

# Dynamic Workforce Scheduling for British Telecommunications plc

---

David Lesaint

*Intelligent Systems Research  
BT France  
11, place des Vosges  
92061 Paris La Defense, France*

Christos Voudouris

*Intelligent Systems Research  
BT Adastral Park  
MLB1-PP12  
Ipswich IP5 3RE, United Kingdom*

Nader Azarmi

*Intelligent Systems Research  
BT Adastral Park*

---

British Telecommunications plc (BT) employs thousands of field engineers across the UK to maintain networks, repair faults, and provide service to customers. To allocate work efficiently, BT developed Work Manager, an information system that automates work management and field communications. In 1996, the Intelligent Systems Research group of BT enhanced Work Manager with a dynamic scheduler (DS) based on a combination of heuristic search and constraint-based reasoning. Rolled out in 1997 and reaching 20,000 engineers in 1998, DS with Work Manager is saving BT \$150 million a year on operational costs. When deployed over the targeted workforce of 40,000 people, the system will save an estimated \$250 million a year.

Simply stated, field-workforce scheduling is about sending the right engineer to the right customer at the right place at the right time with the right equipment—at any time and in any operational environment. For BT, which employs over 50,000 field engineers, workforce scheduling is critical, and its ability

to provide high quality service while achieving maximum productivity and low operational costs is vital to the company's success and competitiveness.

Many factors contribute to the complexity of the problem. First, skill requirements vary immensely: putting in a switch on business premises may take several engi-

neers several days while recovering redundant equipment may take an unskilled person a few minutes. Three basic operational areas have distinct skill demands: business and residential customer access, national business communications, and core network. Each area is the responsibility of a separate BT division.

The second factor contributing to the complexity of scheduling is the geographical distribution of the workforce. For example, the 15,000-person workforce in the access division consists of 175 separate groups that operate within nonoverlapping geographical domains. Each group manages its domain independently and controls from a dozen to a few hundred engineers and each group may perform a few hundred to a few thousand tasks per day. This geographical and operational decomposition allows the management of smaller-sized workforces than would be the case if the entire division's workforce were managed centrally.

Complexity also lies in the multiple requirements and objectives of workforce scheduling. These requirements reflect standard allocation constraints and corporate rules, and they determine the feasibility and acceptability of work assignments. Some constraints recur across domains and divisions:

- each engineer operates primarily within a predefined area;
- each engineer has off-hours and predefined breaks during the day (for example, lunch, time for personal appointments) that must be inserted into his or her schedule;
- engineers can perform only one activity at a time;

- the last task of the day may require overtime, if allowed;
- some tasks must be performed within an agreed-upon time window;
- some tasks must be started or completed by an agreed-upon time;
- engineers must be matched to tasks;
- access to customer premises may be granted only to certain individuals at certain times;
- task execution may be broken according to rules that take into account task and break details;
- the duration of a task depends on the engineer's experience and skills, and the duration of a journey depends on its start time;
- some tasks must be sequenced in time, for example, collecting equipment before making a visit; and
- some tasks must be performed in parallel by different engineers in different places at the same time.

(Geographical) domain-dependent requirements are rare. In fact, domains differ only on nonessential features, such as their road networks, their distribution of customers, the size and skills of their workforces, or their task-notification patterns.

Scheduling decisions within domains are driven by the same set of corporate objectives and customer preferences:

- maximize the productivity of the workforce;
- improve service quality, measured according to the type of customer (residential, business), the priority of the work (provision, installation, maintenance, repair), and any service-level agreement (on appointment windows and due dates);
- make best utilization of skills, for exam-

ple, assign the most suitable engineer to a task;

—minimize the operational costs resulting from travel time, waiting time, and work in overtime; and

—satisfy work controllers' and engineers' preferences as to location and types of assignments, for example, avoid assigning low-level work to specialized individuals.

