Research Statement/Dissertation Abstract

Supriyo GHOSH
School of Information Systems, Singapore Management University
Tel: (65) 83763144; Email: supriyog.2013@phdis.smu.edu.sg
March 8, 2015

Background

Rapid "urbanization" - more than 50% of worldś population now resides in cities - coupled with the natural lack of coordination in usage of common resources (ex: shared bikes, ambulances, taxis, attractions in tourism and leisure industry, rescue team in disaster response, sensors, meeting rooms) has worsen a wide variety of quality of life metrics such as satisfaction in issuing shared cars, waiting times in queues, response time for emergency needs, number of traffic accidents etc., in cities of today. Thus, I am broadly interested in solving large-scale multi-agent planning problems to provide decision support in urban environments, by optimizing the quality of life metrics, using well-known techniques from Operation Research and Artificial Intelligence machinery. More specifically, I am intended to dynamically match supply of resources to the demand (that is both stochastic and dynamic) for resources in urban environments. According to behavior of the problem, we can classify this domain into two sections (a) Dynamic Matching in urban systems with cooperative agents (ex: shared transportation systems (Ghosh et al. 2015a; Shu et al. 2013), emergency medical services (Saisubramanian, Varakantham, and Chuin 2015; Yue, Marla, and Krishnan 2012), personal rapid transports (Lees-Miller, Hammersley, and Wilson 2010)), where all the entities belong to a centralized authority and their goal is to optimize a single global objective (b) Dynamic Matching in urban systems with competitive agents (ex: taxi fleet optimization (Seuken, Cavallo, and Parkes 2008; Varakantham et al. 2012), customer route guidance in theme park, stadium, landmarks (Varakantham et al. 2015)), where individual agents are mainly interested to optimize their own objective.

At a high level, cooperative problems can be solved using optimization techniques such as Linear Programming (LP) or Mixed Integer programming (MIP), while in competitive settings, we need intelligent techniques from game-theory and behavioral economics.

Dynamic Matching in Cooperative Settings

A perfect example of large-scale multiagent planning problem in cooperative settings is shared transportation systems (ex: bike sharing, car sharing). vehicle sharing systems are widely adopted in major cities of the world due to the concerns of extensive usage of private vehicles that led to increased traffic congestion, carbon emissions, and usage of non- renewable resources. In vehicle-sharing systems, base stations (ex: docking stations for bikes) are strategically placed throughout a city and each of the base stations contain a pre-determined number of vehicles at the beginning of each day. Due to the stochastic and individualistic movement of customers, there is typically either congestion (more than required) or starvation (fewer than required) of vehicles at certain base stations. For example, in the worse case, there are about 750 cases of empty stations and 330 cases of full stations per day, in a major bikesharing company (CapitalBikeshare) in USA. We propose to dynamically match the customer demands to number of bikes in source station by redeploying idle vehicles using carriers so as to minimize lost demand or alternatively maximize revenue for the vehicle sharing company. Minimizing the lost demand can result in a significant cost due to carrier vehicles. Thus, we contribute an mixed integer linear programming (MILP) formulation to jointly address the redeployment (of vehicles) and routing (of carriers) problems(Ghosh et al. 2015a; 2015b).

Scalability: Due to the complex structure of the optimization model, it poorly scales with number of stations and thus a black-box solver like CPLEX cannot solve more than 20 stations problem. Therefore, we provide two approaches that rely on decomposability and abstraction of problem domains to reduce the computation time significantly. We have employed dual decomposition approach that computes strategies for the decomposed parts and combine the strategies to obtain a joint strategy. The joint problem of minimizing lost demand decomposes into redeploying of bikes and routing of carriers. To further improve the scalability, we propose an abstraction mechanism, where base stations that are located nearby, are clustered into one abstract station. Lastly, we retrieve the base station level redeployment and routing strategy from the abstract solutions.

Key Results: We validated the utility of our approach on real-world dataset of two bikesharing companies (i) CapitalBikeshare (ii) Hubway. Few of the key results obtained through the use of intelligent systems for dynamic matching of demand and supply of resources in STS are as follows:

- Revenue increased by 3% while lost demand reduced by 20% on CapitalBikeshare data.

