectives often conflict with each other. If the traffic regulations reduce congestion inside the CBD, then they often increase the congestion outside the CBD, and vice versa. The decision makers must decide a scenario from the multiple Pareto-optimal solutions based on this trade-off relationship, and they need a multi-objective optimization method for the decision.

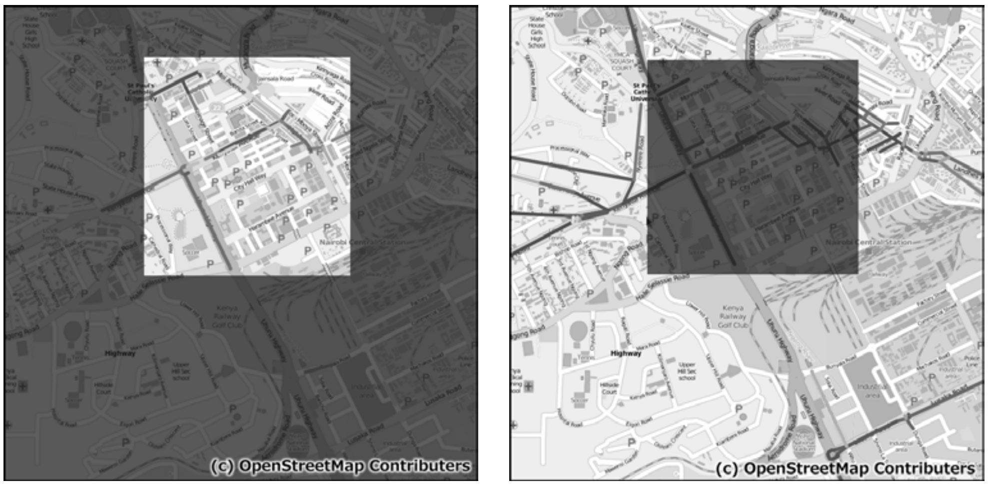

Figure 3: Areas to calculate objective (congestion length) inside and outside of the CBD of Nairobi

2.4 Multi-Objective Optimization

Population-based evolutionary optimization is popular and effective to quickly search for (near) optimal solutions of multi-objective optimization problems. They are easily implemented in parallel on an SMP server or multi-node cluster system. In our work, we used a Genetic Algorithm, or GA, (Goldbergs 1989) which is one of the most used population-based evolutionary algorithms. A GA evolves randomly generated individuals for highly optimized environmental fitness. It recombines, mutates, and selects individuals over many generations in a way that is quite similar to natural selection. The individuals encode explanatory variables in their genes. When we evaluate the fitness value of an individual, the explanatory variables are decoded from the genes, and then applied to an objective function, simulation, or other system with objectives.

In the work described here, we used an agent-based traffic simulation to evaluate the objective values of simulation scenarios. When there are multiple and conflicting objective values, then the individuals will eventually reach Pareto-optimal solutions. The ultimate goal of the multi-objective GA is to provide a range of solutions that are close to the Pareto-optimal front. The diversity of the solutions is very important, since the decision maker is required to understand the trade-off relationships of the objectives based on these solutions, and then select one of them based on priorities or preferences that often change over time. Unfortunately, it is difficult to find diverse solutions with a standard GA. Without diversity maintenance techniques, the solutions tend to converge at the center of Pareto-optimal front when we accelerate the convergence by prioritizing parents.

3 INTERMEDIATE FEATURES FOR QUICKLY GENERATING DIVERSE NEAR OPTIMAL MULTI-OBJECTIVE SOLUTIONS

In a multi-objective optimization problem, finding diverse (near) optimal solutions is important to balance various criteria of optimization. The diversity maintenance of populations is a key to improve the performance of a multi-objective genetic algorithm, and several ideas have been proposed (Konak, Coit, and Smith 2006; Deb 2012). Most of them based on these two methods and their combination: Pareto ranking (Fonseca et al. 1993) and Sharing (Horn et al. 1994). The *Pareto ranking* is used to rank solutions from inside (or, outside) of the objective space, while *Sharing* is used to discard solutions that resemble, or are near to, other solutions that are already in the population. To our best knowledge, all known ideas rank and measure the proximity of solutions based on the end results, namely, the congestion length at the end of traffic simulations for our cases.

