

FieldPlan: Tactical Field Force Planning in BT

Mathias Kern, George Anim-Ansah, Gilbert Owusu, and Chris Voudouris

Intelligent Systems Research Centre
BT Group CTO Office
Adastral Park, Martlesham Heath
Ipswich IP5 3RE, UK

{mathias.kern,george.anim-ansah,gilbert.owusu,chris.voudouris}@bt.com

Abstract. In a highly competitive market, BT¹ faces tough challenges as a service provider for telecommunication solutions. A proactive approach to the management of its resources is absolutely mandatory for its success. In this paper, an AI-based planning system for the management of parts of BT's field force is presented. FieldPlan provides resource managers with full visibility of supply and demand, offers extensive what-if analysis capabilities and thus supports an effective decision making process.

1 Introduction

BT is a leading provider of telecommunication solutions servicing customers in the UK and throughout the world. Like any other service organisation, BT is faced with the stern challenge of delivering services optimally to its customers. The effective management of resources is a fundamental and critical part of this challenge. A proactive resource management approach, i.e. an approach that provides full visibility of the service chain, that offers extensive analysis capabilities, that is automated and user-friendly, is required. [1, 2] are examples of the considerable research and development effort put into the automation of resource management.

BT's Intelligent Systems Research Center has developed a fully integrated suite of applications for the management of field force resources [3, 4]. Among other components, this suite includes:

- *FieldForecast* allows the forecasting of expected demand volumes.
- *FieldPlan* is an application for resource planning.
- *FieldSchedule* is a resource scheduling system.
- *FieldExchange* facilitates the redistribution and exchange of resources.
- *FieldPeople* is a tool for gathering and managing all people-related information such as availability and working patterns.
- *FieldReserve* is a reservation system for incoming work requests.

¹ British Telecommunications plc

The main focus of this paper is FieldPlan. FieldPlan is a planning system for the management of field force resources such as engineers and technicians. It incorporates a planning algorithm based on heuristic search that efficiently and effectively aligns the expected demand for services with the available field force supply. While doing so, resource managers can actively engage in scenario modeling and thus can analyse the problem from different perspectives in order to reach better informed decisions.

The paper is structured as follows. In section 2, the scenario for field force planning in BT is described. Section 3 is a detailed account of the FieldPlan planning approach. Initial results are presented in section 4, followed by concluding remarks in section 5.

2 The BT Field Force Planning Scenario

With regard to resource planning, one has to distinguish between three main types: scheduling, tactical planning and strategic planning. Scheduling is resource management very close to the actual time of service provision, and may look a few hours up to one or two days ahead. Tactical resource planning is short to mid-term planning as it deals with time windows of a few days up to several weeks. If planning is performed for longer time phases, i.e. up to several years, this is then viewed as a strategic planning task.

The resource planning scenario to address here is a tactical one. A workforce of field engineers, the resources, have to be optimally deployed to serve expected demand in form of jobs. This means that plans have to be generated for the field force, on a per day or per week basis, for up to 12 weeks. FieldPlan has to deploy up to 300 engineers to fulfil up to 20,000 job requests in a single planning run. The main planning objective is to optimally utilise the field engineers to complete jobs while reducing operational costs. Jobs and engineers are characterised by the attributes listed below:

Job	Engineer
geographical location	list of geographical locations, with preferences
skill	list of skills, with preferences
product	list of products, with preferences
time window	availability over time
type (actual or forecasted)	type (standard, loan-in, contractor, etc.)

Each job is either an actual job, i.e. a real customer demand, or is forecasted by the FieldForecast system. A job involves the provision, repair or cease of a product at a particular location and requires a specific skill. This should be accomplished within a given time frame. The importance of a job indicates the value of the respective work to the business.

Engineers, on the other hand, possess sets of skills, can provide various products and are flexible in their choice of working area. Preferences for particular areas, skills and products² are possible. Different types of engineers can be deployed, e.g. standard BT engineers or contractors. The engineers show varying daily availabilities, in terms of working hours, during the planning time window.

General importance and productivity rates of skills and products are also known.

