

Intelligent Resource Exchanges: Solutions and Pathways in a Workforce Allocation Problem

Botond Virginas

(Intelligent Systems Lab, British Telecom, UK
botond.virginas@bt.com)

Marian Ursu

(Goldsmith College, University of London, UK
m.ursu@gold.ac.uk)

Edward Tsang

(University of Essex, Colchester, UK
edward@essex.ac.uk)

Gilbert Owusu

(Intelligent Systems Lab, British Telecom, UK
gilbert.owusu@bt.com)

Chris Voudouris

(Intelligent Systems Lab, British Telecom, UK
chris.voudouris@bt.com)

Abstract: This paper considers the problem of resource allocation in the service industries approached from an agent-based perspective. Agent technologies seem to be well suited to this domain by providing a distributed environment, are network centric, semi-autonomous and collaborative and can communicate with each other to achieve better optimisation with little human intervention. The paper describes the context of this solution, a general power model and several pathways with corresponding example implementations with results and discussion. The novelty of the solution resides in the fact that it is a natural and versatile formulation that combines an agent-based model with various artificial intelligence and operations research techniques such as rule-based expressions of allocation strategies and multi-criteria optimisation expressions of allocation objectives.

Keywords: multi-agent planning, e-business agents, case studies and reports on deployments

Categories: D.0, D.1.1, D.1.5, D.2.2, D.2.3, D.2.6, D.2.7

1 Introduction

The effective management of resources is critical to optimal service delivery in service organisations such as British Telecom (BT). BT manages the largest access network in the United Kingdom (UK) with a workforce of around 25,000. Resource

management is a complex process, usually involving the analysis of large amount of information. The complexity increases when more than one objective is being evaluated and the number of variables to consider is large. Clearly there is the need for automating resource planning, which is well recognised and has been the subject of considerable research and development [Aarts and Lenstra, 97], [Buchanan, 95a]. The case for automating resource management is motivated by the drive to maximise profits, improve quality of service and reduce costs.

The problem of resource allocation occurs in very many contexts/applications and is approached, in computational terms, within various theories, models and architectures. For example, in terms of applications, it arises in telecommunication networks (e.g., allocation of bandwidth), operating systems (e.g., memory allocation), and manufacturing (e.g., production scheduling). In terms of computational models, the problem may be formulated using constraint satisfaction (e.g. [Le Pape, 94]), heuristic search methods (e.g., [Lesaint, et al., 00a], [Reeves, 93a]), evolutionary algorithms [Coello], and fuzzy-logic (e.g., [Andersson, 00b]) approaches.

In order for a company like BT to best serve its customers, resource managers within the organisation have to ascertain how best to plan and deploy the company's field engineering workforce on an operational/tactical basis (i.e. for the following day and up to 14 days ahead). A Field Optimisation Suite (*FOS*) has been developed for this purpose within BT [Voudouris et al., 06a]. *FOS* incorporates workload forecasting, optimised workforce planning, as well as advanced tools for visualising and communicating the outputs to end users. *FOS* employs a variety of advanced Operations Research (OR) techniques such as constraint satisfaction for problem modelling, and heuristic search methods for problem solving ([Voudouris et al., 01a], [Shen and Norrie, 99a]).

This paper is based on FieldExchange, a component of *FOS* in charge of monitoring and supporting resource re-distribution decision-making in BT's Operational Resource Management units. In contrast to the other components of *FOS* FieldExchange provides an experimental base for applying agent-based techniques to the problem.

The motivation for applying agent technologies [Wooldridge and Jennings, 95b] draws back to B2B commerce, in particular the use agent-based technologies in supply chain management. Intelligent agents have been demonstrated to be useful structures and technologies in the context of supply chain management ([Fox, et al., 00c], [Walsh and Wellman, 99b]). We consider resource management under the broader term of service chain management. Service chain management is concerned with the optimal provisioning and management of services to customers in order to satisfy customer service demand, whilst minimising the operational costs of supplying services. Could intelligent agents be also utilised in service chain management to trade services and associated resources? Although supply chain management has many similarities to service chain management there are significant differences between the two domains due to several key differences between the nature of service and supply. Most notably, services are not tangible objects. Supply chain management is concerned with the flow of materials to and from suppliers eventually to the end customer. In the supply chain, materials may be measured by physical quality and quantity, however in the service chain the definition of quality of service may become abstract. The key difference in this sense is quality of service. Services are

inconsistent and dynamic and the human factor could be a hindrance in assessing quality of service. Furthermore, this problem may become even more apparent when the acquisition of services involves supply from contractor service providers.