In this paper, we propose a novel diversity-maintenance technique that discards solutions based on the distance of the intermediate features of the simulation process, as illustrated in Figure 4, where we use the traffic simulation to obtain time-series of objective values (the congestion lengths) and evaluate the individual fitness. Suppose we find two solutions that have similar final objective values but in different locations. Even though they have different temporal characteristics, the standard Pareto ranking and sharing approaches will likely discard one of them.

In contrast, our proposed method compares the two solutions in detail in the direction of the temporal axis so that we can reduce the risk of discarding a unique solution. Notice that the simulator runs each scenario for one hour. In our approach, we implement a diversity-distance calculation method that compares time-series of historical objectives (the congestion lengths) obtained at every 10 minutes during a 1-hour simulation among individuals. The distance between two individual is the sum of the absolute differences of their objectives. We can assign different weights to objectives based on their time stamps. For example, assigning weight value 1.0 to final objectives and zero to other (intermediate) objectives is similar to the standard approaches. We assign non-zero weights to all objectives, final and intermediate, so that those with later time stamps are weighed more.

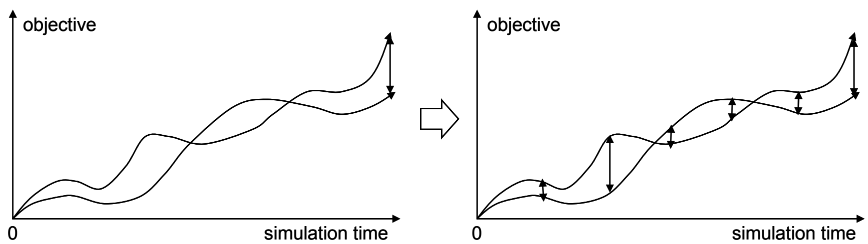

Figure 4: Measuring distance: by final objective only (left) versus historical objectives (right)

4 RESULTS

We will see how the new comparison method supports a variety of Rank1 solutions that are closer to Pareto-optimal solutions, and finds those solutions earlier than standard comparison methods. We denote the standard GA as the *standard GA*, the GA that uses the final objectives to compare its individual as the *final-objective GA*, and the GA that uses our comparison method as the *historical-objectives GA*.

The final-objective GA utilized the Pareto ranking and Sharing by distance between final objectives, which are the congestion length of roads inside and outside the CBD. The historical-objective GA, our proposed method, also utilized the Pareto ranking, but Sharing is now by the sum of absolute differences of historical objectives, namely, the congestion lengths at 10-minute intervals, weighted by the time stamps of the historical objectives. Equation (1) is the formula to compute the distance between the individual i and the individual j, or $d(i,j)$. Equation (2) denotes the (normalized) weighting scheme of the distance, with the differences of the final objectives are weighted thrice as large as the first objectives, while the rest are interpolated between them.

$$d(i,j) = \sum_{t=t_1}^{t_n} | f(i,t) - f(j,t) | \, w(t) \quad (1) \qquad w(t) = \frac{1}{n}\frac{t-t_1}{t_n - t_1} + \frac{1}{2n} \quad (2)$$

[New Measure of the Quality of Multi-Objective Solutions]

To measure both the diversity of a solution set S and its distance from the Pareto-optimal solution set T at the same time, we propose the following measure that we call the *uniform-cover distance* of S with regards to T, denoted by *ucd(S|T)* as follows.

$$ucd(S \mid T) = \max\{dist(t, S) \mid t \in T\}, \text{ where } dist(t, S) = \min\{\| s - t \| \mid s \in S\} \quad (3)$$

In Equation (3), the set *S* is the solution set that is obtained by running a GA after several changes of population, or generations, while the set *T* is the true Pareto-optimal solution. For the small case with all the objectives of all scenarios can be evaluated, we can obtain *T*. However in general, *T* is unknown and hard to obtain. We will show later how we can approximate *T* for the large case where evaluating all scenarios is nearly impossible.

