

The Application of Market-based Multi-Robot Task Allocation to Ambulance Dispatch

Eric Schneider¹, Marcus Poulton², Archie Drake¹, Leanne Smith³,
George Roussos², Simon Parsons¹, and Elizabeth I Sklar¹

¹ Dept of Informatics, King’s College London, WC2R 2LS, UK

{eric.schneider,archie.drake,simon.parsons,elizabeth.sklar}@kcl.ac.uk

² Dept of Computer Science & Information Systems, Birkbeck College, University of
London, WC1E 7HX, UK marcus.poulton@gmail.com,g.roussos@bbk.ac.uk

³ London Ambulance Service, UK

leanne.smith@lond-amb.nhs.uk

Abstract. Multi-Robot Task Allocation (MRTA) is the problem of distributing a set of tasks to a team of robots with the objective of optimising some criteria, such as minimising the amount of time or energy spent to complete all the tasks or maximising the efficiency of the team’s joint activity. The exploration of MRTA methods is typically restricted to laboratory and field experimentation. There are few existing real-world models in which teams of autonomous mobile robots are deployed “in the wild”, e.g., in industrial settings. In the work presented here, a market-based MRTA approach is applied to the problem of ambulance dispatch, where ambulances are allocated in respond to patients’ calls for help. Ambulances and robots are limited (and perhaps scarce), specialised, mobile resources; incidents and tasks represent time-sensitive, specific, potentially unlimited, precisely-located demands for the services which the resources provide. Historical data from the London Ambulance Service describing a set of more than 1 million (anonymised) incidents are used as the basis for evaluating the predicted performance of the market-based approach versus the current, largely manual, method of allocating ambulances to incidents. Experimental results show statistically significant improvement in response times when using the market-based approach.

Keywords: Multi-robot team; task allocation; auction mechanism

1 Introduction

The well-studied problem of *multi-robot routing* involves assigning a team of robots to travel to a set of locations such that collisions are avoided and some performance metrics are optimised, such as minimising travel time or distance. The real-world challenge of emergency vehicle dispatch bears a number of key similarities to multi-robot routing. *Emergency medical services (EMS)* agencies receive calls (*incidents*) which frequently result in the dispatch of one or more vehicles (*responses*) to the location of the incident. The determination of which

ambulance should respond to which incident is highly complex, involving consideration of traffic conditions, knowledge of road infrastructure, patient situation and requirement for any specialised healthcare equipment or specially trained personnel. Sub-optimal travel time or improperly equipped emergency response crews can contribute to loss of life, and extraneous travel distance can result in unnecessary costs for what are typically financially challenged public agencies.

In the study described here, we apply our earlier work on *market-based mechanisms* for *multi-robot task allocation (MRTA)* to the real-world problem of ambulance dispatch. We present three novel contributions:

1. The use of historical data from the London Ambulance Service (LAS) that serves as the set of incidents considered—instead of previously engineered or randomly chosen task and robot starting locations, as is the case in most MRTA literature;
2. The use of the next-generation *routing engine* developed for the LAS — instead of classic routing engines such as A*, which is used by many MRTA approaches, or Google Maps, which is used by many recent studies involving traffic management; and
3. Experimental evaluation comparing predicted response times when using market-based mechanisms versus the current, largely manual, method of task (incident-to-ambulance) allocation—showing statistically significant differences and improved response times when the market-based mechanisms are utilised.

Although this proof-of-concept study is conducted in collaboration with LAS, and so focuses on London (UK), and on emergency medical services response, we believe that more general conclusions can be drawn for ambulance services in other parts of the world, as well as other types of emergency response services, such as police or fire.

The structure of this paper is as follows. Section 2 covers technical background and prior work on aspects of multi-robot task allocation. Section 3 describes the specific application domain addressed in our work: the problem of ambulance dispatch facing the LAS. Section 4 outlines a set of experiments we conducted using data provided by LAS, where we applied market-based multi-robot routing techniques to the ambulance dispatch problem. Section 5 presents and discusses the results of these experiments. Then, Section 6 highlights related work on ambulance dispatch. Finally, in Section 7, we conclude with a summary of current work, next steps and directions for future work.

2 Multi-Robot Task Allocation

In a multi-robot routing problem, a team of mobile robots must collectively visit some number of *task locations* where they will perform activities such as site inspection or object pick-up and/or delivery. Solutions to this class of problem entail distributing tasks to robots and planning routes, or *paths*, from robots'

current positions to their assigned task locations in order to optimise some criteria such as travel distance or time. A multi-robot routing problem is similar to a multiple depot, multiple Travelling Salesperson Problem (mTSP) [3] or Vehicle Routing Problem (VRP) [21], where vehicles need not return to their depots [19]. The primary challenge of solving a multi-robot routing problem is *multi-robot task allocation (MRTA)*: deciding which tasks should be assigned to which robots so that the overall execution of a mission is, by some measure, efficient. While there are several kinds of approaches to solving task allocation problems, we focus on *market-based* methods of task allocation, and *auctions* in particular, because they can be flexible, distributed and, in some cases, scalable.

Market-based approaches to task allocation frame the assignment problem as a multi-agent systems (MAS) problem. Rather than having a centralised planner be responsible for computing the costs or utilities of potential allocations, a market-based approach to MRTA relies on the fact that robot team members are each capable of planning subsets or sub-problems of the mission (i.e., planning to execute individual tasks or groups of tasks) and can express the costs or utilities of these plans in a way that is simple and efficient to communicate. Task allocation is governed by a *mechanism*, a set of rules that dictate how tasks should be assigned, and a protocol for communicating the availability of tasks to robots and the values robots have for them. A mechanism enables a virtual marketplace in which tasks can be distributed to robots or exchanged among them. A common kind of market-based mechanism for MRTA is an *auction*, which compares bids for resources from interested parties and awards them to the highest (or lowest) bidder according to the particular rules of a mechanism. It can be expensive to compute an allocation that is optimal for some performance objective, so most auction mechanisms strive for *approximately* optimal allocations. Designers of auction mechanisms must make trade-offs between the costs of computing an allocation and the performance of the execution of a mission that results from the allocation.