However, effective service chain management is vital in today's economy and organisations have started investigating the agility and automation offered by agent-based approaches. Although agent-based solutions have been extensively reported on partnership formation, brokering and negotiation, primarily in the context of traditional supply chain management, they also hold promise for service chain management too. The reason why they are most useful in partnership formation, brokering and negotiation is because these stages all involve complex issues related to decision making, searching, and matchmaking that agent technologies are well suited to [Kurbel and Loutchko, 01b]. Moreover, agent technologies seem to be well suited to this domain by providing a distributed environment, are network centric, semi-autonomous and collaborative and can communicate with each other to achieve better optimisation with little human intervention [He et al., 03a].

The breakdown of the remainder of this paper is as follows. In Section 2, we describe a generic resource management problem. In Section 3 we present a power model formulated around a multi-agent abstraction. In sections 4, 5, and 6 we provide an overview of three different types of resource exchanges which could be implemented based on the common power model with proposed solutions and relevant implementations. In section 7 we discuss results of existing implementations. Finally section 8 provides future pathways and concludes the paper.

2 Problem Description

In this section we describe the FieldExchange problem domain focusing on its main characteristics. In terms of their management, large companies are partitioned in organisational units. These units could take the form of a number of divisions or domains (which can represent geographical service areas or defined otherwise) participating in a resource allocation process or resource exchange trying to optimise the workforce allocation between them by exchanging resources. This can be in the context of electricity, telecom, water or other utility which wants to move resources between divisions as and when required to improve its customer service. Each participating domain or area could have a number of idle or surplus resources (i.e. service personnel such as field engineers) and a number of jobs that need to be resolved. The idle resource cannot resolve the jobs within their own domains because of conflicting preferences or constraints. For example a field engineer might not have the necessary set of skills required to complete a local job however he/she could be deployed to a neighbouring area with pending jobs requiring his/her skill.

The planning scenario consists of jobs which have to be processed by field engineers from other domains. Each job could be characterised by a set of parameters such as its geographical location, a required skill, a processing time window and a priority level, which is an indication of how quickly it needs to be resolved. Field engineers on the other hand can work in multiple areas, can offer multiple skills, can have preferences with regards to the skills they normally wish to employ, the locations where they normally wish to travel, and the days when they wish to travel; the list of preferences can be longer. The main relationship on the basis of which the

allocation of engineers to jobs is carried out is that of a match: a set of resources must match the set of requirements to which they are allocated.

In FOS we follow a two-phase approach: phase 1 – local allocation and phase 2 – global allocation of surpluses and shortages left after phase 1. In the first phase, for each domain, local engineers are allocated to local jobs. Phase 2 considers the “leftovers” from phase 1 – surpluses (idle engineers) and shortages (unresolved jobs). Each allocation within the second phase has an extra cost, compared to local allocations, due to longer travelling and, possibly, overnight accommodation. An effective solution, here, not only maximises the number of resolved jobs, but also minimises this extra cost [see Fig. 1].




Figure 1: A two-phase workforce allocation process

A solution to local allocation (phase 1) has been developed and reported elsewhere ([Voudouris et al., 03b], [Owusu et al., 06b]). FieldExchange is in charge of global allocation monitoring and supporting resource re-distribution decision-making in BT’s Operational Resource Management units. We will describe the problem model in the next section.

3 Centralised and Distributed Power Structures

A common model has been introduced to capture an enterprise’s distribution of intelligence and mode of operation related to service chain management. Each domain or organisational unit described in section 2 has a set of resources (field engineers, for BT), a set of requirements (jobs, for BT) and is managed by a domain manager. The domain manager is the local repository of intelligence/knowledge – decision criteria and decision-making strategies – and power – the ability to enforce decisions – regarding the domain’s requirements and resources. Local power may be recognised by other domain managers or its enforcement may require the intervention of central company managerial mediators. A computational model appropriate for such an organisation consists of a set of software agents. We have decided to formulate this model using a multi-agent abstraction, primarily because each domain manager pursues different objectives using different strategies and has access to different privileged information. Other contributing factors to this decision were that

these domains are geographically distributed and managed by independent human stakeholders.

In this model each domain has a Service Agency. The Service Agency is responsible for the management of resources and requirements for the particular region. These Service Agencies are managed by a Central Agency. For the purpose of FieldExchange each Service Agency consists of a Service Buyer and a Service Seller. Each Service Buyer has a set of jobs to be completed. Each Service Seller manages a set of field engineers, who can be assigned to jobs. Service Sellers and Service Buyers enact the geographic region's interest and priorities via the decisions/choices they make. These local interests and priorities may be expressed via a set of criteria/objectives that are to be optimised (e.g., "minimise the field engineers' overall travelling distance"), constraints that are to be satisfied (e.g., "a field engineer should not travel for more than 600 miles a week") and rules of operation or strategies that are to be followed (e.g., attend to jobs of higher priority first).