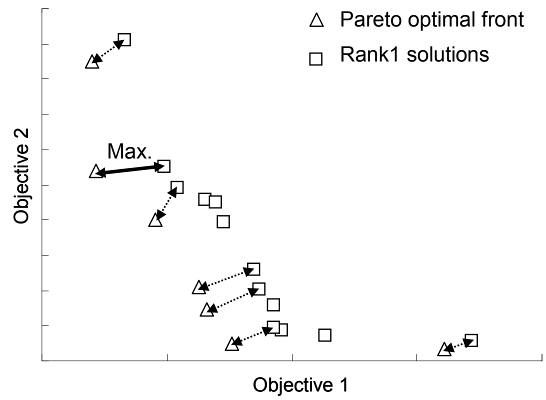

Figure 5: Uniform-Cover Distance (UCD) from Pareto-optimal front

Let us give an intuitive explanation of the uniform-cover distance by using Figure 5. Assume that S is the set of Rank1 solutions obtained from a GA. For each Pareto-optimal solution, the uniform-cover distance guarantees the existence of a nearly-optimal solution in S. Namely, if $ucd(S|T) = \delta$, then the GA can generate solutions whose multi-objective values are within δ from the optimal ones. On the other hand, to achieve a uniform-cover distance δ, a GA must generate a solution set that contains at least one solution within the distance δ uniformly for every Pareto-optimal solution. We can also replace the "max" in Equation (3) with the average for all $t \in T$, which we call the *averaged UCD* for simplicity

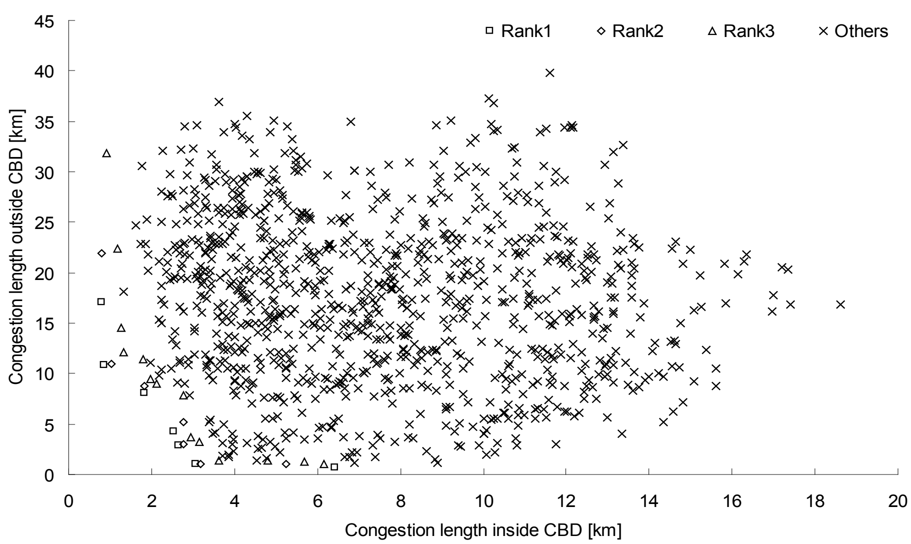

Figure 6: Full search result for the small case

4.1 The Small Case: Four Types of Actions on Five Locations

We compared the standard GA, final-objective GA, and the historical-objectives GA on the small case, which has only 5 locations to impose 4 actions to regulate traffic. Because there are only 1024 different scenarios, we can evaluate all simulation scenarios and obtain the objective values that are the congestion length after imposing the corresponding action of the scenario for one hour. Figure 6 depicts the distribution of the multi-objective values of all 1024 scenarios. The coordinate of each point in the figure denote the congestion length inside and outside CBD. It is easy to obtain the Pareto-optimal solutions from this case, as confirmed in the figure. (Notice that the traffic simulator is actually a stochastic one, and therefore the objective values each time it runs on the same setting differ. However, by a preliminary test, we found out that their differences are small. Figure 6 only depicts the value of objectives for one run per simulation scenario.)

