Complete description of the problem

The ultimate goal of this research thrust is the development of robust and adaptable transportation networks for a safer and more efficient transportation system. Our approaches are motivated by the pervasiveness of information technology and the growing network of cyberinfrastructure, resulting in large volumes of data (including some in real time) contributed by users and vehicles, which can be used to sense, respond and shape system behaviors in a variety of ways. The context is that of novel and evolving technologies beginning to be adopted for instrumented networked systems, such as vehicle-to-vehicle (v2v) sharing of locally sensed data; which are layered upon existing fixed sensors to result in a dynamic network of interconnected communication systems.

We present a two-pronged approach to designing robust and dynamically reconfigurable systems. Our first research direction is the robust design and planning of systems, to incorporate flexibility against future uncertainty. Our second and parallel research direction is the dynamic adaptation or reconfiguration of systems as uncertainty is realized in real-time; in order to minimize associated negative impacts.

We plan to use two case studies - (1) improving emergency response and disaster management, and (2) providing navigational assistance to users for finding safe and fast routes on the National Highway system. In year one, we focused on Case Study 1 due to quicker availability of data.

Data gathered

Case study1: Emergency Medical Systems

We gathered usage data from an emergency medical system (EMS) operating in a large, densely populated and representative city with a number of challenges. The dataset includes approximately ten thousand logged emergency requests over the course of one month. Each logged request contains the type and location of the call, the ambulance (if any) that was dispatched, and the various travel times (e.g., base to scene, scene to hospital, hospital back to base, etc.). We use this representative data to further sample call logs into the future and build our modeling and solution approach.

Case Study 2: Finding fast and safe routes on the National Highway System.

We conversed with multiple state transportation agencies regarding the availability of data of two kinds: (i) traffic data and (ii) crash data. We obtained data about all reported crashes in the year 2011-2012, as well as aggregated daily traffic count data from the Penn DOT. We had hoped to synthesize both these data sets to infer and learn conditions (for example, weather, traffic, behavior, etc.) that that are correlated to crashes, to predict and provide better navigation systems to users. However, because the traffic count data was not disaggregated by time of day, etc. we could not use that data in our models. We hope that data obtained through counts from video will, in time, provide us the disaggregate data needed for this purpose.

In the following sections, we will focus on the results approach, findings and results from Case Study 1.

Approach and Methodology

Leveraging the pervasiveness of information technology, we obtained call log information from an EMS provider, described previously. We use a *data-driven modeling approach* that embeds within it the information available from these call logs. Using the call logs and a learning model gleaned from the dispatchers at the EMS, we infer missing data from the call logs. We built a simulation testbed to within which we generate several sequences of calls following the same pattern as seen in the original data and embed them within a *simulation-optimization* approach.

The simulation-optimization approach we propose offers four main advantages compared to existing approaches. First, we can easily incorporate features such as time-dependent travel times and (street and call) congestion patterns. Second, we can also easily optimize for a variety of metrics, such as preparedness, survival probabilities and others. Third, our approach is scalable to problem instances much larger than the ones in the literature. Fourth, our simulation approach has the ability to capture customer behavior, such as inferring limits on customer patience and waiting time.

Our model can be used for pre-positioning ambulances based on expected call patterns, as well as to dynamically re-position ambulances as peaks occur in expected demands for the service.

Findings

We find, from our customer behavior model, that there are patterns in behavior and patience, that vary between rural and urban areas, and primarily depend on base-to-scene time of the EMS service.

Mathematically, we also show that the structure of the problem as modeled using our simulation optimization approach is an approximately sub-modular optimization problem. We thus solve it using a greedy algorithm. We also prove mathematically that the results obtained from our algorithm produce results that are close to optimal.

We test our approach and solutions to changes in data and show that they are robust to changes in the data, that is, if the demand is higher than expected, or if travel time is longer than expected.

Conclusions and results

Our results indicate that the performance of the EMS can be significantly improved through our modeling and solution approach. The number of requests that are not serviced is decreased by about 50%, improving responsiveness. Our approach also improved multiple performance metrics, including base-to-scene time, fraction of requests with service time less than 30 min, etc. We also show that using our dynamic re-positioning model, which can be solved in seconds, allows for further improved results beyond the static pre-positioning model.

Our modeling and optimization approach is applicable in other settings such as disaster response, humanitarian service logistics and facility location. While the static allocation problem is well studied in such settings, significant gains are possible via efficient and effective dynamic redeployment.

Recommendations developed

We stress the benefits of using a data-driven approach that remains faithful to the occurrences seen in the real world. Our data-driven approach offers the flexibility of using a similar approach to both planning (static allocation) and real-time operations (redeployment), facilitating real-world implementation. Additionally, our approach offers flexibility in balancing between competing performance metrics. However, such flexibility also implies the need to tune the penalty function via simulation for different operational goals.

One potential limitation is the need for a reliable simulator (and data sampler), which typically requires a substantial amount of historical data as well as model engineering. However, simulation is typically used for evaluation in these types of settings, implying the ready availability of such simulators in practice to use for optimization.

Because our setting lies within a regime where worst case problem instances do not arise in practice (i.e., our objective is “close" to submodular), we can employ greedy algorithms to arrive at an efficient algorithm for real-time redeployment. Nonetheless, it may be beneficial to develop more sophisticated algorithms that can offer further improvements, especially for objectives that are far from submodular.

We see that different customer bases behave differently based on the service. It may be a worthwhile endeavor to explore if different dispatch algorithms can improve the reach of the EMS service, thus resulting in different service strategies for different communities.