- Revenue increased by 4% while lost demand reduced by 40% on Hubway data.
- Duality gap was less than 0.5% for Dual Decomposition.
- Outperforms current practice even with small variation in mean demand

Dynamic Matching under Uncertainty

Robustness is a key aspect in any real-life multi-agent planning problem because of the uncertainty in the problem domain. In the case of bike-sharing the uncertainty arises because the exact probability distribution of the customer demands is unknown. Instead, we can characterize the uncertainty sets for the demands based on historical data. Thus, we are interested to improve the worst case scenario (maximize the minimum reward), while respecting the bounds on the demands. We propose a fictitious play based approach between two players (decision maker and nature), each solving a mixed-integer programming model(Ghosh et al. 2015c). The decision maker selects the routes of the carriers and his initial intention of redeployment (i.e., the number of bikes picked up and dropped off at the stations). Then, nature generates a demand scenario that results in the lowest bike usage given that the trucks are routed as proposed. The demand scenarios generated by nature at each iteration are added incrementally to the decision makers problem to ensure the convergence of the algorithm. Given the growing complexity of the decision makers model, Lagrangian Relaxation is applied to solve this problem efficiently.

Dynamic Matching in Distributed Environment

Distributed constraint optimization (DCOP) is an important and widely adopted framework for coordinated multiagent decision making. We address a practically useful variant of DCOP, called resource-constrained DCOP (RC-DCOP), which takes into account agents consumption of shared limited resources. In RCDCOP, we need to match the supply/capacity of resources to the agents requirement/demand for resources in a distributed fashion. RCDCOP has been utilized in applications such as distributed management of smart grids(Kumar, Faltings, and Petcu 2009), distributed meeting scheduling (Bowring et al. 2009). In several real world applications, agents consume multiple shared resources with limited capacity. For e.g., in distributed meeting scheduling, agents schedule is constrained by their travel budget; in sensor networks, sensors may have limited battery. The coordination problem is now to optimize the global objective, while also respecting the resource limit for each resource. We present a promising new class of algorithm for RCDCOPs by translating the underlying coordination problem to probabilistic inference(Ghosh, Kumar, and Varakantham 2015). Using inference techniques such as expectation-maximization (EM) and convex optimization machinery, we develop a novel convergent message-passing algorithm for RCDCOPs. However, addressing resource constraints within the EM framework proves challenging as EM for RC-DCOP does not admit closed form solutions. Therefore, we combine several tools from convex optimization machinery (such as dual optimization, block coordinate descent) and algebra (polynomial root finding) in a novel way to derive the EM algorithm for RC-DCOPs. EM is easily implementable using local message-passing among agents, and is highly scalable. Unlike traditional and practically best DCOP algorithm like Max-Sum (MS), EM is guaranteed to converge. Empirically, we show that EM provides significantly better quality than MS, has low failure rate even under tight resource constraints and proves highly competitive to an efficient centralized constraint solver.

Key Results: We evaluated the utility of our approach on multiple synthetic and standard benchmark (distributed graph coloring) problems. We compare our approaches with two state-of-the-art DCOP solvers (a) Best approximate DCOP solver called Max-Sum (b) An efficient centralized solver called toulbar2. Few of the key results obtained through the use of intelligent systems for matching of demand and supply of resources in a distributed environment are as follows:

- EM almost always provide better solution quality than MS. EM was able to achieve a near-optimal solution which was very close to toulbar2. Indeed, for harder instance with higher edge density (ex: density=0.9) EM provided better quality than toulbar2.
- EM has less than 10% failure rate for random graph coloring problems, while failure frequency of MS is almost 90%.
- While toulbar2 cannot optimally solve the hard problem instances within 1 hour, EM converges to a near-optimal solution within 180 seconds only.

Dynamic Matching in Competitive Settings

My future plan is to work on dynamic matching problem in the presence of competitive agents. A few related application domains in urban environments that can be modeled as dynamic matching problem are:

- Customer route guidance (matching of customers to attractions) in theme park, landmarks to maximize the customer satisfaction by minimizing the waiting time in queue.
- Taxi fleet optimization (matching of customer demands to resources/taxis) to reduce the waiting time for customers as well as individual taxis.