We ran the three GAs for 100 repetitions of different initial populations. All of the GAs started from the same initial population that contained 50 individuals. The populations of the GAs were evolved with crossover rate 0.2 for 40 generations, and their sizes were kept constant at 50. We set the mutation rate to zero because with non-zero mutation it would be difficult to distill the impact of selection from the mutation.

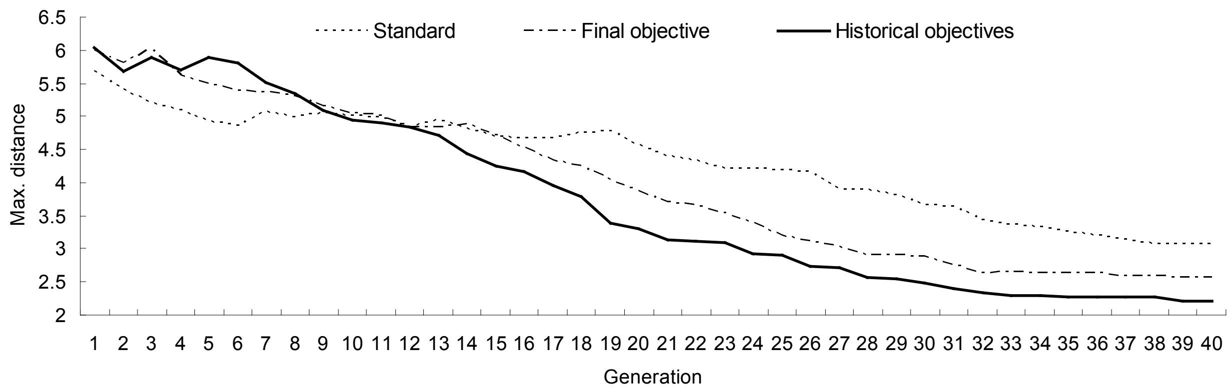

Figure 7: Uniform-Cover Distance (UCD) comparison (average of 100 trials using the small case)

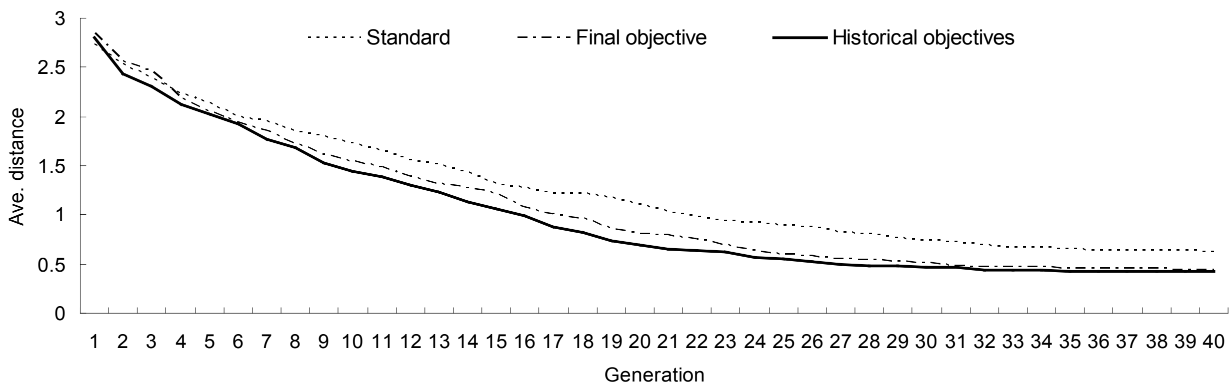

Figure 8: Average UCD comparison (average of 100 trials using the small case)