In contrast to scheduling systems [5] which assign engineers to particular jobs, the goal here is to determine the best arrangement of engineers in terms of areas, skills and products. For every time period within the planning horizon, each engineer has to be assigned to an area or a set of areas, a skill or a set of skills, and a product or a set of products. Furthermore, the respective numbers of engineer working hours have to be established. The decision about what jobs engineers will actually work on is left to a scheduler at a later point.

Results of the planning process are presented in two forms: capacity and deployment plans. A capacity plan is a coarse-grained summary of the service situation purely in volumes: how many jobs are expected to be cleared/uncleared per area, skill and product; how many engineers are available and utilised per area, skill and product; in short, can the demand be met with the available supply. The deployment plan refines this information by giving explicit area, skill and product deployment recommendations together with expected utilisation percentages for each engineer.

3 The FieldPlan System

In this section, the planning algorithm and its main components are described and important issues are discussed. A schematic outline of the planning approach is given in figure 1.

FieldPlan Algorithm

```

FOR all planning time periods:
  1. Aggregate jobs
  2. Generate baseline plan
  3. Optimise baseline plan
  4. Decompose job aggregates
  5. Generate output
Return all outputs

```

Fig. 1. The FieldPlan planning algorithm

² In the current version of FieldPlan in BT, engineers are deployed by default to all products.

Planning Time Window and Periods The process of resource planning is performed for a given time frame. This time window is usually further divided into shorter intervals, and plans have to be generated for all individual stages. A typical tactical example is the planning for two weeks on a per day basis.

The FieldPlan system captures this notion by introducing planning time periods. A planning time period is defined by a start and end time and thus represents a time interval. The overall planning time window is, consequently, a sequence of non-overlapping time periods. By basing the core planning mechanism in FieldPlan purely on the generic concept of time periods, we are able to express the algorithm without any reference to actual time intervals like days or weeks.

A central characteristic of our planning algorithm is the iterative construction of the overall plan out of partial plans. For each individual time period, a partial solution is constructed. Rather than constructing a single plan for the whole time window in one step, a number of steps yields a number of partial solutions which are finally combined to form the complete plan.

Initial Planning Time Period The first planning time period represents the current time period, i.e. it contains the current time at the point of planning. It differs therefore from all subsequent time intervals in two aspects:

- At the point of planning, engineers might already be working. Consequently, they are already assigned to specific areas, skills and products for the first time period. These restricted choices have to be considered when evaluating the future course of the first period.
- Furthermore, parts of the initial time period might have already past at the point of planning. This must be reflected by reduced working hours of the engineers. If, for instance, 50% of time period 1 have already past then the available engineer capacities must be reduced by half.

Job Aggregation and Decomposition In contrast to scheduling systems, FieldPlan does not assign engineers to work on particular jobs. The aim of tactical planning is of a more general nature, namely to decide what kind of engineers have to be deployed to which location to fulfil what kind of jobs. This generalisation allows the planning process to work with job categories rather than single jobs. Hence our planning system merges single jobs into more general objects referred to as job aggregates and bases all planning decisions purely on these aggregations. The advantage of this approach is clear: reduced complexity while maintaining a sufficient level of granularity. Planning becomes easier, faster and more scalable.

In our scenario, a job aggregate is described by the following information:

Job Aggregate
geographical location
skill
product
state
volume

This means that jobs that match in area, skill, product and state are aggregated. Instead of a time window, job aggregates possess a state: backlog, current or future. Backlog job aggregates should have been finished before the current planning time period, current aggregates should be completed within the current time period, and future job aggregates have a later target completion. The volume indicates the number of aggregated jobs.

In the main planning algorithm, all jobs are initially aggregated during each time period. After generating the baseline plan and optimising it, the cleared and uncleared job aggregates are decomposed back into cleared and uncleared jobs.

Objective Function The objective function is the central instrument of evaluating the quality of a particular resource deployment. By assigning a numeric value to each such deployment, comparing candidate solutions and thus selecting the better one is made possible. The compositions of the FieldPlan objective is illustrated below.