Our previous work in the MRTA domain led to the development of our MRTEAM framework [30,29], which was designed to evaluate a range of task allocation mechanisms in simulation and on physical robots. While other research in multi-robot routing has concentrated largely on discovering optimal assignment mechanisms for a single type of environment, our work using MRTEAM has focussed on evaluating a range of performance metrics in a variety of complex task environments [12,20] and analysing both task assignment *and* task execution—which makes this framework particularly relevant for application to a real-world domain where we are especially concerned with measuring *response* or *travel times*.

In MRTEAM, a *map* specifies the extent of a geographical space and the arrangements of free space and obstacles within it. A *team* is a set of n robots $R = \{r_0, \dots, r_{n-1}\}$. A *starting configuration*, S , specifies the location on a map of each robot in the team at the beginning of a *mission*. A *scenario* is a set of m tasks $T = \{t_0, \dots, t_{m-1}\}$ situated on the map. Each task $t \in T$ has the following properties: $t.pos$, a fixed position on the map; $t.arr$, the arrival time

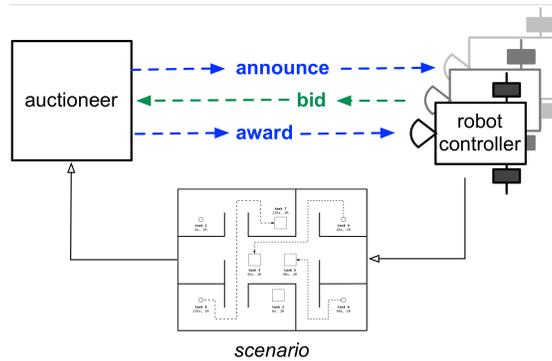


Fig. 1: Interaction of an *auctioneer* agent with robots in MRTEAM. *Robot controller* agents compute bids and are responsible for autonomous navigation to task locations after allocation. In the ambulance domain, the auctioneer is a proxy for a dispatcher and robot controllers are proxies for ambulance vehicles.

of the task; and $t.req$, the number of robots required to complete the task. A *mission* comprises the map, a scenario, and a robot team with a starting configuration: $M = \{map, T, R, S\}$. An auction mechanism allocates tasks to robots over a number of *rounds*. In an auction round, a coordinating *auctioneer* agent announces tasks to the team, team members compute and submit *bids* to the auctioneer, and the auctioneer awards one or more tasks to team members according to the rules of the mechanism (Figure 1).

A bid for a task is computed by a robot as:

$$bid = w_0 f_0 + w_1 f_1 + \dots + w_{\varphi-1} f_{\varphi-1}$$

where f_i is a quantitative factor, w_i is a weight associated with that factor, and φ is the number of factors to consider in the bid. For example, f_0 could be the estimated travel distance from the robot’s starting location at the time of placing the bid to the location of the task it is bidding on, and f_1 might be the priority of the task. If we deem distance more important than task priority, then $w_0 > w_1$.

The metrics we use to evaluate performance in MRTEAM measure the *distance travelled* by robots as well as various time-based performance measures such as: *deliberation time* (the time taken to compute a set of task assignments); *execution time* (the time taken to execute the assigned tasks); *movement time* (the time robots spend actually moving towards task locations); *idle time* (the time a robot sits idly waiting for other robots to complete their tasks after it has completed its last task); and *delay time* (the time each robot spends waiting for other robots to pass safely by in order to avoid collisions). As described in Section 4, *movement time* is the most relevant measure for the experiments described here, though *idle time* will be a key metric to consider in future work.

The next section explains how the problems faced in ambulance dispatch are related to those explored in the MRTEAM project.

3 Ambulance dispatch in London

Greater London (UK) has a population of approximately 8.2 million people, according to the 2011 national census [25], and its medical emergencies are handled by the London Ambulance Service, which is governed by a UK state agency, the National Health Service (NHS). In order to manage the needs and resources of such a large city, the NHS divides London into 33 *Clinical Commissioning Groups* or *CCGs*. These CCGs are grouped into five *sectors* (Figure 2). Ambulances are allocated to "home" locations within each CCG, indicating where crews report at the start and end of a shift. In performing assignment of an ambulance to an incident, there is an attempt to keep ambulances within their "home" CCGs for a number of reasons, such as reducing travel distances (and associated petrol costs) and travel times (not only because distances are shorter but also because crews are more familiar with the roadways in their home CCG).

LAS receives emergency medical service calls 24 hours a day, 365 days of the year. For example, in 2016, LAS handled approximately 5,000 emergency calls on a daily basis. If a call requires a responder vehicle, then the call is logged as an *incident*. The main responder vehicles are *accident and emergency units* (AEUs), the large "truck" ambulances which are capable of transporting a patient to hospital, and the *fast response units* (FRUs), i.e., cars, which can get to the scene in a shorter amount of time. Every responder vehicle that is sent to an incident is defined as a separate *response*, so often there can be multiple responses for a given incident. It is not unusual for an FRU to be dispatched first, in order to give help as quickly as possible, followed by an AEU, to provide the means to convey a patient to hospital.

LAS receive many different types of call, and these are categorised based on severity, which is defined according to the nature of a patient's *chief complaint*. Calls are categorised from A (highest priority) to C, each with subcategories referring to the target response time. For Category A, the highest subcategory is **red1**, referring to a life threatening incident with a target response time of 8 minutes. The subcategories for Category C range from **green1**, with a response time of 20 minutes to **green4** with a target time of 4 hours. Depending on the nature of the call, the LAS definition of different time measures vary. *Response time* is always defined as the difference in time between the *clock start time* and the first responder arriving at scene of the incident, however the clock start time is measured differently for different patients. The clock start for the highest priority cases begins when the call is answered by the control room. For other calls, the clock start time begins at the earliest of: the first vehicle being dispatched; the type of incident being determined; or 240 seconds after the call is answered.