Constructing feasible and good-quality work schedules under these conditions is hard! In fact, the problem may be viewed as a complex variant of the vehicle-routing problem featuring multiple vehicles, multiple depots, various compatibility constraints, time constraints, operational constraints, synchronisation constraints, and conflicting quality objectives.

The problem is further compounded by the inherent instability of the environment. Indeed, much schedule information is uncertain, imprecise, and incomplete:

—BT and its customers may request, cancel, or amend jobs unpredictably;

—engineer availability is subject to last-minute changes, and estimates of task duration and travel time cannot be totally accurate;

—work controllers may modify provisional work assignments and review business objectives at any time; and

—the environment itself (weather, traffic conditions) is unpredictable.

Therefore, any scheduling process must continuously incorporate new data to generate valid work assignments but it must also minimize the impact of these data on the current work schedule. Indeed, anyone interacting with the work-allocation system must be provided with schedule information that is as stable as possible over

time. For instance, a resource manager who elaborates short-term plans based on the schedule information provided must be confident that this information will resist future changes.

### **Towards a Unified Work-Management Process**

To improve and unify its work-allocation operations, BT decided in 1989 to automate its work-management and field-communication processes. In 1993, it introduced an information system called Work Manager [Garwood 1997]. Work Manager provides (Figure 1):

—an interface to other BT systems that feeds it with work;

—task-dispatch and task-progress-and-closure functions;

—a task-jeopardy management system to ensure that work controllers are alerted to potential failure;

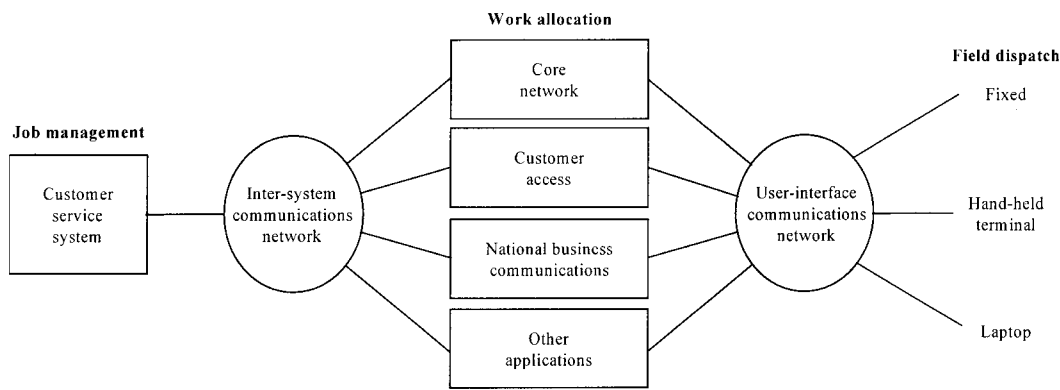
—feedback services to keep the system that originated a task informed of its progress;

—management of sequences and lineups;

—an interface to inventory management systems; and

—a real-time and historic management information system.

In developing Work Manager, BT wanted to enable and benefit from automated scheduling. It initially equipped Work Manager with a real-time allocation algorithm (RTA), and each geographical domain ran an instance of this system. The RTA is an extension of the Hungarian algorithm that operates on an assignment-cost matrix to determine the minimal-cost assignment set [Laithwaite 1995]. Within Work Manager, the RTA ran every five minutes to identify and assign the next



**Figure 1: Work Manager is the system that manages the allocation of customer orders and their execution by field engineers in different operational environments (core network, customer access, and national business communications). Field engineers communicate with the system through various interfaces (laptops, hand-held terminals, and so forth).**

task to each engineer. Rolled out in 1994, Work Manager with the RTA made tangible improvements in productivity and quality of service.

