Hedging the Risk of Delays in Multimodal Journey Planning

Adi Botea

■ *Traditional multimodal journey planners are deterministic. However, uncertainty in a transportation network can often lead to missed connections, causing delays and hurting the reliability level of the service. This article provides an overview of what is possibly the first multimodal journey advising system that is capable of reasoning under uncertainty and that provides more reliable journey plans.*

Multimodal travel, which combines several transport modes within a trip, can help fight congestion, carbon emissions, and pollution, important problems that affect many cities globally. These could become even more pressing in the future, as the global urban population will increase.

Multimodal trips exhibit challenges that, depending on how effectively they are addressed, can affect the attractiveness of multimodal travel. One example is a potential lack of reliability in some multimodal transportation networks. Intuitively, a more reliable transport network allows reaching the destination in time. Connection points, which are inherent in multimodal trips, make a journey more prone to delays because of potential missed connections.

Traditional multimodal journey planners compute deterministic, sequential plans, implicitly assuming that the input data is accurate. However, in real life, a multimodal transport network can feature many types of uncertainty. The exact arrival and departure times of scheduled vehicles, such as buses or trains, can differ from the published schedules. When using a city's shared-bike network, if a bike is not available straight away, hiring a bicycle can involve a variable waiting time. So can returning the bicycle into a free slot or parking a car. Even small variations of such uncertain timings can result in missed connections. Replanning can sometimes help, but more generally, early bad decisions can lead to states where replanning has little value if no good alternatives exist after missing a connection.

At IBM, staff from the Dublin Research Lab, the Zurich Research Lab, and other locations have developed a multimodal journey advisor capable of reasoning about uncertainty (Botea, Nikolova, and Berlingerio 2013; Nonner 2012) to provide more reliable journey plans (Botea and Braghin 2015). The topic was selected as a differentiator from existing multimodal journey planners.

 ISSN 0738-4602

As known, in AI research, shifting from deterministic to uncertainty-aware planning is a substantial change. These require different approaches, such as A* search in the former and AO*-based search in the latter. Typically, nondeterministic planning is computationally more difficult. The challenge in our domain is to solve real problems in real time. "Real problems" refer to modeling the domain in a detailed fashion and encoding the transport network of a (large) real city. Finding plans within a fraction of a second in most cases arguably passes as real-time performance.

A textbook AO* does not scale in such a domain. Building a fast, scalable journey planner based on AO* requires a substantial research effort to construct a detailed and yet effective domain modeling as an AND/OR search problem and to develop efficient speedup enhancements. The modeling and part of the enhancements are reported in a conference paper (Botea, Nikolova, and Berlingerio 2013). Others are planned for a future report. Below I list ideas that worked well, many of which resemble in high-level terms existing strategies from the AI literature.

The heuristic to guide AO* uses two lookup tables with admissible estimates (that is, not overestimating the true value), one for the travel time and one for the number of legs in the journey. When several buses could possibly come in any order at a given stop, a form of partial-order reduction allows pruning equivalent combinations. When several bus stops are within walking distance from a current location, generating a walking action to each of them might be unnecessary (for example, imagine that all stops are served by the same route). The system leverages rules for walking-action pruning and proves their correctness in the presence of uncertainty. Furthermore, users can specify maximum acceptable amounts for the walking time, the cycling time, and the number of legs in a journey. States are pruned when the amount spent so far, plus a precomputed admissible estimate of the amount needed from here on, exceed the maximum acceptable amount. The system performs state-dominance pruning. In nondeterministic planning, the correct application of dominance pruning depends on the types of the nondeterministic branches on the path to each of the two states considered.

The quality of the input data is important. Static bus timetables are more broadly available, but their accuracy suffers when actual arrival and departure times do not respect the timetable. Actual arrival and departure times at stops along a route, extracted for instance from GPS traces, are useful, but their availability is not very common yet. Historical GPS data can be used to build uncertainty models for arrival and departure times. Real-time updates are useful, for instance, to check whether trips in progress are invalidated by recent changes in the network. This allows notifying the affected users and replanning.

The system can read public transport updates, such as fresh arrival and departure times, from various sources, including the IBM IIT product. For bike, bike parking, and car parking availability, efforts have focused on prediction models based on historical and current data (Chen et al. 2013).

Our asset (Botea et al. 2016) has been featured in multiple client demonstrations, sometimes with a specific focus stemming from client needs. In one demonstration, the scope was extended from a single city to a regional, multicity area. Other extensions integrate car pooling (Berlingerio et al. 2015) and private cars.

Feedback received included the important question whether, in a realistic scenario, the performance of nondeterministic plans would differ significantly from standard plans. Part of our testing has focused on differences in terms of arrival times (Botea and Braghin 2015). A larger-scale evaluation is under way.[1] Actual journey plan requests and dynamic updates on transport network data from Rome, Italy, are passed in real time to our system, in parallel with the existing production system.[2] This would allow comparison of nondeterministic and deterministic multimodal journey planning in real-life conditions.

Other future directions include integrating electrical vehicles, and further research in AI speedup techniques for optimal multimodal journey planning under uncertainty.

Notes

1. Work partially supported by the EC under the FET-Open Project n. FP7-ICT-609042, PETRA.

2. www.agenziamobilita.roma.it/en/journey-planner.html

References

Berlingerio, M.; Bicer, V.; Botea, A.; Braghin, S.; Lopes, N.; Guidotti, R.; and Pratesi, F. 2015. Mobility Mining for Journey Planning in Rome. In *Proceedings of the 2015 International Workshop on Meta-Learning and Algorithm Selection Colocated with European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2015* (ECMLPKDD 2015), CEUR Workshop Proceedings 1455, 222–226. Aachen, Germany: RWTH Aachen University

Botea, A., and Braghin, S. 2015. Contingent versus Deterministic Plans in Multi-Modal Journey Planning. In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling,* ICAPS-15, 268–272. Palo Alto, CA: AAAI Press.

Botea, A.; Berlingerio, M.; Bouillet, E.; Braghin, S.; Calabrese, F.; Chen, B.; Gkoufas, Y.; Laummans, M.; Nair, R.; and Nonner, T. 2016. Docit: An Integrated System for Risk-Averse Multi-Modal Journey Advising. In *Smart Cities and Homes: Key Enabling Technologies,* ed. Mohammad S. Obaidat and Petros Nicopolitidis. San Francisco: Morgan Kaufmann Publishers.

Botea, A.; Nikolova, E.; and Berlingerio, M. 2013. MultiModal Journey Planning in the Presence of Uncertainty. In *Proceedings of the Twenty-Third International Conference on Automated Planning and Scheduling,* ICAPS-13. Palo Alto, CA: AAAI Press.

Chen, B.; Pinelli, F.; Sinn, M.; Botea, A.; and Calabrese,F. 2013. Uncertainty in Urban Mobility: Predicting Waiting Times for Shared Bicycles and Parking Lots. In *Proceedings of the 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013),* 53–58. Piscataway, NJ: Institute for Electrical and Electronics Engineers. dx.doi.org/10.1109/ITSC.2013.6728210

Nonner, T. 2012. Polynomial-Time Approximation Schemes for Shortest Path with Alternatives. In *Algorithms — ESA 2012: The 20th Annual European Symposium,* Lecture Notes in Computer Science 7501, 755–765. Berlin: Springer. dx.doi.org/10.1007/978-3-642-33090-2_65

Adi Botea is a researcher in the IBM Ireland Research Lab in Dublin, Ireland.