Performance is measured by the proportion of first-responder response times that fall within an incident category's maximum allowable time. Performance is measured for Greater London overall, but also for each CCG. In 2016, ap-

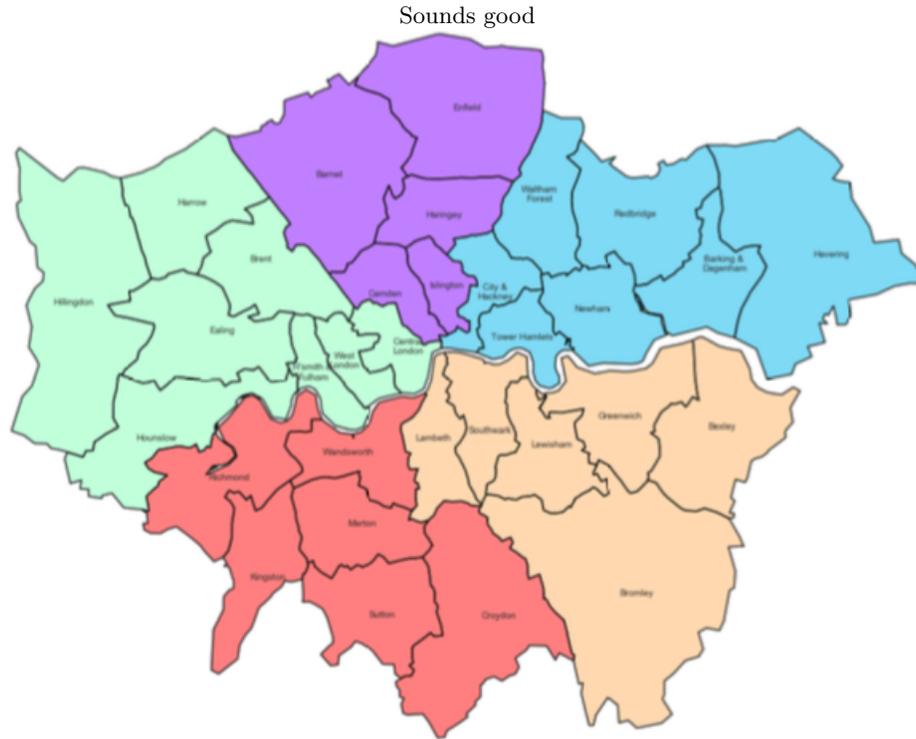


Fig. 2: London Ambulance Service sectors and NHS Clinical Commissioning Groups. The Service sectors are the large coloured regions: North Central (purple), North East (blue), North West (green), South East (ochre) and South West (red). Clinical Commissioning Groups are the areas within the service sectors.

proximately 65% of Category A incidents received responses within the target 8-minute time limit. Since 2017, the LAS have adopted new categories of severity and performance measurement through the Ambulance Response Program.⁴

The notion of computer-assisted dispatch (CAD) was first introduced in the London Ambulance Service in 1992 and quickly became a lesson in software engineering mishaps. The early version of CAD included two key components: “an automated vehicle locating system (AVLS) and mobile data terminals (MDTs) to support automatic communication with ambulances” [10]. Within hours after deployment, the AVLS lost track of vehicles’ whereabouts, so the CAD database became inaccurate. Ambulances were dispatched non-optimally: some calls received multiple ambulances; others received none. The CAD software started issuing error messages and overloaded the system. Ambulance crews stopped sending status reports via MDTs because the system was too slow. This catas-

⁴ <https://www.england.nhs.uk/urgent-emergency-care/arp/>

trophic failure led to deficient patient care, possible loss of life, and loss of employment for the LAS chief executive. Since then, the road has not been smooth: a software upgrade in 2006 led to systemic failure [22] and the initial introduction in 2011 of the current CAD system, *CommandPoint*,⁵ was delayed for technical reasons. However, since 2012, CommandPoint has successfully been providing dispatch support for LAS [14]. The LAS are currently in the process of re-evaluating and upgrading their CAD software modules to exploit new and emerging technologies, such as pervasive mobile computing, and sources of data such as real-time traffic and weather information [9].

The long term vision for the work described here is the integration of our auction-based multi-robot routing methodology into the computer-assisted dispatch system. The experiments described in the next section and results that follow will help us demonstrate to LAS the predicted advantages of our approach.

4 Experiments

This section describes the series of experiments that we conducted to compare the results when vehicles are allocated to incidents using our auction-based mechanism versus the manual allocation process currently employed, where human dispatchers in the LAS control room can consult the CAD system for recommendations but ultimately execute selections themselves.

4.1 Experimental Setup

The auction-based mechanism employed in our experiments was taken from the MRTEAM framework, described earlier. For simplicity, bids were derived using one quantitative factor: $f_0 = \textit{estimated travel time}$. This will provide us with a baseline for future work in which we can consider additional factors in the bidding. The experiments conducted here demonstrate that even using just the one factor, the auction-based methodology predicts significantly shorter response times (detailed in Section 5).

Our experimental evaluation was facilitated by an historical data set provided by the LAS, which records, for each incident that occurred in 2016, the location and call time of the incident, the locations of vehicles at the times they were dispatched to the incident, and the vehicles' travel times to the incident location, as well as other information about the incident, such as chief complaint and category. The data set contains 1.1 million incident records and 1.5 million response records. In order to keep the data anonymised, location coordinates are quantised to the nearest vertex on a 100m-precision grid.⁶ For the experiments described here, we only considered Category A incidents.

⁵ <https://www.northropgrummaninternational.com/capabilities/command-point/>

⁶ <https://www.ordnancesurvey.co.uk/resources/maps-and-geographic-resources/the-national-grid.html>

Because estimated travel time is taken as the basis for bidding, it is important for us to compute that carefully. We considered two methods for computing routes and estimating travel times between vehicle and incident locations: one makes use of a publicly available route planner (the Google Maps Directions API,⁷ referred to here as GMAPS), and the other makes use of a proprietary routing engine, called QUEST [27]. Thus we can compare three different response times: (i) the *historically observed* response time (taken directly from the LAS data set); (ii) the *QUEST-simulated* response time, taking the vehicle start and end locations from the LAS data set and using the QUEST routing engine to estimate travel time; and (iii) the *GMAPS-simulated* response time, again taking the vehicle start and end locations from the LAS data set, but using the Google Maps route planner to estimate travel time. For privacy reasons, we do not have access to the actual routes taken by emergency response vehicles, so using the simulated response times based on historical start and end locations gives us a fair basis for comparison between actual vehicle choices and simulated choices taken by our auction-based mechanism.