RTA originally complemented an overnight batch-scheduling process called TourBuild which scheduled the more predictable and stable provision work. However, business needs were changing rapidly, and TourBuild could not support a more integrated workforce with dynamic repair and provision. BT then successfully enhanced RTA to automatically allocate engineers both types of work, separately and as a mixed workload. However, it was clearly not optimal for situations that called for many scarce skills, for long tasks (a day or more), or for tasks with interdependencies—all of which required greater look-ahead to schedule and allocate efficiently. In these situations, work controllers found it difficult to minimize ineffective time and balance work across the workforce because of RTA’s short-term view, and they had to intervene manually to get optimal results. Indeed, work con-

trollers can still schedule tasks manually if they are not satisfied with the assignments the RTA proposes.

BT needed a scheduler that would be as good as the RTA for very dynamic workloads with short response times but much better for more complex work and mixed workloads requiring greater look-ahead. In April 1995, BT decided to explore alternative scheduling algorithms that could replace the RTA within Work Manager and retain the benefits and savings already achieved. BT UK customer service division set up a cross-divisional team technically led by the intelligent systems research group of BT advanced communications engineering division to design and develop a new dynamic scheduling system called Dynamic Scheduler (DS).

**Dynamic Scheduling**

BT had the following objectives for Dynamic Scheduler: to extend the scope of Work Manager to all work and people, to maintain or improve quality of service, to reduce travel time, to improve field and control productivity, to reduce control in-

terventions, to improve management control through data visualization, and to improve resource planning.

The computational problem DS faces is to generate feasible and high-quality schedules and to adapt them over time to accommodate the dynamics of reality. Based on our experience with the RTA and our analysis of the problem, we built DS based on two principles:—coupling loosely an on-line allocator and a predictive scheduler to preserve responsiveness while benefiting from global optimization (Figure 2); and—using a uniform constraint optimization approach to provide a generic and efficient underlying computational model.

Because of the complexity of the feasibility problem, we decided to exploit the most appropriate technology for handling constraints, constraint programming [Tsang 1993]. This could ensure generating realistic schedules using optimal constraint-propagation techniques and encountering no limitations in scope if problem specifications evolve (as they did for the RTA).

However, a fully constraint-based scheduler is not appropriate for the optimization dimension of the problem. The ideal approach would combine proven optimization techniques with constraint-based reasoning to tackle all aspects of the problem comprehensively and efficiently. We chose heuristic search because of its track record in vehicle routing [Laporte and Osman 1995]. After running comparative tests on simulated annealing, tabu search, and genetic algorithms over large collections of simplified workforce scheduling problem instances [Brind, Muller,

and Prosser 1995], we elected simulated annealing [Kirkpatrick, Gelatt, and Vecchi 1983] as the best algorithm based on efficiency, simplicity, and robustness criteria.

The challenge in integrating simulated annealing and constraint-based reasoning is to preserve the former's performance without losing the exactness and extensibility of the schedule model as implemented by the constraint-based representation. For this reason, we developed a systematic algorithm to preschedule the most-constrained tasks: the long-duration tasks (typically, over eight hours) and the interdependent tasks (tasks that must be synchronised to satisfy a complex work request). This algorithm performs a tree search and is guided by domain-dependent heuristics defined over the choice of tasks, engineers, and tour positions for tasks.

After the prescheduling stage, we use simulated annealing to schedule the other tasks (the independent or short-duration tasks) without any performance deterioration. Both algorithms operate with the same schedule horizon, which is typically two days in advance. Simulated annealing does most of the work since most tasks are independent (more than 90 percent in the access domain, more than 80 percent in the core domain). The output of the two-stage scheduling is a provisional schedule defining a sequence of tasks for each engineer. This schedule complies with all the constraints specified and minimizes a cost function reflecting the various objectives of the problem (we describe the algorithms in detail in the Appendix and in Lesaint et al. [1998]).