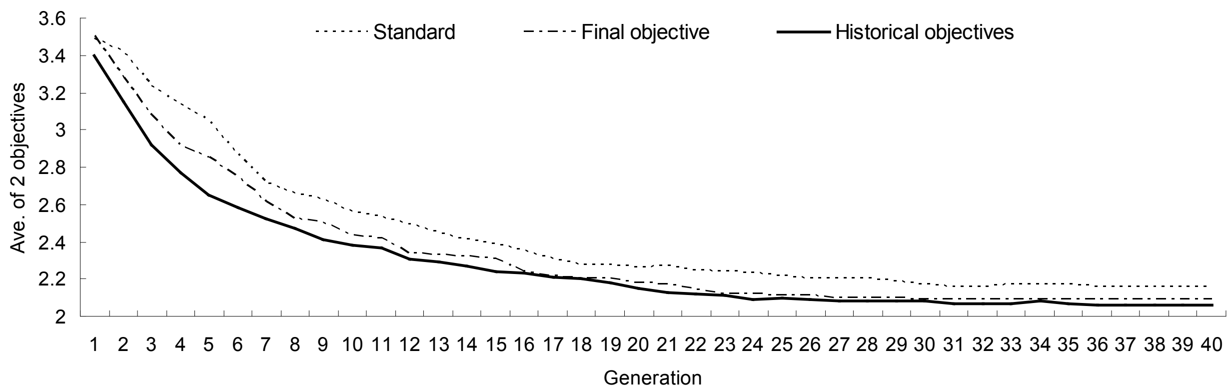

Figure 9: Average of 2 objectives in search process (average of 100 trials the small case)

Figure 7 and Figure 8 show, respectively, the Uniform-Cover Distance (UCD) and the averaged UCD of the solution sets of the three GAs at each generation. Figure 9 shows the intermediate changes of the 2 objectives' values (the average of the congestion length inside/outside the CBD) at each generation. These figures illustrate the advantages of the proposed method (the historical-objectives GA). For example, from Figure 7 we can see that after 12 generations, the UCDs of the populations in the proposed method were always lower than those of other GAs. Since these results are average of 100 trials, we believe they reach statistical significance.

4.2 The Large Case: Four Types of Actions on Sixteen Locations

We compared the three GAs on finding optimal scenario for the large case, which has 16 locations to impose the four actions. In contrast to the small case, the number of scenarios is too large and running all scenarios to obtain the Pareto-optimal solutions is nearly impossible so that we only computed their approximation. To obtain the approximated Pareto-optimal solutions, we ran all GAs for many times and recorded the multi-objective values of their solutions from the whole individuals whose scenarios were simulated. We then extracted the Rank1 solutions and regarded them as the approximated Pareto-optimal solutions.

We ran the three GAs started from the same initial population that contained 160 individuals. The populations of the GAs were evolved with crossover rate 0.2 for 47 generations, and their sizes were kept constant at 160. We set the mutation rate to zero similarly to the small case.

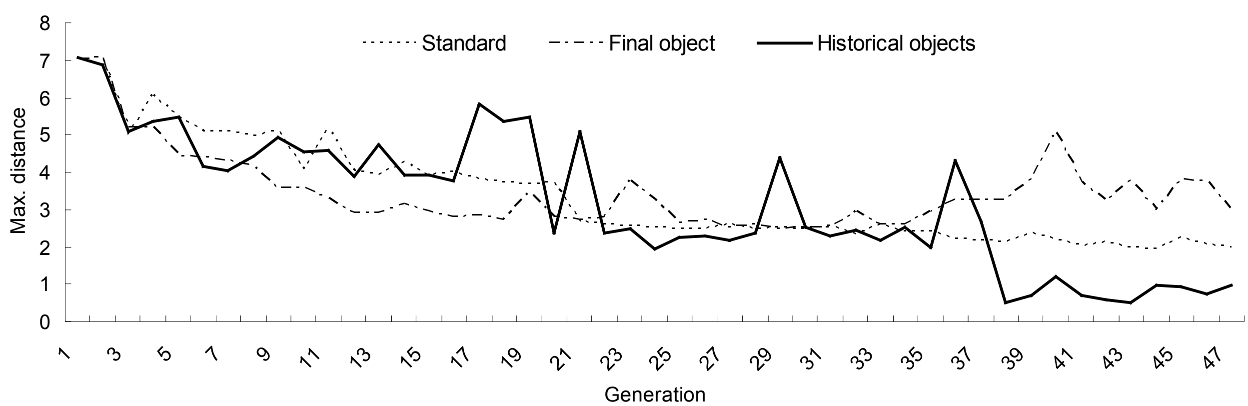

Figure 10: Maximum distance comparison (1 trial using the large case)