We designed a set of simulation experiments to compare two independent variables: (1) *vehicle selection* ("historical choice" (HIST) or "auction mechanism choice" (AUCT)) and (2) *routing engine* (QUEST or GMAPS). **Our hypothesis is that the auction mechanism choices will predict shorter response times than the historical choices, for either routing engine.** We evaluate this hypothesis in two steps. First, we produce a benchmark measure by comparing QUEST-simulated and GMAPS-simulated travel times with historically observed travel times, using the same (historic) start and end locations from the LAS data set for all three metrics. Second, we evaluate the efficacy of the auction mechanism by comparing simulated travel times for pairs of start and end locations: the historically recorded vehicle in the LAS data set (the benchmark) versus the vehicle chosen by the auction mechanism. Details of the two steps are provided below.

4.2 Benchmark Generation

To provide a benchmark for evaluation, we compared the historically observed response time for a chosen vehicle from the LAS data set with the simulated response times computed by each routing engine. In both cases, we used the start and end locations of the historically chosen vehicles to compute routes. This step also serves to demonstrate the advantage of using QUEST-simulated travel times, which are derived from historical road speeds of emergency service vehicles. A sample of 2000 Category A incidents was drawn uniformly randomly from the data set. We identified the first response vehicle assigned to each incident and queried each of QUEST and GMAPS for a route between the vehicle's location at the time the vehicle was dispatched and the incident's location, along with an estimated travel time. We then compared both estimates to the historical travel time observed in the sample. The results, shown in Figure 4 and discussed in the

⁷ <https://developers.google.com/maps/documentation/directions>

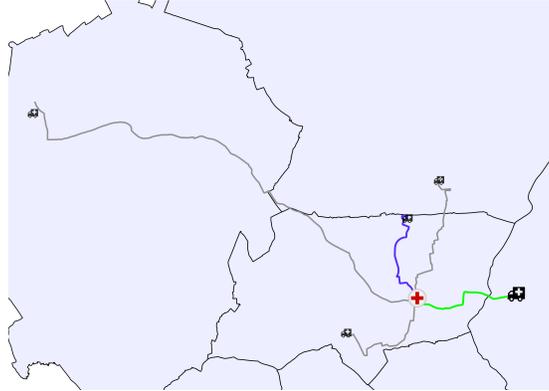


Fig. 3: Locations of idle vehicles and their routes to an example incident location (red cross). The route (369 seconds) of the historical first-responder vehicle dispatched to the incident is shown in blue. The route of a first-responder vehicle chosen via auction (duration: 272 seconds) is shown in green.

next section, show that the estimates computed by QUEST are more accurate than GMAPS with respect to actual travel times of emergency vehicles. Following this demonstration of the effectiveness of the QUEST engine we carried out the experiments that are the main contribution of this paper.

4.3 Adaptation of Auction Mechanism Framework

The MRTEAM framework was adapted for these experiments in the following ways. The *map*, represented in the QUEST routing engine, is based on the ITN Road Layer⁸ map produced by the UK’s Ordnance Survey mapping agency. Each “robot” agent in R represents an emergency vehicle (ambulance) and is capable of planning a route between two locations on the map and computing the distance and estimated time to travel along the route. A *scenario* comprises a single emergency incident task, described in detail below. The *mechanism* employed in all experiments is the *sequential single-item* auction, which has been shown to produce allocations that are close to optimal [18] while scaling better than e.g., combinatorial auctions [4]. A *bid* comprises a single bid factor, the estimated travel time between a vehicle’s location and an incident location. The *auctioneer* agent functions identically to its robot setting.

The data set did not provide the locations of *idle vehicles*—vehicles not *en route* to an incident or otherwise assigned (for example, for conveyance from an incident location to a hospital). However, the data set did provide the locations of vehicles both at the time they completed their assignments and at the time they were next dispatched following a completed assignment. Thus, the locations

⁸ <https://www.ordnancesurvey.co.uk/business-and-government/help-and-support/products/itn-layer.html>

of idle vehicles for a given time instant were interpolated along a route between these two locations.

For every incident in each experimental condition, idle vehicles were identified within a 20 km² neighbourhood around the incident and their locations at the call time of the incident were estimated as described above. The auctioneer agent announced the “task” (incident) to agents representing idle vehicles in their neighbourhood. These agents computed bids representing their estimated travel times to the incident location, and submitted them to the auctioneer. The auctioneer then aggregated the bids received and assigned the lowest-bidding vehicle to the incident. Figure 3 depicts an example in which idle vehicles have planned routes to an incident location (red cross). The route of the lowest-bidding vehicle is shown in green while that of the vehicle dispatched to the incident, historically, is shown in blue.

4.4 Experimental Conditions

Since evaluating a dispatch decision for every incident from the data set (> 1 million) was infeasible, we defined four experimental conditions that drew samples of incidents. In each condition, 100 Category A incidents were sampled uniformly randomly from a temporal and geographic range. For each sample incident, the travel time of the historically-assigned first-responding vehicle from its location at dispatch time to the incident location was compared to that of a (possibly different) vehicle chosen by the auction mechanism. Condition 1M-1C sampled incidents that occurred over one month in one arbitrarily selected Clinical Commissioning Group (CCG); 12M-1C sampled from 12 months (all of 2016) in the same CCG; 1M-nC sampled from one month and all CCGs ($n = 33$); and 12M-nC sampled from 12 months and all CCGs. Table 1 lists the four conditions under which experiments were conducted in order to evaluate the effectiveness of our approach.

	one month	12 months
one CCG	1M-1C	12M-1C
all CCGs	1M-nC	12M-nC

Table 1: Experimental conditions

4.5 Metrics

We computed two types of metrics, both of which are analysed in the next section. The first metric is *simulated response time*, discussed above, where a routing engine takes as input a start and end location for a vehicle and then estimates

the amount of time needed for the vehicle to travel from one location to the other. Shorter response times are better. This metric is the equivalent to *movement time* from the MRTEAM framework. The second metric is *vehicle choice*. During experiments, we record the identity of the vehicle chosen by the auction mechanism and then compare that to the historically observed vehicle choice. We count how many choices were made differently by the auction mechanism as opposed to the human-in-the-loop CAD-advised process currently employed in the LAS control room. We express these as percentages: higher values indicate more differences in vehicle choice.