For a given job aggregate $\mathcal{A} = (area, skill, product, state, volume)$, the *job aggregate clearance score* $jacs(\mathcal{A})$ is defined as

$$jacs(\mathcal{A}) = importance(\mathcal{A}) \cdot \frac{volume}{productivity(skill, product)}. \quad (1)$$

It is a measure for the value of clearing the job aggregate. The formula normalises the volume with regard to the productivity: the longer work takes, the higher is the score. The score is also better if the job aggregate is more important. In our current implementation, $importance(\mathcal{A})$ is calculated as $importance(skill) + importance(state)$, but more complex measures that consider area and skill as well are possible.

For an engineer \mathcal{E} with assigned job aggregates $\mathcal{A}_1^{\mathcal{E}}, \dots, \mathcal{A}_n^{\mathcal{E}}$, the *engineer clearance score* $ecs(\mathcal{E})$ is given as

$$ecs(\mathcal{E}) = \sum_{i=1}^n jacs(\mathcal{A}_i^{\mathcal{E}}). \quad (2)$$

While the engineer clearance score only evaluates the (quality of) service delivered by an engineer, the *engineer score* $es(\mathcal{E})$ takes also service costs into account as it considers the actual deployment of the engineer:

$$\begin{aligned}
es(\mathcal{E}) &= ecs(\mathcal{E}) \cdot \\
&\quad area_penalty(\mathcal{E}) \cdot skill_penalty(\mathcal{E}) \cdot product_penalty(\mathcal{E}) \cdot \\
&\quad utilisation_penalty(\mathcal{E}) \cdot technician_type_penalty(\mathcal{E}).
\end{aligned} \tag{3}$$

Above penalty functions return values between 0 and 1. The original engineer clearance score is reduced by this formula if the engineer is not assigned to the preferred area, skill or product. Further penalties might be applied in case of poor utilisation or if the engineer has to be brought in as additional workforce like contractors or on overtime.

The overall *objective score* obj of a deployment of the engineers $\mathcal{E}_1, \dots, \mathcal{E}_n$ can be calculated as

$$obj(\mathcal{E}_1, \dots, \mathcal{E}_n) = \sum_{i=1}^n es(\mathcal{E}_i). \tag{4}$$

This objective is a balanced measure of the quality of service delivered by all engineers and the incurred service costs.

Baseline Plan Generation The initial baseline plan is generated by assigning all engineers to their default choices, i.e. to their preferred areas, skills and products. The least flexible engineers in terms of available areas, skills and products are deployed first to ensure that even highly constrained engineers are able to pick up work. Assigning engineers to such choices involves the assignment of matching uncleared job aggregates according to their capacities.

Plan Optimisation The strategy employed in FieldPlan to optimise an initial baseline plan is a hill-climbing algorithm. The current plan is iteratively modified in small steps which are referred to as moves. Each new modification of the solution is characterised by an improved, i.e. higher, objective function evaluation. Three types of moves are considered by FieldPlan:

- *Move-to moves*: Deployments of single engineers are altered by assigning them to new areas, skills and/or products.
- *Swap moves*: The deployments of two engineers are exchanged.
- *Replacement moves*: This move type involves a chain of two or more engineers. The first engineer replaces the second engineer, the second engineer replaces the third, and so on. The final engineer is thus freed and can be moved elsewhere. Currently, only chains of length two and three are considered.

The optimisation routine repeatedly applies the three types of moves to the engineers. As long as improvements are found, the process continues. Only when none of the considered moves offer any improvements anymore, the optimisation is halted as a local maximum is reached.

Heuristic search algorithms such as Simulated Annealing [6], Tabu Search [7], Genetic Algorithms [8] or Guided Local Search[9] employ strategies to escape such local extrema and to improve the generated solution even further. Although these advanced search methods can be easily incorporated in FieldPlan using the defined objective function and move operators, we use the more basic

hill-climber due to computation time restrictions. As interactivity is application critical, fast computation and response times are of uttermost importance. The less time intensive hill-climbing approach can meet these requirements and is thus an appropriate choice.