The tree-search and simulated-annealing

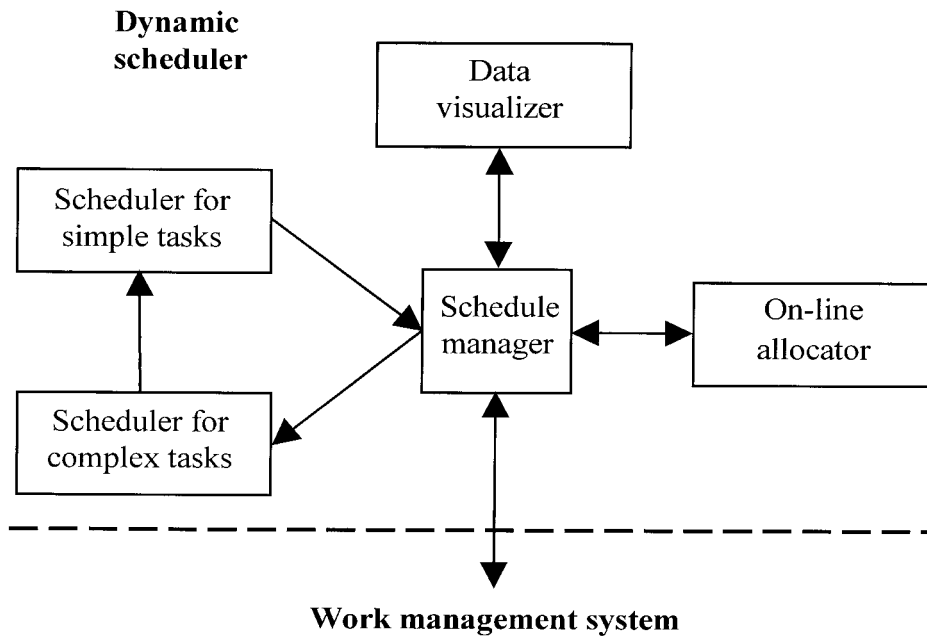


Figure 2: The dynamic scheduler consists of a schedule manager that stores and communicates the provisional schedule to the work-management system, a scheduler for complex tasks and a scheduler for simple tasks that periodically produce the long-term schedule, an on-line allocator that adjusts the schedule just before dispatch and a schedule and data visualizer (a desktop computer) that field-resource controllers and managers use to access problem and schedule information.

algorithms form the predictive component of DS. They are run every five to 15 minutes to generate the provisional schedule, which is then fed to the on-line allocator. We adopted this activation strategy to match the dynamic nature of the problem.

The on-line allocator dispatches work assignments to engineers: it is the ultimate decision maker (apart from work controllers, who can always interfere). It is triggered by external events, for example, tasks requested by on-line engineers and tasks that require immediate work before they become failed business targets. By default, the allocator simply forwards the assignments prescribed in the provisional schedule to the engineers, but it may modify these assignments if unforeseen events

have caused infeasibility or suboptimality, for example, late notification of a high-priority task or early arrival of an engineer. The allocator’s decision making is rule based, and its scope is limited to a subset of the resources and tasks so it can deliver answers within a few seconds.

Provisional assignments often prove inappropriate at dispatch time and the allocator must interfere. This is unavoidable given the uncertainty and the pace of change in the problem. Ideally, one would tackle instability by capturing and reasoning about elements of change using such techniques as stochastic programming. However, it is very difficult to forecast the evolution of BT’s workforce-scheduling environments accurately. Although we

have identified general patterns (for example, on customer calls), no forecasting model is accurate enough to be exploited during scheduling [Simpson et al. 1995].

DS is equipped with an interface that allows one to view provisional schedules and any information generated by the predictive scheduler. Both temporal and spatial dimensions of the schedules are displayed through a Gantt chart and a geographical map. Pie charts, bar charts, and graphs provide additional statistics and measures on algorithms and solutions. This interface can also be used to implement what-if scenarios by running the predictive algorithms off-line on selected scenarios and with different settings.

### **Broad Impact**

BT first deployed DS in November 1996 in South London. After two successful months, it decided on a national rollout starting in 1997. In May 1998, the coverage reached 20,000 of the 28,000 engineers under Work Manager's control. DS has helped BT in its aim of repairing "today's faults today" and achieving high customer satisfaction.