5 Results

5.1 Benchmark Generation Results

Figure 4 compares distributions of travel times for 2000 journeys between vehicle and incident locations. Historical travel times are shown in blue ($\mu = 426$ s); travel times for the same journeys estimated by QUEST are shown in green ($\mu = 441$ s) and those estimated by GMAPS are shown in red ($\mu = 768$ s). The Wasserstein distance from the historically observed distribution of journey times is 40.9 for QUEST and 336.39 for GMAPS. These results show that QUEST produces travel time estimates that closely agree with historical travel times while GMAPS tends to overestimate them. Both routing engines employ traffic models that are tuned for specific journey times (a given hour- or minute-of-the-week). However, QUEST’s estimates are based on road speeds of emergency vehicles, which obey different traffic rules and tend to be higher than those of passenger or commercial vehicles, which GMAPS targets. These results validate and extend previous work that demonstrated the accuracy of the QUEST routing engine when compared to a simple model of computing travel times based on straight-line Euclidean distances [27]. The results presented here demonstrate QUEST’s accuracy even when compared with GMAPS, a state-of-the-art routing engine. These benchmark results also provide a measure of confidence in the accuracy of our auction-based results as compared with related work that employed GMAPS [23].

5.2 Auction Results

The results of auction-based allocation under the four experimental conditions are shown in Tables 2–3 and Figure 5. Focusing on results obtained using the QUEST routing engine, under the 1M-1C condition, response times were reduced from 396 to 205 seconds; under the 12M-1C condition from 460 to 155 seconds; under the 1M-nC condition from 437 to 170 seconds; and under the 12M-nC condition from 407 to 187 seconds.

Table 3 shows the proportion (percentage) of times auction-based allocation chose a vehicle to dispatch that was different to the vehicle that was historically dispatched to an incident. With both the QUEST and GMAPS routing engines, a different vehicle was chosen 89+% of the time across all experimental conditions.

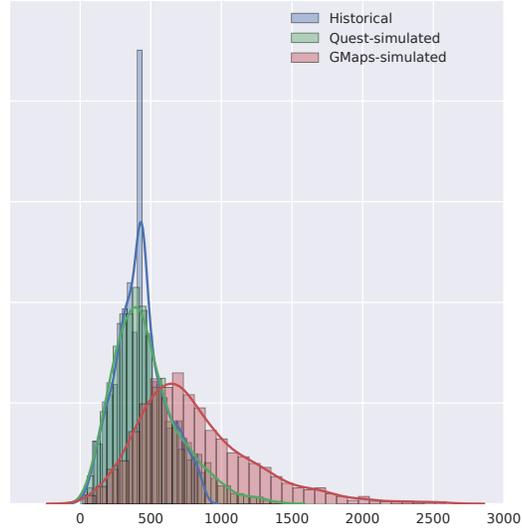


Fig. 4: Distributions of travel times to incidents (seconds). Historical travel times (blue) compare with those of the same journeys estimated by the Quest routing engine (green) and the Google Maps Directions API (red).

This indicates that, historically, there was often an alternative vehicle that could have reached an incident location sooner given our assumptions about idle vehicle locations (discussed below). Note that there is no value judgement inherently attached to the percentage difference values, but it is interesting to be able to consider that a high percentage of differences implies that the methodology evaluated here is predicted to behave differently from the current system, as borne out in the improved response times.

As a whole, these results indicate a potentially large reduction in response times when using an auction-based approach to dispatching. The auction mechanism in these experiments produced allocations in a way similar to the “closest available vehicle” strategy currently employed by dispatchers at the LAS, but using a different method to assess what “closest” means. One factor that may explain the difference in average response times is the accuracy of the QUEST routing engine when compared with the routing engine used by the LAS at the time that the data set was recorded. The LAS estimate vehicle travel times using a method that considers the types of road segments along a proposed route (i.e., number of lanes) but not current or historical traffic conditions. The QUEST routing engine considers historical traffic conditions and right-of-way rules that apply to emergency services vehicles when estimating travel times, and so produces more accurate estimates than state-of-the-art non-specialised alternatives such as GMAPS.