The Service Buyers buy services from the sellers from other Service Agencies by expressing their requirements (e.g. inviting them to bid for jobs). The Service Sellers base their responses on which engineers could service which job depending on a set of constraints and preferences. Example of these include availability and skills of engineers, the engineers' preferences and their travelling distances to the jobs. Fig. 3 illustrates an allocation process where each domain has associated a Service Seller Agent (SerSelA) and a Service Buyer Agent (SerBuyA) and the region has a Central Agency. Agents interact (<see> Fig. 2) in the process of resource allocation. The global allocation of resources is accomplished through communication and exchanges between agents. This has been approached from two perspectives: centralised approach and distributed approach. Obviously, they represent extreme points of view; many combinations between them are possible.




Figure 2: An allocation model where each domain has a Service Buyer and Service Seller Agent associated with them with a Central Agency also employed in the allocation process

In the centralised approach, the focus is on the company's global interests and priorities. They are expressed via a set of global criteria. There is no decision power expressed at domain/local level. In the distributed approach, the focus is on local interests and priorities. In a purely distributed approach, all the intelligence is

located/represented in the constituent agents – i.e. in the leaves of the system. The overall behaviour of the system emerges completely from the agents' interaction from different service agencies. The company's overall interests and priorities are not explicitly expressed – there is no central manager to enforce them. Their accomplishment should emerge from the individual accomplishment of the local ones. Examples of interaction structures include explicit asynchronous co-ordination (e.g., collaboration, competition and negotiation) and synchronism (e.g., request for resources must be received by a certain deadline).

Embodiments of possible systems implemented based on the common model may be configured to act in different ways to assist the redistribution of resources between entities within an organisation or across the service chain (e.g. automating the interactions of a company with its service subcontractors). These configurations can be grouped in various types of models based on a number of criteria. A number of scenarios linked to different types of resource exchanges can be identified. The next section will describe these scenarios with relevant implementations. These solutions and implementations are based on BT's operational requirements.

4 Centralised Resource Exchange

The first scenario we are presenting will require a collaborative model of the overall centralised system which will have a common objective to fulfil. Each domain could have a number of idle resources and the optimisation process might expand over a shorter or longer planning period. One important aim of this process would be to maximise the number of jobs resolved throughout a region while minimising the travel and lodging costs associated with moving personnel. Two solutions have been created for this scenario.

4.1 FieldExchange

One solution was driven by a concrete business need of centrally balancing the failure rates for individual geographic areas. Each individual domain will have a specific demand/capacity ration which is called failure rate. The main objective of FieldExchange is to balance this failure rate across an entire region or geographic area. An optimisation software toolkit has been employed for this purpose to program the internals of the Exchange Agent [Voudouris et al., 01a]. We have implemented a system which is currently under trial in a number of operational Resource Management units within BT.

In this solution the Central Agency takes a very important optimisation role. It is using a search framework based on BT's iOpt optimisation toolkit, which was built for modelling and solving combinatorial problems using invariants (one-way constraints) and heuristic search methods. The Central Agency collects all job requests (shortages) from all Service Buyers and field engineer data (surpluses) from Service Sellers and it's overall aim is to balance the demand/capacity ratio across a geographic area. Once proposals are broadcasted back to Service Sellers they will allocate individual engineers to requests based on domain preferences/constraints and respond with resource offers.

To solve the balancing problem, we've designed a variable neighbourhood search (VNS) based framework and six individual heuristics. The reason behind selecting the use of a variable neighbourhood search was to do with the real time nature of the system where heuristics are required. In our example the variable neighbourhood search has shown some superiority to methods such as tabu search, simulated annealing and BestCNS all provided within iOpt. Central to this approach is an objective function which determines the quality of each candidate solution.

Since engineers in each domain can be moved to another domain for just one day or a couple of days, we consider two kinds of moves: a daily move which is only for one day, and an accommodation move which is for more than one day. The cost for moving one engineer from one domain to another varies according to the distance between the two domains and other constraints. Thus, sometimes moving engineers from one domain to another directly is impossible due to the cost. To solve this problem, we've designed another kind of move which is called shuffling move. The shuffling move tries to move some engineers from domain A to domain B via the help of domain C. The cost of moves between domains A and B is high, while the cost between domains A and C, domain C and B is low. Thus, the shuffling move will move a number of engineers from domain A to C, and then move the same number of engineers from domain C to B. Thus, in total we have three kinds of moves: the daily move, the accommodation move and the shuffling move.

The six heuristics that we have designed provide for daily moves, accommodation moves, and shuffling moves for clearing the surpluses and evenly distributing failures across domains. The VNS pre-defines a sequence for the six low level heuristics. It starts by applying the first heuristic in the sequence. When a local optimum is met, the VNS will go to next heuristic in the sequence. If a better solution is found after one iteration, the VNS will go back to the first heuristic to continue the search, otherwise it jumps to the next heuristic in the sequence to search for a better solution. The search will continue until the stopping condition is being met.