Output At the end of each optimisation cycle, the compiled planning solution is analysed by FieldPlan to provide resource managers with vital information on the state of the supply chain. In addition, FieldPlan computes data that can be utilised as inputs by further systems. The following list gives an overview of the different categories of analysis:

- *Capacity plan:* On the one hand, this plan analysis provides data about the available capacity in the form of engineers. On the other hand, it lists required capacity in form of jobs. Resource managers are able to see how well both match, how many cleared and uncleared jobs can be expected and how well the engineers will be utilised.
- *Deployment plan:* This output makes explicit area, skill and product recommendations for each engineer for all time periods. This deployment information can consequently be used as direct input to a scheduling system.
- *Under-utilisation:* FieldPlan identifies engineers who are not fully utilised. If the utilisation level fall below a certain threshold, reasons for the poor efficiency are compiled and recommendations for improving the utilisation are given.
- *Training:* By analysing the demand for each skill, bottlenecks can be identified and according training recommendations can be given.
- *Resource redistribution:* A resource redistribution system like FieldExchange aims to balance resources across a wider geographical area. Resource shortages and surpluses indicated by FieldPlan can be used as a data feed.
- *Appointment booking:* The analysis of available engineer capacity can be used by appointment booking systems.

Plan Types and Planning Options FieldPlan distinguishes three levels of plan optimisation: baseline plans, optimised plans and user-specified plans. The main differentiator is the level of flexibility the planning mechanism is given in deploying the resources. On one side of the scale, baseline plans do not offer any flexibility at all. All technicians are restricted to their default areas, skills and products, and no overtime is available. On the other side of the scale, optimised plans do not restrict technician deployments in any way, e.g. technicians can be moved to any of their areas, skills and products, and allow the full utilisation of all available overtime. User-specified plans enable resource managers to vary the level of flexibility between the two extrema via a set of planning controls. The most important parameters for such what-if analysis in FieldPlan are:

- The amount to penalise engineers, in terms of the engineer score, for not being deployed to their default areas, skills and products can be varied.
- Productivity rates can be set by the user.
- The importance of skills can be varied.

- The importance of job aggregate states can be varied.
- The usage of overtime can be switched on or off.
- A minimum target utilisation can be specified.
- Jobs can be excluded from the planning by area, skill, product and type.
- Resource managers can choose to not deploy engineers to backlog, current or future work.

All these options are calendarisable, i.e. can be set independently for each planning time period. They provide resource managers with rich scenario modeling capabilities and thus support the decision making process.

4 Results

We present planning results for four real-life scenarios for a single planning time period and compare the performance of FieldPlan’s baseline plans (BL), plans produced by FieldPlan’s hill-climbing method (HC) and plans generated by standard Tabu Search³ (TS) and Simulated Annealing⁴ (SA) implementations which are based on the same neighbourhood structure and objective function as the hill-climber.

For each of the test cases, the following table shows the objective score for each plan type, together with the number of area/skill moves⁵, total number of cleared jobs and number of cleared backlog/current/future jobs. Backlog jobs had the highest importance in all four scenarios, followed by current jobs and future jobs. The number of deployed engineers ranged from 35 to 150.

We can observe that the baseline plans have the lowest objectives in each of the scenarios. The hill-climber produces improved objective scores, but the best results are always achieved by Tabu Search and by Simulated Annealing. The improvements from baseline plans to hill-climber plans are much more dramatic than the improvements from hill-climber plans to Tabu Search/Simulated Annealing plans.

While the baseline plans disallow any kind of area/skill moves, the other algorithms redeploy engineers in terms of area or skill. This increased flexibility results in better plans. The hill-climber, Tabu Search and Simulated Annealing produce plans which show higher total job clearance rates in all scenarios except case two. However, notably more high importance jobs (backlog) are cleared by those algorithms in that particular scenario.

Overall we can state that plan optimisation through hill-climbing, Tabu Search and Simulated Annealing leads to the clearance of more and more important jobs in our experimental study with the latter two approaches consistently

³ The Tabu Search algorithm employs aspiration, marks technicians who are involved in a move as taboo for $\frac{\text{number of engineers}}{5} \pm 2$ cycles and runs for 500 iterations.

⁴ The Simulated Annealing method uses an exponential cooling scheme and runs for 40,000 moves (accepted and refused).

⁵ Engineer deployments away from default choices.