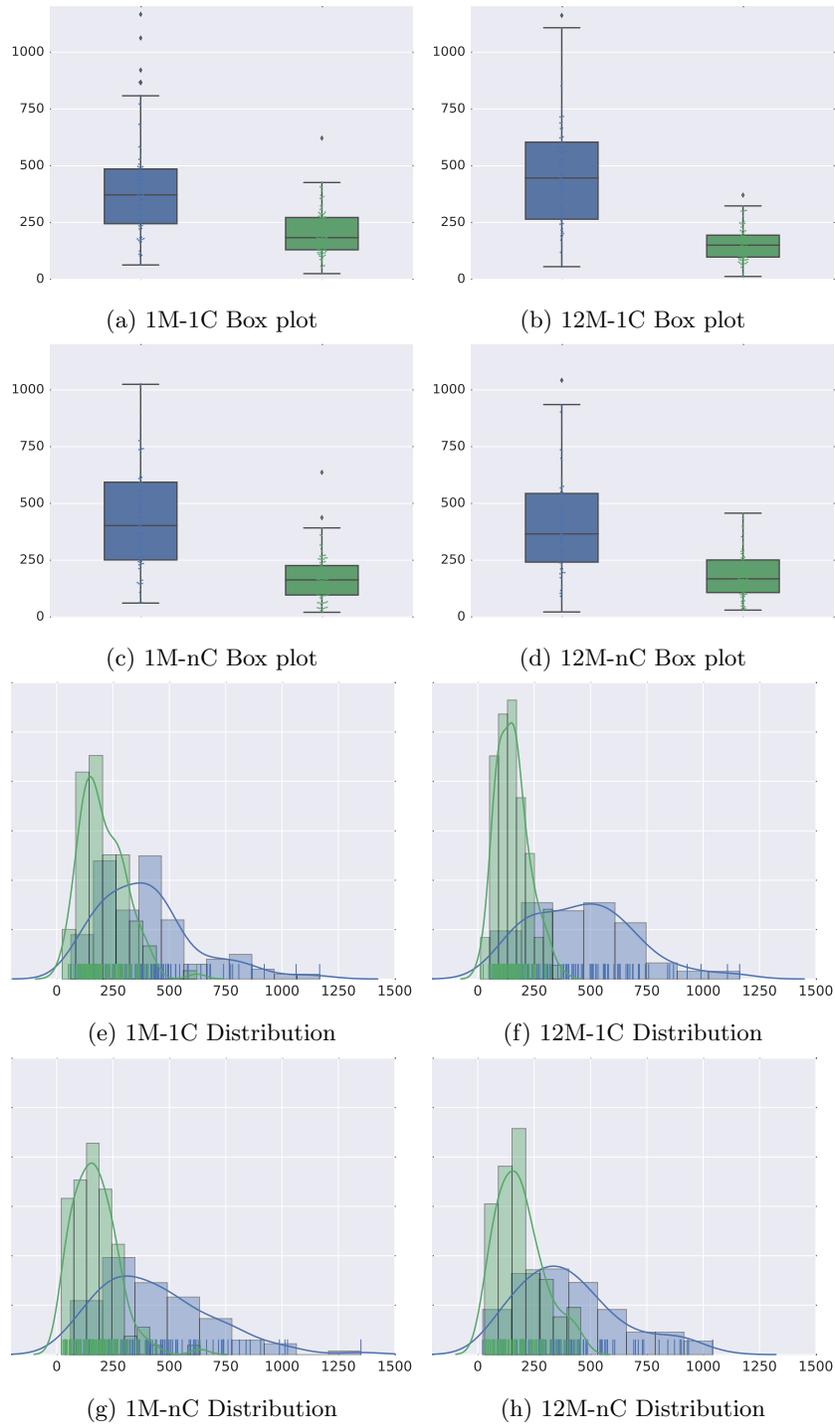


Fig. 5: Auction results for QUEST across the four conditions: 1M-1C, 12M-1C, 1M-nC and 12M-nC. For each condition, we present both a box plot of travel times, and a representation of the distribution. For both box plot and distribution, blue indicates simulated historical travel times (HIST), and green indicates the estimated travel time of the auction winner (AUCTION).

Exp. condition		GMAPS	QUEST
1M-1C	HIST	682.81	396.27
	AUCT	263.36	205.41
	<i>t</i> -statistic (<i>p</i> -value)	10.5 (8.7e−21)	9.7 (1.9e−18)
12M-1C	HIST	807.37	460.04
	AUCT	265.38	154.50
	<i>t</i> -statistic (<i>p</i> -value)	7.99 (1.1e−13)	12.28 (9.8e−26)
1M-nC	HIST	730.03	437.44
	AUCT	272.83	170.28
	<i>t</i> -statistic (<i>p</i> -value)	5.29 (3.3e−7)	9.02 (1.9e−16)
12M-nC	HIST	741.97	407.45
	AUCT	279.92	186.69
	<i>t</i> -statistic (<i>p</i> -value)	3.51 (5.5e−4)	4.83 (3.0e−6)

Table 2: Historical and auction-based response times compared. Values are average response times in seconds with a 2-tailed *t*-statistic.

Experimental Condition	GMAPS	QUEST
1M-1C	93%	89%
12M-1C	96%	97%
1M-nC	91%	92%
12M-nC	92%	94%

Table 3: Percentage (%) of times that auction-based allocation chose a different vehicle to that which was dispatched historically.

These results are based on several assumptions. The identities and locations of idle vehicles were not present in the data set provided by the LAS and needed to be estimated. Incidents were assumed to be independent: the effect of assigning a vehicle to an incident, possibly moving it away from responding to subsequent incidents in its idle area of coverage, were not modelled. Nevertheless, the auction-based model is attractive because the bid each vehicle agent computes can be extended to consider factors other than estimated distance or travel time, factors such as the cost of removing a vehicle from an area of service (decreasing the equity of coverage), crew fatigue, the ability of a vehicle to convey a patient, or the presence of specialist equipment or skills of personnel on board the vehicle. A key factor remains the ability of a routing engine to accurately estimate travel times, possibly enhanced by real-time traffic data. These

factors that comprise the suitability of a vehicle to respond to an incident can be clearly presented to a human dispatcher who makes an ultimate assignment decision.

6 Related work

There are four main areas of related work on ambulance dispatch: applying new information technologies (IT); predicting demand; predicting response time; and identifying the optimum location of emergency services.

Applying new technologies to support emergency response includes a wide range of data-centric modelling and decision-support solutions. Zhou *et al.* [39] created a geo-temporal model of ambulance demand in Toronto (Canada) and demonstrated that such modelling could lead to more accurate predictions of operational results than current industry standards. This is one of a number of studies that have investigated application of various modelling methodologies to better understand the range of factors that influence emergency response [26,6,8,15]. A number of approaches for decision-support systems to aid emergency services have been explored, primarily by analysing data from past incidents [1,2,33]. The problem of providing information to citizens and responders during incidents has been studied by [28], who focussed in particular on ways to communicate with citizens via mobile devices to provide live updates and instructions. Zadorozhny & Lewis [37] consider the problem of *information fusion*. Although their example scenario concerns robot-aided urban search-and-rescue, they address the question of data reliability and propose a *crowdsourcing* approach to mitigate the adverse effects of inaccurate or incomplete information, an approach also taken in [5]. Collectively, these studies demonstrate the potential of non-traditional data-backed, technology-based methods to improve ambulance response.

Predicting incident demand is perhaps the largest area of EMS research, and focuses on predicting demand of a population across a day, week or year. It is important for EMS personnel to understand demand in order to have appropriate numbers of ambulances on shift. The moving average method, which is commonly used by ambulance services in the US to predict demand [32], is based on an average of the call volume of one hour time periods on a specific day in four consecutive weeks over the previous five years. This can be used to predict demand for a specific location as well as for an entire city.