After two years of exploitation, the results have confirmed the superiority of the approach over the short-term scheduling strategy of the RTA. First, the use of DS in Work Manager annually saves \$150 million on engineers' and controllers' pay costs and other workforce related costs, such as vehicles, equipment, tools, training, and administration. This is a 10-percent increase in savings over Work Manager used with the RTA. The savings realized with the DS are based on a four-percent productivity increase nationwide,

which comes from reductions in travel time and waiting time and better allocation of work.

In addition, DS has paved the way for the automated management of another 12,000 engineers now managed outside Work Manager because their work was not suitable for RTA's allocation. BT expects a 15-percent productivity increase for this workforce, saving BT an additional \$100 million a year. The savings for the full coverage of 40,000 engineers should reach \$250 million a year. The expected 15-percent increase is made up of an 11-percent productivity increase achieved by the introduction of the Work Manager and an additional four-percent increase resulting from DS deployment.

DS has also improved field-control operations. Because it provides a forward view (of up to five days ahead), work controllers can be more proactive in managing task jeopardy and in planning short-term resources. Using the visualizer and its what-if capability, they can identify and resolve resource bottlenecks and potential task failures in advance. The accuracy and quality of the schedule information the system provides has given work controllers confidence in its performance. As a result, they have reduced their manual interventions dramatically: for example, the percentage of manual allocation dropped by 40 percent during the first field trial.

DS has helped BT to meet due dates, appointment windows, and skill requirements more consistently than it used to. Furthermore, field engineers' work schedules have improved. They travel less; their absences (for lunch, appointments, sick-

ness, and so forth) are smoothly inserted within their work schedules; work interruptions tend to be less frequent; and legal rules agreed between BT and unions on the management of overtime or travel in own time are rigorously implemented.

DS has also enhanced work controllers' interactions with Work Manager. The interface for configuring the predictive scheduler is much simpler than the RTA's. The user has to select a few heuristics to "drive" the prescheduler, for example, to order jobs based on priority, urgency, or skill levels. Regarding the simulated-annealing algorithm, the user can configure the cost function through intuitive parameters, for example, a penalty for minutes spent on travel or a weight to balance the importance of one task category over another. To operate the visualizer itself, users need basic knowledge of the window-environment. They take national training courses to familiarize themselves with the system and to learn how to drive it. However, because each domain is unique, DS needs extensive tuning before going live, and work controllers also have to get accustomed to the peculiarities of these domain-dependent adjustments.

Because of the generic nature of its constraint-based schedule model, DS requires little reengineering for use in scheduling within BT. In fact, DS can facilitate the removal of geographical and operational boundaries (domains and divisions) that divide BT's work-management operations. BT is investigating this simplifica-

tion of its work-management organization. DS (as any other optimization software) is limited only by the size of problems in relation to the expected response-time levels. Overall, DS has:

- reduced BT's operational costs;
- improved the integrity, quality, and consistency of work allocation for the whole workforce;
- improved customer satisfaction;
- enhanced resource control;
- improved field engineers' and work controllers' perception of Work Manager;
- and
- offered BT the potential of simplifying its work-management organization.

From a purely technical viewpoint, DS demonstrated that integrating heuristic-search and constraint programming is a successful approach for tackling large-scale industrial vehicle routing problems. In fact, we believe that such hybridization is the key to solving a wide range of hard combinatorial problems.

#### **Acknowledgments**

We acknowledge our colleagues in BT UK Customer Service and Advanced Communications Engineering Divisions: Ian Alleston, Bob Laithwaite, Andy Noble, John O'Donoghue, and Paul Walker.

---

Executive summaries of Edelman award papers are presented here. The complete article was published in the INFORMS journal *Interfaces* [2000, 30:1, 45-56]. Full text is available by subscription at <http://www.extenza-eps.com/extenza/contentviewing/viewJournal.do?journalId=5>