In recent past, a wide range of research papers have addressed these problems to find a Nash equilibrium (Varakantham et al. 2012; Ahmed, Varakantham, and Cheng 2012) or by incentivizing customers using a novel dynamic mechanism design approach (Seuken, Cavallo, and Parkes 2008). However, (Aumann 1974) has shown that a Nash equilibrium can be arbitrarily bad in terms of global/social objective as optimizing individual goal can lead to an unsatisfactory joint strategy. For example, taxi drivers are always intended to visit crowded place like airport or shopping malls, which can lead to a congestion of taxis in busy areas and starvation of taxis in remote areas. A promising direction is to find a Correlated equilibrium (CE), where we find a probability distribution over all the joint strategies that optimize the global objective as well as ensure that individual objective is optimized. Unfortunately, as the joint strategy space grows exponentially with the number of agents, solution approach for CE suffers with a poor scalability issue. Therefore, I am interested to develop techniques for large-scale planning problems by exploiting the following key characteristics of the problem

- **Anonymity in interaction:** Typically, in urban environments, interactions between agents are anonymous. That is to say, outcome (reward or transition) of an interaction is dependent on how many agents and not on which specific agents are interacting. Thus, we can reduce the strategy space by partitioning the strategy space into a group of equivalent strategies.
- **Exploiting homogeneity of agents:** While there are large numbers of agents in urban environments, there are typically only a few types, where all agents of a type have similar decision models. We can compute the same mixed strategy for all agents of a type in order to exploit homogeneity without resulting in congestion.
- **Generate strategies incrementally:** Specifically in problems where mixed strategies are the outcomes, we can incrementally add relevant pure strategies. Although this kind of column generation approach does not provide any quality guarantee, it proves to be very effective in real-life problems with large action space.

Conclusion

Given the ever increasing urbanization rate, lack of available resources for cities and the above mentioned real world challenges, dynamic matching of supply and demand is an extremely important and challenging problem area that requires rapid development of next generation models, solution concepts and techniques to develop efficient smart-cities of tomorrow. Therefore, I believe that dynamic matching problem is a rich and high impact application area for dynamic decision making under uncertainty. The inherently high interdisciplinary nature of this field excites me and motivates me to cut across multiple areas of Artificial Intelligence, Operation Research and Machine Learning.

References

Ahmed, A.; Varakantham, P.; and Cheng, S.-F. 2012. Uncertain congestion games with assorted human agent populations. *arXiv preprint arXiv:1210.4848*.

Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics* 1(1):67–96.

Bowring, E.; Yin, Z.; Zinkov, R.; and Tambe, M. 2009. Sensitivity analysis for distributed optimization with resource constraints. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 633–640. International Foundation for Autonomous Agents and Multiagent Systems.

Ghosh, S.; Varakantham, P.; Adulyasak, Y.; and Jaillet, P. 2015a. Dynamic redeployment to counter congestion or starvation in vehicle sharing systems. In *ICAPS*.

Ghosh, S.; Varakantham, P.; Adulyasak, Y.; and Jaillet, P. 2015b. Dynamic redeployment to reduce lost demand in vehicle sharing systems. Under submmission, European Journal of Operation Research.

Ghosh, S.; Varakantham, P.; Dominik Jena, S.; and Jaillet, P. 2015c. A two player game approach for robust redeployment in bike sharing systems. Under submmission, 4th Informs TSL Society Workshop.

Ghosh, S.; Kumar, A.; and Varakantham, P. 2015. Probabilistic inference based message-passing for resource constrained dcops. Under submmission, IJCAI.

Kumar, A.; Faltings, B.; and Petcu, A. 2009. Distributed constraint optimization with structured resource constraints. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 923–930. International Foundation for Autonomous Agents and Multiagent Systems.

Lees-Miller, J. D.; Hammersley, J. C.; and Wilson, R. E. 2010. Theoretical maximum capacity as benchmark for empty vehicle redistribution in personal rapid transit. *Transportation Research Record: Journal of the Transportation Research Board* 2146(1):76–83.

Saisubramanian, S.; Varakantham, P.; and Chuin, L. H. 2015. Risk based optimization for improving emergency medical systems. In *AAAI*.

Seuken, S.; Cavallo, R.; and Parkes, D. C. 2008. Partially synchronized dec-mdps in dynamic mechanism design. In *AAAI*, volume 8, 162–169.

Shu, J.; Chou, M. C.; Liu, Q.; Teo, C.-P.; and Wang, I.-L. 2013. Models for effective deployment and redistribution of bicycles within public bicycle-sharing systems. *Operations Research* 61(6):1346–1359.

Varakantham, P.; Cheng, S.-F.; Gordon, G. J.; and Ahmed, A. 2012. Decision support for agent populations in uncertain and congested environments. In *AAAI*.

Varakantham, P.; Mostafa, H.; Na, F.; and Chuin, L. H. 2015. Direct: A scalable approach for route guidance in selfish orienteering problems. In *AAMAS*.

Yue, Y.; Marla, L.; and Krishnan, R. 2012. An efficient simulation-based approach to ambulance fleet allocation and dynamic redeployment. In *AAAI*.