Separate models for both daily and hourly demand have been developed by [7] based on data for Calgary (Canada) during 2000–2004. This suggests that there is an overall increase in demand over the four years, with larger volume in July and December. Special Days, where the demand is unusually large, can be identified—these include New Year’s Day and the annual Calgary Stampede event. Call arrival data from Toronto is the basis of the Poisson-based model developed in [24]. Here New Year’s Eve and New Year’s Day were Special Days. Vile *et al.* [36] analysed demand data from the Welsh Ambulance Service Trust (UK), once again showing that there are daily, weekly and yearly periodicities as well as Special Days (in this case, all Special Days were New Year’s Day in

different years). There was also an overall positive increase in demand across the 57 months for which data was available.

As well as understanding when demand is expected to be particularly high or low, research also investigates the distribution of demand across geographic regions. For example, Kamenetzky *et al.* [17] developed a model to predict demand across any area of Southwestern Pennsylvania (US), using regression analysis based on 1979 data from 82 ambulance services in the region. Spatio-Temporal analysis provides more precise demand models by combining the two techniques discussed above, predicting demand based on the time of day for specific areas of a population. Setlzer *et al.* [32] developed such a model based on Artificial Neural Networks to improve prediction forecasts for the Charlotte-Mecklenburg region of North Carolina (US) beyond the accuracy and precision of the MEDIC model. Other work has to develop more accurate methods of predicting ambulance demand for more precise areas [38,39,40] has also developed models that are significantly more accurate than MEDIC.

As well as simply predicting ambulance demand, further work investigates developing models which directly *predict ambulance response time*, based on a prediction of call arrivals. In this line, Scott *et al.* [31] developed a probabilistic model that was fitted to a random 28-day sample from data for Houston, Texas (US) between July 1973 and June 1974. Similarly Taylor [34] modelled response times for the London Fire Brigade using survival analysis, and Thornes *et al.* [35] list factors affecting response times which include: the number of ambulances; congestion in A & E; and weather and consequent road conditions.

Further considerations in optimising the ambulance services include the *location of EMS facilities* to enable adequate coverage across the city. For example, Gendreau *et al.* [11] looked for optimum ambulance locations in Montréal (Canada) under a double coverage model, and Higgins *et al.* [13] describe a spatial model to identify communities most at risk from fires around Merseyside (UK) based on fire station location.

Poulton & Roussos [27] developed a routing engine and simulation framework to evaluate the performance of ambulance dispatching and relocation methods. Their relocation model, which seeks to provide geographic coverage for current and anticipated emergency incidents, improved on historical response times on a sample of emergency incidents drawn from the Greater London (UK) area in 2011.

Lujak & Billhardt showed that auction-based approaches to ambulance allocation, applied to a sample of emergency incidents drawn from Madrid in 2009, led to reductions in ambulance travel distance and response times as compared with a first-come-first-served approach [23]. In contrast to the work presented here, their approach used the GMAPS routing engine with a combinatorial auction mechanism, the computational costs of which scale exponentially with the number of tasks and agents [4] and is unlikely to be able to municipal-sized dispatching problems.

7 Summary

This paper has described our work applying techniques from multi-robot routing to the problem of ambulance dispatch at the London Ambulance Service. Our results strongly suggest that a combination of accurate route plan estimation and auction-based vehicle selection has the potential to significantly reduce response times—which was the case for all four experimental conditions we evaluated. Of the four experimental conditions that we examined, the worst average performance across 100 incidents was for the 1M-1C condition (which corresponds to January 2016 in the Harringey CCG), where the average simulated response time for the auction mechanism choice was 48% faster than the average simulated response time for the historically chosen vehicles. For the 12M-nC condition (2016 across the whole of London), the average simulated response time for the auction mechanism choice was 54% faster than that of the historically chosen vehicle.

The type of auction-based resource allocation mechanism presented here could also be applied to the task of an ambulance crew deciding which hospital to transport a patient to, termed *conveyance* by the LAS. It is not always sensible to bring a patient to the nearest hospital due to factors such as a patient’s need for access to special equipment, services or medical specialists, the proximity to a patient’s home for ease of family visits, the location where a patient has previously been treated, the current waiting time at the hospital’s emergency room (termed “A & E” in the UK), or the number of available beds in the hospital. In the case of stroke patients, for example, it has been shown [16] that minimising time to treatment — what is known as “door to needle time” — is best achieved not by conveying patients to the nearest hospital, but by taking them to a specialist stroke unit. A post-response, pre-conveyance auction could take place, where the ambulance is the “auctioneer” and the hospitals are the “bidders” to address exactly this issue.

Acknowledgements This work was partially funded by ESRC through the Data Awareness for Sending Help (DASH) grant ES/P011160/1. We are grateful to our collaborators from the King’s College London Policy Institute, especially Dr Alexandra Pollitt.

References

1. Bartels, M.: Communicating Probability: A challenge for decision support systems. In: ISCRAM. Pennsylvania, US (May 2014)
2. Barthe-Delano, A.M., Truptil, S., Benaben, F.: Agility of crisis response: gathering and analyzing data through an event-driven platform. In: ISCRAM. Pennsylvania, US (May 2014)
3. Bektas, T.: The multiple traveling salesman problem: an overview of formulations and solution procedures. *Omega* **34**(3), 209–219 (2006)