This system has been engineered as an enterprise application using the three-tier software architecture model. The three layers are the visualisation layer, which is implemented as a thin client in a web browser; the business logic layer, which is the layer for generating the demand forecast and optimising the workforce usage against the demand; and the data layer, which deals with the data storage, retrieval and manipulation. The three-tiers means loose coupling between the different application entities which means easier maintenance — it is easier to modify or replace any tier without affecting the other tiers. Also the application is more flexible and new requirements can be easily accommodated without significantly disrupting the other tiers. Three-tier architectures lend themselves to be easily distributed to improve performance and load balancing/sharing. The business logic and data layers are implemented as Enterprise Java Beans (EJB) and are hosted on BEA's WebLogic application server. The data is stored on an Oracle 10 database server.

4.2 Centralised Collaborator (CC)

Another example solution is where a common goal for the system could be established in order to try to optimise the workforce allocation for the entire region. However, while the agents will have this as their main objective the system will take into account conflicting objectives of the entities. Examples of conflicting criteria

could be minimise travelling distance and maximise use of skill proficiency for resources. In this solution, the Central Agency will act as a central matchmaker that tries to satisfy requests by performing a multi-objective optimisation using hard constraints and soft constraints provided by the region's different areas. The Central Agency implements global interests and priorities and these are expressed as a set of criteria animated via a multi-criteria optimisation algorithm [Ehrgott and Gandibleux, 02]. User preferences (soft constraints) will then be used to select the best-preferred solution out of this subset. The intelligent agents representing the areas could then use local constraints to filter out requests or offers for resources that do not satisfy local interests and priorities.

We have implemented a prototype system – Centralised Collaborator (CC) – on the basis of the completely centralised model. Initially, the Central Agency collects all the job requests from all Service Buyers and ranks them according to their importance. The Central Agency, then, enters an iterative process whereby, at each iteration, a job is allocated with appropriate resources. Each iteration deals with the most important job from the ones that remained unresolved.

The selection is implemented as a Pareto optimisation – the chosen optimal set is the set of non-dominated solutions or the Pareto front. We use a greedy algorithm to construct the Pareto front. Currently, the optimisation algorithm uses two criteria: minimise travelling distance and maximise use of skill proficiency. The optimal set is sent to the Service Buyer whose request is under current consideration.

The Service Buyer uses local constraints to filter out offers that do not satisfy local interests and priorities. The Service Buyer, then, selects one offer from the filtered set using specified selection strategies and communicates its choice to the Central Agency. Finally, the Central Agency notifies the Service Seller whose offer was selected and the process is resumed with the next most important job request.

This prototype system has been engineered as a Java application. The GUI agents in the prototype are the Global Monitor (stakeholder: region manager) and Local Monitors (stakeholder: domain manager). The control features provided by the Global Monitor include parameter setting (such as, planning period, and global criteria) and assistance towards the region manager regarding monitoring the allocation of resources (e.g. visualisation of results). The control features provided by local monitors include setting of local constraints, manual selection of requests or resources to be submitted to CC, manual overriding of individual allocations and data visualization. This implementation has been reported in detail elsewhere [Virginas et al., 03c].

5 Distributed Resource Exchange

Our second scenario relates to a number of domains participating in a resource exchange process and trying to optimise the workforce allocation between them, but in this instance the emphasis is placed on local interests and priorities. In the same way as in the previous scenario, each domain could have a number of idle resources and the planning period could expand over shorter or longer periods. The solution will require a competitive model of the overall distributed system. The distributed scenario is a promising one where multiple entities have to work together and centralised resource planning looks difficult for a variety of reasons. For example different entities could have potentially conflicting interests or individual entities could have many constraints to consider, and these constraints could change dynamically. One

possible solution that could be envisaged for this scenario is where the region's overall interests are not explicitly expressed and their accomplishment emerges from the individual accomplishment of the local ones.

5.1 Distributed Collaborator (DC)

We have implemented an example prototype system - Distributed Collaborator (DC) - as an almost isomorphic implementation of the completely and uniformly distributed model. In this solution the workforce allocation is modelled as an iterative communication process, based for example on a 4-step communication contract net protocol, between Service Buyers, on one hand, and Service Sellers on the other. Decisions regarding the allocation are made locally, within each agent. A global allocation would emerge from this interaction/communication.

The communication protocol has two stages: information gathering and contract establishment. The information exchanged during the information gathering stage bears no legal obligations. For example, a job's agent may request engineers for double the amount of jobs it actually holds. Therefore, various strategies, employing bogus or incomplete information, may be employed here, in order to attract favourable contracts in the following stage. The information exchanged in the contract establishment stage is legally binding.

In step one, all Service Buyers broadcast requests, based on their job requirements. In the simplest case, each Service Buyer broadcasts requests for all its jobs to all Service Sellers. This is the solution we initially adopted in DC. We are now experimenting with more elaborated strategies like preferential and bogus broadcasts.

In step two, Service Sellers respond with offers. Each Service Seller offers its best matching set of engineers to each Service Buyer.