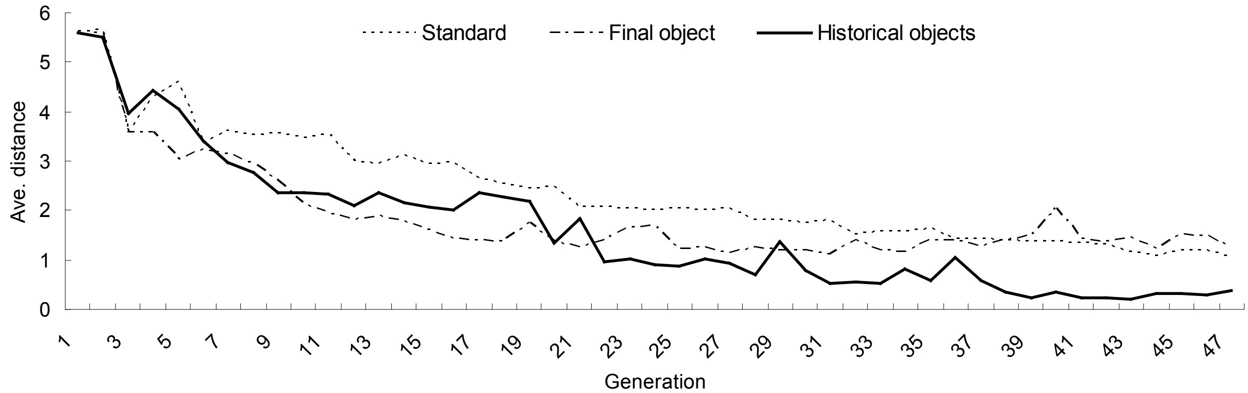

Figure 11: Average distance comparison (1 trial using the large case)

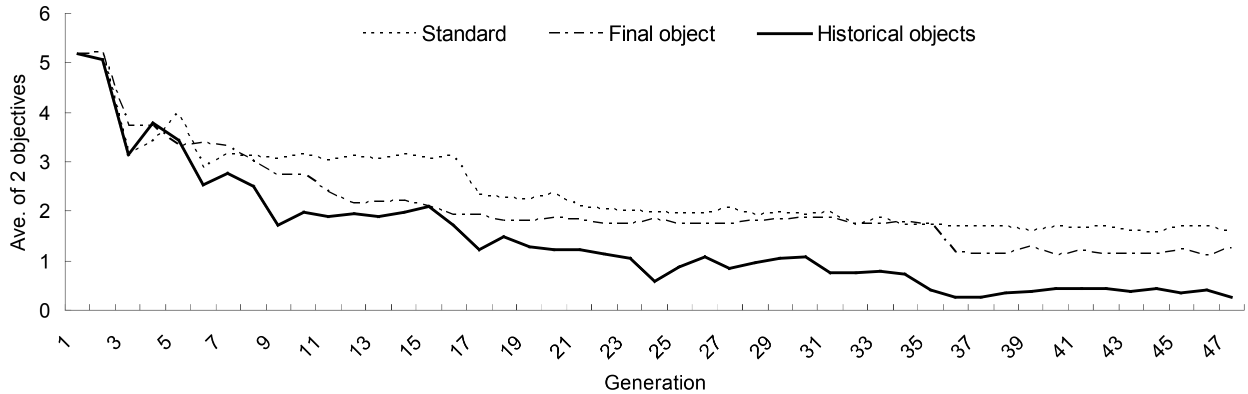

Figure 12: Average of 2 objectives in the search process (1 trial using the large case)

Figure 10 and Figure 11 show the comparison of the (approximate) UCDs and the average UCDs of the standard, final-objective, and historical-objectives GAs at each generation. From Figure 10, we can see that although the UCDs of the historical-objectives GA before the 37-th generations were higher than the other GAs, their values were significantly lower at the later generations. The averaged UCDs of the proposed GA were better at earlier generations, as evidenced by Figure 11. Figure 12 shows the intermediate changes of the 2 objectives' values (average congestion length inside/outside CBD) at each generation.