Plan Type	Objective	# Moves	#Cleared Jobs	#Backlog	#Current	#Future
BL	27.91	0	227	88	139	0
HC	30.44	3	236	100	136	0
TS	31.07	4	245	101	144	0
SA	31.09	3	246	100	146	0
BL	112.40	0	776	405	222	149
HC	124.41	17	746	489	233	24
TS	125.65	23	746	525	204	17
SA	126.11	24	751	517	227	7
BL	98.85	0	781	351	244	186
HC	116.44	22	822	442	271	109
TS	117.57	28	815	464	266	85
SA	117.56	27	803	459	268	76
BL	107.55	0	793	354	323	116
HC	119.91	24	822	516	235	71
TS	122.02	33	841	568	194	79
SA	122.85	36	834	571	189	74

performing best. But these two algorithms come with higher computational costs. Both Tabu Search and Simulated Annealing take about 50 times longer than the hill-climbing method. In real-life scenarios, the improvements achieved by Tabu Search and Simulated Annealing are not significant enough to justify their much slower response times. Computation times of more than 5 minutes on typical server hardware for seven day plans are not acceptable for an interactive planning system. FieldPlan therefore employs the hill-climbing algorithm for plan optimisation to balance between performance and computation time.

After a successful trial period, the FieldPlan system is currently used nationwide (UK) for the deployment of 4,000 engineers of a BT sub-division. Reports indicate improvements in all their key performance measures, examples of which include:

- 8%-14% more engineers are assigned a job first thing in the morning, i.e. do not have to wait. The respective travel time has been reduced from 95min to 85min due to the better initial placement of engineers.
- Because of automation gains with FieldPlan, manual intervention by controllers has been limited to 18% (down from 31%).
- Quality of service has been improved, e.g. 1.1% more business provision jobs are completed on time.

The application of FieldPlan to a wider field force of 25,000 BT engineers is currently in preparation. Deployments are expected to start in February/March 2006.

5 Conclusions

The Aberdeen Group [10, 11] argues that service organisations have to optimise and automate the operations of their field services in order to improve efficiency, profitability and customer satisfaction. Without the constant drive for such improvements, survival in the highly competitive market places of the 21st century seems impossible. BT, as a provider of telecommunication services, faces these challenges every day.

In this paper, we have introduced FieldPlan, a tactical resource planning solution for the optimisation of parts of BT's field force. We have outlined the resource management scenario, have described the core planning mechanism and have discussed the extensive scenario modelling options. Initial results of the application of FieldPlan to the management of 4,000 field engineers are promising.

The current development work focuses on the large-scale application of FieldPlan to 25,000 BT field engineers. Future research will include the development of a generic planning framework and its extension to strategic planning scenarios.

References

1. Azarmi N. and Smith R. (1995) Intelligent Scheduling and Planning Systems for Telecommunications Resource Management. In *BT Technology Journal*, 13(1):7–15.
2. Laithwaite R. (1995) Work Allocation Challenges and Solutions in a Large-Scale Work Management Environment. In *BT Technology Journal*, 13(1): 46–54.
3. Voudouris C. et.al. (2005) ARMS: An Automated Resource Management System for British Telecommunications plc., In *European Journal of Operational Research*.
4. Owusu G. et. al. (2003) ARMS: Application of AI and OR Methods to Resource Management, In *BT Technology Journal*, 21(4):27–32.
5. Lesaint D., Voudouris C. and Azarmi, N. (2000) Dynamic Workforce Scheduling for British Telecommunications plc., In *Interfaces*, 30(1):55–66.
6. Kirkpatrick S., Gelatt C. D. and Vecchi M. P. (1983) Optimization by Simulated Annealing, In *Science*, 220(4598):671–680.
7. Glover F. (1989) Tabu search: Part I. In *OSRA Journal on Computing*, 1(3):190–206.
8. Holland J. H. (1975) *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, Michigan.
9. Voudouris C. and Tsang E. (1999) Guided Local Search and its application to the Travelling Salesman Problem. In *European Journal of Operational Research*, 113(2):469–499.
10. Aberdeen Group Report (2004) *The Field Service Optimization Benchmark Report - Tapping the Service Supply Chain for Profit and Competitive Advantage*.
11. Aberdeen Group Report (2005) *Optimising Field Service to Achieve Profitability Goals*.