In step three, Service Buyers, faced with various offers, must decide to whom to propose contracts (legally binding). In DC each Service Buyer aims to maximise the number of jobs attended to and to minimise the travelling costs it has to support. This could be achieved via a local multi-objective optimisation algorithm. DC takes the "strategy route" – it implements strategies of compiling contract proposals, mimicking the behaviour of domain managers.

In step four, Service Sellers decide whom to contract their resources. In DC, each Service Seller attempts to maximise the number of jobs to attend, whilst maintaining the overall travelling distance within reasonable limits. Here, too, DC takes the "strategy route".

The synchronisation of the agents' behaviour, necessary for the realisation of the 4-step communication protocol, is achieved via the Central Agency which would take the role of a *Monitor*. This ensures that a step is initiated only after all the agents have indeed completed the previous step. This prototype system has been engineered as a Java application using the Java messaging service on a BEA Weblogic Service for communication purposes. This implementation has been reported in detail elsewhere [Ursu et al., 04a].

6 Hybrid Power Resource Exchange

Our third scenario relates to a situation where individual Service Buyer and Service Seller agents have potentially conflicting interests but there is also a central objective. The business objective in this case is to allow for a distributed market to operate where multiple service providers have to serve multiple service buyers while the central objective is a multi-objective optimisation problem: the central manager has to strike a balance between job completion rates, service quality, travelling distances and other objectives. The Service Sellers and Service Buyers all attempt to maximise their own utility. The overall problem is considered as a multi-objective optimisation problem. We have used distributed constraint satisfaction for the solution which is highly relevant to distributed planning. Most of the work in distributed constraint satisfaction involves cooperative agents, where agents work together to achieve some common goal. In this solution we study distributed constraint satisfaction problems where agents may have conflicting goals.

6.1 ASMCR: An Open Constraint Optimisation System

We have implemented a prototype system, ASMCR, as an open constraint optimisation system where individual agents have potentially conflicting interests. We have introduced a retractable contract net protocol, which we call RECONNECT, that supports hill-climbing in the space of solutions. It is built upon a job-release and compensation mechanism. The problem of each Service Buyer and Service Seller was formulated as an open constraint optimisation problem [Tsang, 93b] where some constraints are not entirely within the control of the problem solver itself.

The Service Buyer's Model: the problem of buyer b can be formulated as an open constraint satisfaction model:

$$(Z_b, D_b, C_b, E_b, f_b, Ag_b, EtA_b, CP_b)$$

$Z_b = \{s[1], s[2], \dots, s[n_b], p[1], p[2], \dots, p[n_b], d[1], d[2], \dots, d[n_b]\}$, where n_b is the number of jobs that b has, $s[i]$ represents the service seller that b appoints to do job i , $p[i]$ represents the preference and $d[i]$ represents the distance for serving job i , which are proposed by the service sellers and accepted by the buyer;

Ag_b is the set of seller agents who b has contact with;

D_b is a function that defines the domain of the variables in Z_b , as in constraint satisfaction. For all i , $D_b(s[i]) = Ag_b$ plus ϕ , which means $s[i]$ could be assigned one of the service sellers, or assigned no seller at all (which is represented by the value ϕ); $D_b(p[i]) = \{0, 1, \dots, 9\}$, which means $p[i]$ could be assigned a value of 0 to 9, with 0 meaning the job is not being served, 1 to 9 are preferences in the service. For all distance variables $d[i]$, $D_b(d[i]) = R$;

C_b represents a set of internal constraints, which is $\{\}$ in this case, i.e. there are no constraints on what value b assigns to the variables;

$E_b = \{E_b(s[i], p[i], d[i]) \mid i = 1..n_b\}$, where $E_b(s[i], p[i], d[i])$ is a constraint on the values of $s[i]$, $p[i]$ and $d[i]$, restricting the values that they can take simultaneously; the values of $p[i]$ and $d[i]$ are to be determined by external agents, $s[i]$ indicates the seller that b assigns job i to; it is assigned by b , depending on the bids by service sellers;

f_b is the objective function for b . It is a multi-objective function.

The buyer b attempts to maximise its utility:

Utility = $k_1 * reve_b - k_2 * failure_b - k_3 * pref_b - k_4 * dist_b - k_5 * comm_b + Trade$
 where the weights k_1 to k_5 are given by the manager; Trade is income from other buyers – payment to other buyers for contract-release;

EtA_b is the mapping from each external constraint E_b to a service seller; i.e. $EtA_b(E_b(s[i], p[i], d[i]))$ returns a value in Ag_b ; this indicates that the values of $p[i]$ and $d[i]$ are to be determined by all Seller Agents (Ag_b) in communication with b ;

CP_b is the communication protocol. Here we assume the following protocol:

1. The buyer b sends a set of invitation to bid to seller s ; each invitation is

(Job_ID, Job_information)

where Job_information is a tuple as defined above:

(Location, Min_Skill, Duration, StartDay, Price)