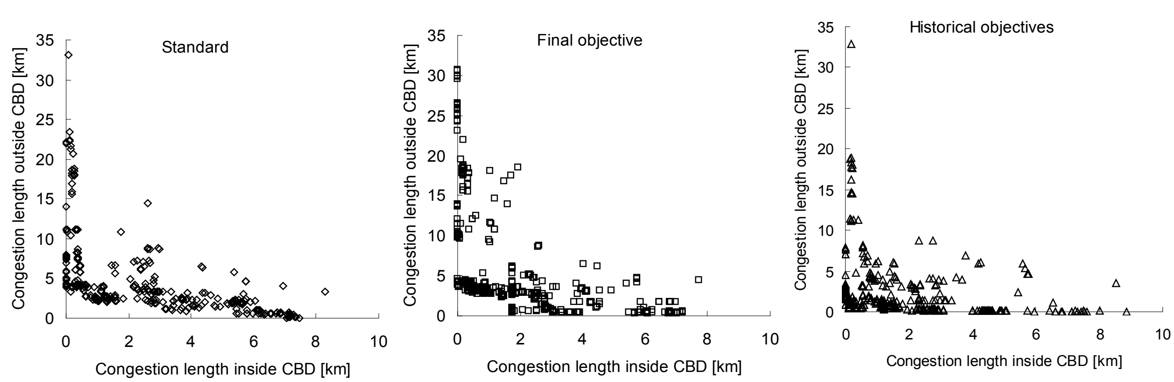

Figure 13: Rank1 solutions in 47 generations (the large case)

Figure 13 shows the Rank1 solutions of the three GAs from 1 to 47 generations. Figure 14 shows the Rank1 solutions of the three GAs in the 47-th (final) generation. Those figures show qualitative evidences of the superiority of the proposed GA to generate solution sets that are diverse and close to Pareto optimal quickly within several number of generations.

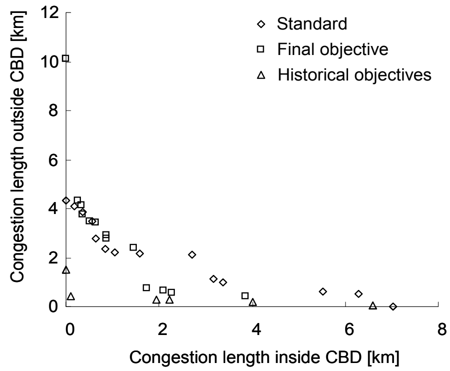

Figure 14: Rank1 solutions in the final generation (the large case)

5 CONCLUSIONS

We showed how to use the changes of objective values of simulations to assist a genetic algorithm measure diversity of simulation scenarios. By experiments on a small-scale simulation scenario where the Pareto-optimal scenarios were known, we confirmed that the proposed genetic algorithm could produce better solutions in less number of generations. Similar results were also observed on a large-scale simulation scenario.

There are several directions for future research. An agent-based traffic simulation can produce more aggregated states of its agents during the simulation. Finding out if such additional data can deliver better evolutionary algorithms is interesting. Other promising direction is to see if similar results hold for evolutionary algorithms on other types of simulations.

ACKNOWLEDGMENTS

A part of this research was supported by JST, CREST. We would like to thank Dr. Takashi Imamichi for his comments on previous work and arrangement of traffic regulation scenario for the simulation.

REFERENCES

Behrisch, M., L. Bieker, J. Erdmann, and D. Krajzewicz. 2011. “SUMO - Simulation of Urban Mobility:An Overview.” In *Proceedings of International Conference on Advances in System Simulation*, 55–60.

Brown, N., J. Gearhart, D. Jones, L. Nozick, N. Romero, and N. Xu. 2013. “Multi-objective optimization for bridge retrofit to adress earthquake hazards.” In *Proceedings of the 2013 Winter Simulation Conference,* edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 2475-2486, Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Delahaye, D., S. Oussedik, and P. Stephane. 2005. "Airspace congestion smoothing by multi-objective genetic algorithm." In *Proceedings of the 2005 ACM symposium on Applied computing*, 907-912.

Fonseca, C. M., and P. J. Fleming. 1993. "Multiobjective genetic algorithms." *IEE colloquium on Genetic Algorithms for Control Systems Engineering*, Digest No. 1993/130.