4. Berhault, M., Huang, H., Keskinocak, P., Koenig, S., Elmaghraby, W., Griffin, P., Kleywegt, A.: Robot exploration with combinatorial auctions. In: Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on. vol. 2, pp. 1957–1962. IEEE (2003)
5. Calderon, A.C., Hinds, J., Johnson, P.: IntCris: A tool for enhanced communication and collective decision-making during crises. In: ISCRAM. Pennsylvania, US (May 2014)
6. Chalk, D., Black, S., Pitt, M.: Which factors most influence demand for ambulances in South West England? *Journal of Paramedic Practice* **8**(7) (2016)
7. Channouf, N., L’Ecuyer, P., Ingolfsson, A., Avramidis, A.N.: The application of forecasting techniques to modeling emergency medical system calls in Calgary, Alberta. *Health Care Management Science* **10**(1), 25–45 (2007)
8. Dietze, P., Jolley, D., , Cvetkovski, S.: Patterns and Characteristics of Ambulance Attendance at Heroin Overdose at a Local-Area Level in Melbourne, Australia: Implications for Service Provision. *Journal of Urban Health: Bulletin of the New York Academy of Medicine* **80**(2) (2003)
9. Drake, Archie and Pollitt, Alexandra and Sklar, Elizabeth and Smith, Leanne and Parsons, Simon and Schneider, Eric: Data for Ambulance Dispatch: New & emerging forms of data to support the London Ambulance Service. Policy Report, Policy Institute, King’s College London (2018)
10. Finkelstein, A., Dowell, J.: A Comedy of Errors: the London Ambulance Service case study. In: IWSSD-8 (1996)
11. Gendreau, M., Laporte, G., Semet, F.: Solving an ambulance location model by tabu search. *Location Science* **5**(2), 75–88 (1997)
12. Gerkey, B.P., Mataric, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems. *IJRR* **23**(9), 939–954 (2004)
13. Higgins, E., Taylor, M., Jones, M., Lisboa, P.: Understanding community fire risk: A spatial model for targeting fire prevention activities. *Fire Safety Journal* **62**, 20–29 (2013)
14. HSCIC: London Ambulance Service case study: CommandPoint assurance (2013)
15. IAS: Alcohol’s impact on emergency services (October 2015)
16. James, M.: Acute management of ischaemic stroke (2016)
17. Kamenetzky, R.D., Shuman, L.J., Wolfe, H.: Estimating need and demand for prehospital care. *Op. Res.* **30**(6), 1148–1167 (1982)
18. Koenig, S., Tovey, C., Lagoudakis, M., Kempe, D., Keskinocak, P., Kleywegt, A., Meyerson, A., Jain, S.: The Power of Sequential Single-Item Auctions for Agent Coordination. In: Proceedings of National Conference on Artificial Intelligence (2006)
19. Lagoudakis, M., Markakis, V., Kempe, D., Keskinocak, P., Koenig, S., Kleywegt, A., Tovey, C., Meyerson, A., Jain, S.: Auction-based multi-robot routing. In: RSS (2005)
20. Landén, D., Heintz, F., Doherty, P.: Complex task allocation in mixed-initiative delegation: A UAV case study. In: PRIMA. pp. 288–303. Springer (2010)
21. Laporte, G.: The vehicle routing problem: An overview of exact and approximate algorithms. *Euro. J. Op. Res.* **59**(3), 345–358 (1992)
22. Long, T.: Oct. 26, 1992: Software Glitch Cripples Ambulance Service. *Wired* (26 October 2009)
23. Lujak, M., Billhardt, H.: Coordinating emergency medical assistance. In: Agreement Technologies, pp. 597–609. Springer (2013)
24. Matteson, D.S., McLean, M.W., Woodard, D.B., Henderson, S.G.: Forecasting emergency medical service call arrival rates. *Annals of Applied Stats* pp. 1379–1406 (2011)

25. Office for National Statistics: 2011 Census aggregate data. UK Data Service (Edition: June 2016) (2016). <https://doi.org/http://dx.doi.org/10.5257/census/aggregate-2011-1>
26. PenCHORD: <http://clahrcpeninsula.nihr.ac.uk/penchord> (2016)
27. Poulton, M., Roussos, G.: Towards smarter metropolitan emergency response. In: Personal Indoor and Mobile Radio Communications (PIMRC), 2013 IEEE 24th International Symposium on. pp. 2576–2580. IEEE (2013)
28. Ruiz-Zafra, A., Nez, A.G., Penads, M.C., Cans, J.H., Borges, M.R.S.: SUCRE: Supporting Users, Controllers and Responders in Emergencies. In: ISCRAM. Pennsylvania, US (May 2014)
29. Schneider, E.: Mechanism Selection for Multi-Robot Task Allocation. Ph.D. thesis, University of Liverpool (2018)
30. Schneider, E., Sklar, E.I., Parsons, S., Özgelen, A.T.: Auction-based task allocation for multi-robot teams in dynamic environments. In: TAROS (2015)
31. Scott, D.W., Factor, L.E., Gorry, G.A.: Predicting the response time of an urban ambulance system. *Health Services Research* **13**(4), 404 (1978)
32. Setzler, H., Saydam, C., Park, S.: EMS call volume predictions: A comparative study. *Comp. & Op. Res.* **36**(6), 1843–1851 (2009)
33. Sutton, J., Spiro, E., Fitzhugh, S., Johnson, B., Gibson, B., Butts, C.T.: Online Message Amplification in the Boston Bombing Response. In: ISCRAM. Pennsylvania, US (May 2014)
34. Taylor, B.M.: Spatial modelling of emergency service response times. *JRSS: Series A* **180**(2), 433–453 (2017)
35. Thornes, J.E., Fisher, P.A., Rayment-Bishop, T., Smith, C.: Ambulance call-outs and response times in birmingham and the impact of extreme weather and climate change. *Emerg Med J* **31**(3), 220–228 (2014)
36. Vile, J.L., Gillard, J., Harper, P.R., Knight, V.A.: Predicting ambulance demand using singular spectrum analysis. *JORS* **63**(11), 1556–1565 (2012)
37. Zadorozhny, V., Lewis, M.: Information fusion for USAR operations based on crowdsourcing. In: FUSION. Istanbul, Turkey (July 2013)
38. Zhou, Z., Matteson, D.S.: Predicting ambulance demand: A spatio-temporal kernel approach. In: KDD. pp. 2297–2303. ACM (2015)
39. Zhou, Z., Matteson, D.S., Woodard, D.B., Henderson, S.G., Micheas, A.C.: A spatio-temporal point process model for ambulance demand. *J. Am. Stat. Assoc.* **110**(509), 6–15 (2015)
40. Zhou, Z., Matteson, D.S., et al.: Predicting Melbourne ambulance demand using kernel warping. *Annals App. Stats* **10**(4), 1977–1996 (2016)