2. The seller s sends a set of pairs of values to b for instantiating ($p[i], d[i]$)
3. The buyer b offers s a contract, which comprises a pair of p and d values
4. The seller s accepts the contract (and commits its resources) or declines the offer, in which case, go back to Step 3 (where b could offer a contract to another seller)

The Service Seller's Model: the problem of seller s can be formulated as a dynamic open constraint satisfaction model:

($Z_s, D_s, C_s, E_s, f_s, Ag_s, EtA_s, CP_s$)

$Z_s = \{e[1], e[2], \dots, e[N], p[1], p[2], \dots, p[M], d[1], d[2], \dots, d[N]\}$, where N is the total number of jobs that s has been invited to bid for and s is still in contention (i.e. the buyer has not yet assigned the job to another seller); $e[i]$ represents the engineer that is assigned to do job i ; p and d represent preferences and distances as defined in buyers;

Ag_s is the set of Service Buyers who s has contact with;

D_s is a function that defines the domain of variables in Z_s , as in constraint satisfaction. For all i , $D_s(e[i]) = \text{the set of engineers plus } \phi$, which means $e[i]$ could be assigned one of the engineers, or assigned no engineer at all (which is represented by ϕ); $D_s(p[i]) = \{0, 1, \dots, 9\}$, which means $p[i]$ could be assigned a value of 0 to 9, with 0 meaning the job is not being served, 1 to 9 are preferences in the service; For all distance variables $d[i]$, $D_s(d[i]) = R$; default values for p and d are 0, which means no engineer is assigned to job i until commitment is made (by s);

C_s represents the internal constraints that are governing the feasibility of the engineers doing the jobs; this involves the availability and skills of the technicians;

$E_s = \{ E_s(e[i], p[i], d[i]) \mid i = 1.. N \}$, where $E_s(e[i], p[i], d[i])$ is a constraint on the values of $e[i]$, $p[i]$ and $d[i]$, restricting the values that they can take simultaneously; the values of $e[i]$, $p[i]$ and $d[i]$ are proposed by s , to be approved by the Service Buyer;

f_s is the objective function for s . Associated to each job i , $Price([i])$ is a constant given by the Buyer. f_s is a multi-objective function.

The seller s attempts to maximise its utility:

Utility = $k_1 * JD - k_2 * (DT)^2 + k_3 * LB - k_4 * (RD) + CR$

where the weights k_1 to k_5 are given by the manager;

EtA_s is the mapping from each external constraint E_s to a buyer; i.e. $EtA_s(E_s(e[i], p[i], d[i]))$ equals the service buyer for job i

The objective is to balance failure rates across all domains. The Manager's overall objective is a multi-objective function. The manager should attempt to produce a Pareto set of solutions. This can be done by giving the Service Buyer/Service Seller different sets of weights for different measures for the buyers and sellers (the k_i 's mentioned above). For example the k_i 's can be used to empower individual buyers to increase their bargaining power so as to reduce their failure rates. The manager may ask the agents to schedule from scratch or improve on a previous schedule.

The ASMCR software allows the management to have full control over the company's multi-objectives. The manager generates a Pareto set of solutions by defining, for each Service Buyer and Service Seller, the weights given to each objective. ASMCR gives Service Buyers and Service Sellers ownership of their problem and freedom to maximise their performance under the criteria defined by the management. This implementation has been reported in detail elsewhere [Tsang et al., 05].

7 Discussion

We have described in this paper various solutions and pathways in a workforce planning problem. The proposed solutions and implementations based on the common model described in sections 4, 5 and 6 have provided us with a large experimental database. Most of the results have been described in more detail elsewhere as highlighted above. Please find below a brief summary of these results.

FieldExchange has recently been in trial within several resource management units, the feedback is very promising and it is in the stage of being rolled out across BT. The optimisation algorithm has demonstrated a good improvement over clearing and balancing surpluses across geographic regions. There has also been a reduction in travel time (from 95min to 85min) since technicians are properly positioned geographically to service customers first thing in the morning. Clearly this has led to an increase in productivity and morale. We are constantly improving the algorithm adapting to the business objectives driving the scope of FieldExchange.

CC and DC are being considered together looking at shifting the power structures from the centralised model towards the decentralised approach with various degrees of central control. Currently richer data models and larger data sets are being considered. The issue of strategies vs. criteria in the expression of local interests and priorities occurs, whether this is done centrally or locally. Thus far, we have experimented with criteria in the central approach and strategies in the distributed approach. We shall continue our experiments using criteria locally and strategies centrally and then combinations of the two. The initial feedback is that strategies give a better sense of control to human agents (managers). The aim is to move towards and investigate more complex decision power structures. We shall explore collaboration strategies and negotiation, which could be achieved with or without Central Agencies.