Goldbergs, D. 1989. *Genetic algorithms in search, optimization, and machine learning.* Addison-Wasley Publishing Company.

Gomes, G., A. May, and R. Horowitz. 2004. "Congested Freeway Microsimulation Model Using VISSIM." *Transportation Research Record: Journal of the Transportation Research Board* (1876): 71–81.

Gonzales, E. J., C. Chavis, Y. Li, and C. F. Daganzo. 2009. "Multimodal Transport Modelling for Nairobi, Kenya: Insights and Recommendations with an Evidence-Based Model". Technical Report UCB-ITS-VWP-2009-5, UC Berkeley Center for Future Urban Transport, A Volvo Center of Excellence.

Horn, J., N. Nafpliotis and D. E. Goldberg. 1994. "A niched Pareto genetic a lgorithm for multiobjective optimization." In *Proceedings of the First IEEE Conference on Evolutionary Computation*, 82-87.

Ikeda, K., H. Suzuki, S. Markon, and H. Kita. 2007. "Traffic-sensitive controllers for multi-car elevators; design, multi-objective optimization and analysis." In *SICE, 2007 Annual Conference*, 2655-2662

Imamichi, T., and R. Raymond. 2013. "Simple and Fast Trip Generation for Large Scale Traffic Simulation". In *Proceedings of the 2013 Winter Simulation Conference,* edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 2202-2031.Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Konak, A., D. W. Coit, and A. E. Smith. 2006. "Multiobjective optimization using genetic algorithms: a tutorial." *Reliability Engineering and System Safety*, 91, 992-1007.

Osogami, T., T. Imamichi, H. Mizuta, T. Morimura, R. Raymond, T. Suzumura, R. Takahashi, and T. Ide. 2012. "IBM Mega Traffic Simulator." IBM Research Report, RT0947.

Osogami, T., T. Imamichi, H. Mizuta, T. Suzumura, and T. Ide. 2013. "Toward simulating entire cities with behavioral models of traffic." *IBM Journal of Research and Development*, Vol.57, No.5.

Schmöcker, J., S. Ahuja, and M. Bell. 2008. "Multi-objective signal control of urban junctions – Framework and a London case study." *Transportation Research Part C: Emerging Technologies*, Vol 16, Issue 4, 454-470.

Suzumura, T., S. Kato, T. Imamichi, M. Takeuchi, H. Kanezashi, T. Ide, and T. Onodera. 2012. "X10-based massive parallel large-scale traffic flow simulation." In *Proceedings of the 2012 ACM SIGPLAN X10 Workshop*.

AUTHOR BIOGRAPHIES

HIDEMASA MUTA is researcher at IBM Research - Tokyo. He received his Ph.D. degree in engineering from The University of Electro-Communications, Tokyo, Japan. He has worked on various projects in IBM for 23 years. His research interests include multi-objective optimization, agent-based traffic simulation, HPC, and computer vision. His email address is hmuta@jp.ibm.com.

RUDY RAYMOND is an analytics and optimization researcher at IBM Research - Tokyo. He studied in Kyoto University and obtained the Bachelor of Engineering, Master, and Doctor of Informatics degrees in 2001, 2003, and 2006, respectively. In IBM Research, he has been working with various clients in telecommunication and manufacturing industries. His email address is raymond@jp.ibm.com.

SATOSHI HARA is an analytics and optimization researcher at IBM Research - Tokyo. He received his Ph.D. in Electrical Engineering from Osaka University in 2013, and a M.Eng. degree in Computer Science from Tokyo Institute of Technology in 2010, and a B.Eng. degree in Applied Physics from Waseda University in 2008. His research interests are machine learning algorithms and their applications in industries. His email address is satohara@jp.ibm.com.

TETSURO MORIMURA is Researcher at IBM Research - Tokyo. He received the M.E. and Ph.D. degrees in engineering from Nara Institute of Science and Technology, Japan. His research interests include reinforcement learning, machine learning, and its applications to neuroscience and business intelligence. His email address is tetsuro@jp.ibm.com.