Furthermore, we are going to experiment with both synchronous and asynchronous communication protocols. So far we have implemented the synchronisation of the agents' behaviour, via the Monitor. This ensures that a step is initiated only after all the agents have indeed completed the previous step. However,

this solution lends itself easily to other methods of timing. An asynchronous model, delegates the decision of when an agent should intervene in the overall environment to the agent itself – this too, then, becomes a matter of a local choice, based on local criteria and/or strategies. A combined solution can be accomplished via a system of deadlines. Agents would be allowed to decide when to initiate a process, but only up to a specific deadline that is enforced by the Monitor.

The Central Agency may be employed in other types of mediation. However, it is possible to further relax the purely distributed approach, and to enforce certain interaction structures via this Central Agency. Thus far we considered static power models – wherein agents are assigned unchangeable decision powers – and implicit or hard-coded representations of power – i.e., as a mode of operation. However, we are now in the process of devising models in which power is an explicit attribute that can be reasoned about. We aim to devise mechanisms whereby power can be negotiated between agents. The general solution in DC is readily extendable towards the inclusion of Central Agency with specialist roles. Central Agencies may be regarded as fulfilling specific roles, such as “emergency expert”, which would intervene in the allocation process only if certain emergency situations arouse. This perspective extends the traditional classification of central agents (within a federation of agents architecture) in facilitators, brokers and mediator.

The novelty of the solution as demonstrated by CC and DC resides in the fact that it is a natural and versatile formulation that combines an agent-based model with rule-based expressions of allocation strategies and multi-criteria optimisation expressions of allocation objectives. Furthermore, the solution provides for natural extensions towards different aspects of dynamic allocation, such as dealing with emergencies. ASMCR has been tested thoroughly using randomly generated problems.

This choice has led to a more general testing model than the real application. The usefulness of hill-climbing has been confirmed. Experiments also confirmed that by changing the weights to the multiple objectives, the manager can reduce the number of jobs not served, travelling distance and preference (lower value in preference means better service quality). ASMCR has demonstrated that it has real potential to be developed into practical solutions to BT’s workforce planning problem. ASMCR took 5 to 15 minutes to complete when tested on real-sized problems. We believe that it has potential to be developed into practical solutions for BT’s workforce planning problem.

8 Conclusions and Future Work

The Aberdeen Group note that winning service organisations are optimising their field service operations by collaborating with their customers and leveraging field service automation technologies to improve operational efficiencies, profitability, and customer satisfaction levels [Vigoro, 04b]. We have presented a novel resource planning approach using a common power model which ultimately aims to fulfill this objective. Although they are presented as disjoint prototypes and applications, they are sharing a common power model and they follow an incremental development path.

One major aspect to be investigated is the use of various market mechanisms and auctions to the resource allocation problem domain. Currently we are looking at

possible scenarios and the feasibility of auctions in this space. One general problem with the proposed solution is derived from the fact that local optimisation often conflicts with global optimisation. Therefore we have devised various reward systems in order to be able to arbitrate between collaborative (global) and competitive (local) objectives of domains and regions. We will be investigating the role of auctions in this respect where agents will have an incentive to give up for example their surplus resources for domains in order to improve the global situation of resource allocation. Therefore we will be looking at direct incentives for agents to release resources from within their domains, rather than keeping hold of them.

Ultimately our endeavour is to integrate all these implementations within a versatile service exchange system which could be fully customized to serve our internal resource market as well as external contractor based exchanges. We would like to create a resource exchange platform where one will have the option to select the various features presented in the component implementations like: dynamic power models with various degree of distributed/central control, strategies, criteria or dynamic constraint satisfaction in the expression of local interests and priorities, various synchronous and asynchronous communication protocols including the retractable contract net protocol RECONNET, and the choice of multi-objective optimisation algorithms as described in ASMCR and CC.

References

- [Aarts and Lenstra, 97] Aarts, E.H.L. and Lenstra, J.K. (eds): *Local Search in Combinatorial Optimisation*, Wiley, New York, 1997.
- [Andersson, 00b] Andersson J.: *A Survey of Multiobjective Optimization in Engineering Design*, Technical report LiTH-IKP-R-1097, Department of Mechanical Engineering, Linköping University, Linköping, Sweden, 2000.
- [Buchanan, 95a] Buchanan, J.T.: *Mathematical Programming and Resource Management*, BT Technology Journal, Volume 13 Number 1/1995, 29-36, 1995.
- [Coello] Coello, C.A.: *List of References on Evolutionary Multi-Objective Optimisation*, <http://www.lania.mx/~ccoello/EMOO/EMOObib.html>.
- [Ehrgott and Gandibleux, 02] Ehrgott, M. and Gandibleux, X. (eds): *Multiple Criteria Optimisation: State of the Art Annotated Bibliographic Survey*, Kluwer's International Series in Operations Research and Management Science, Vol. 52, Kluwer Academic Publishers, Boston, 2002.
- [Fox, Barbuceanu and Teigen, 00c] Fox M.S., Barbuceanu, M, and Teigen R: *Agent-Oriented Supply Chain Management*, *The International Journal of Flexible Manufacturing*, Volume 12/2000, 165-188, 2000.
- [He, Jennings and Leung, 03a] He M., Jennings N.R. and Leung H.: *On agent-mediated electronic commerce*, *IEEE Transactions on Knowledge and Data Engineering*, Volume. 15, Number 4/2003, 985-1003, 2003.
- [Kurbel and Loutchko, 01b] Kurbel K and Loutchko L.: *A framework for multi-agent electronic marketplaces: analysis and classification of existing systems*, In *Proc. of the International ICSC Congress on Information Science Innovations (ISI'2001)*, American University in Dubai, U.A.E., 2001.

- [Le Pape, 94] Le Pape, C.: Implementation of resource constraints in ILOG SCHEDULE: A library for the development of constraint-based scheduling systems, *Intelligent Systems Engineering*, Volume 3, Number 2/1994, 55-66, 1994.
- [Lesaint, Voudouris and Azarmi, 00a] Lesaint, D., Voudouris, C. and Azarmi, N.: Dynamic Workforce Scheduling for British Telecommunications plc., *Interfaces*, Volume 30, Number 1/2000, 45-56, 2000.
- [Owusu, Voudouris, Kern, Garyfalos, Anim-Ansah and Virginas, 06b] Owusu, G., Voudouris, C., Kern, M., Garyfalos, A., Anim-Ansah, G., Virginas, B.: On Optimising Resource Planning in BT with FOS, *IEEE International Conference on Services Systems and Services Management, IEEE/SSSM06*, 2006.
- [Reeves, 93a] Reeves, C.R. (ed): *Modern Heuristic Techniques for Combinatorial Problems*, Blackwell Scientific Publications, Oxford, 1993.
- [Shen and Norrie, 99a] Shen, W. and Norrie, D.H.: Agent-Based Systems for Intelligent Manufacturing: A State of the Art Survey. In: *International Journal of Knowledge and Information Systems*, 1:2 (1999) 129-156, 1999.
- [Tsang, 93b] Tsang, E.P.K.: *Foundations of Constraint Satisfaction*, Academic Press 1993,
- [Tsang, Gosling, Virginas, Voudouris and Owusu, 05] Tsang, E.P.K, Gosling T., Virginas, B., Voudouris C. and Owusu G.: Retractable contract nets for distributed workforce scheduling, In *Proc. of the 2nd Multidisciplinary Conference on Scheduling: Theory and Applications (MISTA 2005)*, New York, USA, 18-21 July, 2005.
- [Ursu, Virginas and Voudouris, 04a] Ursu, M., Virginas, B and Voudouris, C: Distributed Resource Allocation via Local Choices: General Model and a Basic Solution, *Knowledge-Based Intelligent Information and Engineering Systems*, 8th International Conference, KES 2004, Wellington, New Zealand, September 20-25, 764-771, 2004.
- [Vigoroso, 04b] Vigoroso, M.: *The Field Service Optimization Benchmark Report: Tapping the Service Supply Chain for Profit and Competitive Advantage*, Aberdeen Group Report, June, 2004.
- [Virginas, Owusu, Voudouris and Anim-Ansah, 03c] Virginas, B., Owusu, G., Voudouris, C., and Anim-Ansah, G: A two stage optimisation platform for resource management in BT. In *Proc. of the Twenty third Annual International Conference of the British Computer Society's Specialist Group on AI (2003)*, 109-121, 2003.
- [Voudouris, Owusu, Dorne and McCormick, 06a] Voudouris, C., Owusu, G., Dorne, R., and McCormick A.: FOS: An Advanced Planning and Scheduling Suite for Service Operations, In *Proc. IEEE International Conference on Services Systems and Services Management, IEEE/SSSM06*, 2006.
- [Voudouris, Dorne, Lesaint and Liret, 01a] Voudouris, C., Dorne, R., Lesaint, D., and Liret, A.: iOpt: A Software Toolkit for Heuristic Search Methods, Principles and Practice of Constraint Programming, In *Lecture Notes in Computer Science*, Vol. 2239, 716-729, CP, 2001.
- [Voudouris, Owusu, Dorne, Ladde and Virginas, 03b] Voudouris, C., Owusu, G., Dorne, R., Ladde, C and Virginas, B.: ARMS: An Automated Resource Management System for British Telecommunication plc, In: *EURO/INFORMS Joint International Meeting EURO Excellence in Practice Award*, Istanbul, 2003.
- [Walsh and Wellman, 99b] Walsh W. and Wellman M: Modelling Supply Chains Formation in Multiagent Systems, In *Proc. IJCAI-99 Workshop Agent-Mediated Electronic Commerce*, 1999.

[Wooldridge and Jennings, 95b] Wooldridge M., and Jennings N.R.: IntelligentAgents: Theory and Practice, The Knowledge Engineering Review, Volume. 10, Number 2/1995, pp. 115-152., 1995.