

# **EPO: A knowledge based system for wood procurement management**

**Seppo Linnainmaa, Juha Savola and Olli Jokinen**

VTT Information Technology  
P.O.Box 1201, FIN-02044 VTT, Finland  
seppo.linnainmaa@vtt.fi

## **Problem**

Enso-Gutzeit is a Finnish company operating in pulp and paper industry. It has 1200 employees and its annual turnover is 700 million US Dollars. For its factories and sawmills it requires more than 600 million cubic feet of timber as raw material annually, equalling to about 400000 truckloads. Most of the timber is loaded from forests located all around Finland. About 85% of the logs are transported directly to the factories and sawmills by 250 trucks operated by Enso. Alternatively parts of the transportation are carried out by trains or waterways.

The available timber is classified to several tens of categories, depending not only of the species of trees, but also of the quality and length of the logs. Every factory estimates its daily need in each category.

Traditionally the dispatch centers have allocated one or more trucks for each foreman in the forests, depending on the amount of available timber, and the foremen have taken care of the control of trucks. Through increasing competition it became evident that this mode of operation had to be changed. Without more global guidance the trucks drove unloaded 50% of all their driving, and in worst cases two trucks loaded with the same type of timber passed each other on the road, driving into opposite directions.

To guarantee the needs of the factories, rather large piles of logs were stored both in the forests and around factories. This storage period caused waste of money, and also the quality of the timber decreases when it is stored. Also a law was made recently in Finland stating that it is illegal to store timber in the woods more than a couple of weeks in summertime, because they provide an excellent reproduction environment for pests. On the other hand, when the factories have practically no buffer storage, quite exact schedules for arriving timber are mandatory. The existing, mainly manual system for wood procurement was far too inaccurate and slow to face the new challenges.

## **Solution: the EPO system**

In the turn of the decade, Enso decided to develop EPO, a completely new system for its wood procurement, based on the highest available information technology. EPO is the Finnish acronym for Enso-Gutzeit's Wood Procurement Management System.

The EPO system as a whole covers all the operations of the wood procurement from strategic planning to actual deliveries to the factories and sawmills. The necessary data is collected on-line from the computing equipment of the forest foremen, forest harvesters and tractors, trucks, and plant receptions. In addition to data transmission through telecommunication cables, mobile data transmission is used where necessary. The foremen register the locations of forthcoming wood batches, consisting of piles of logs, on their electronic maps. The trucks can see their destinations on the electronic maps in their cabins, as well as the locations of themselves, determined by GPS satellite navigation (see Fig. 1).

The heart and the most challenging part of the system is EPO2 software running in UNIX workstations of regional dispatch centers, each managing about 20 trucks. The main focus of this paper is in EPO2 software, developed by VTT (Technical Research Centre of Finland) as a subcontractor of Carelcomp. Carelcomp is a Finnish software house specialized in information systems for forest industry. It has been responsible of the development of EPO as a whole. The purpose of the EPO2 subsystem is to determine optimal weekly routes for the trucks, knowing the availability of timber and transportation vehicles, and the needs of the plants.

The complexity of the problems solved in EPO2 are caused by numerous reasons. There are lots of conflicting goals that should be optimized, and hundreds of different kinds of strict and desirable constraints are involved in the problem, some of which are very difficult to express in a mathematical format. To give an idea of the nature of the existing constraints, some of them are briefly described here.

Each wood batch has a maximal life span. It cannot be transported before it is piled up, and regulations indicate

the maximum amount of time that a batch is allowed to stay in the woods before transportation. A strict rule is that on each batch a truck must take either its load full or all the remaining logs. Each plant has a special delivery program, determining how much and which kind of raw material the plant wants to receive within certain time intervals. As strictest, a plant may accept e.g. only birch fiber logs with length of 4 meters, but extra computational load is caused also by such facts that some raw material can be replaced partially by a substitute material, provided that the substitute material is essentially cheaper when the transportation cost is included.

Many plants require in addition that the trucks arrive in even intervals within the given opening hours. Regulations exist also e.g. for the working hours of the drivers, and naturally to weight limits of the loads. Agreements may exist that a certain driver is granted to collect the batches in a certain area. One main goal is to minimize the amount of empty driving, and thus special care must be taken in cases where (typically very rare) material is transported very far from its original location. Of course the most important goal is the minimization of truck driving as a whole.

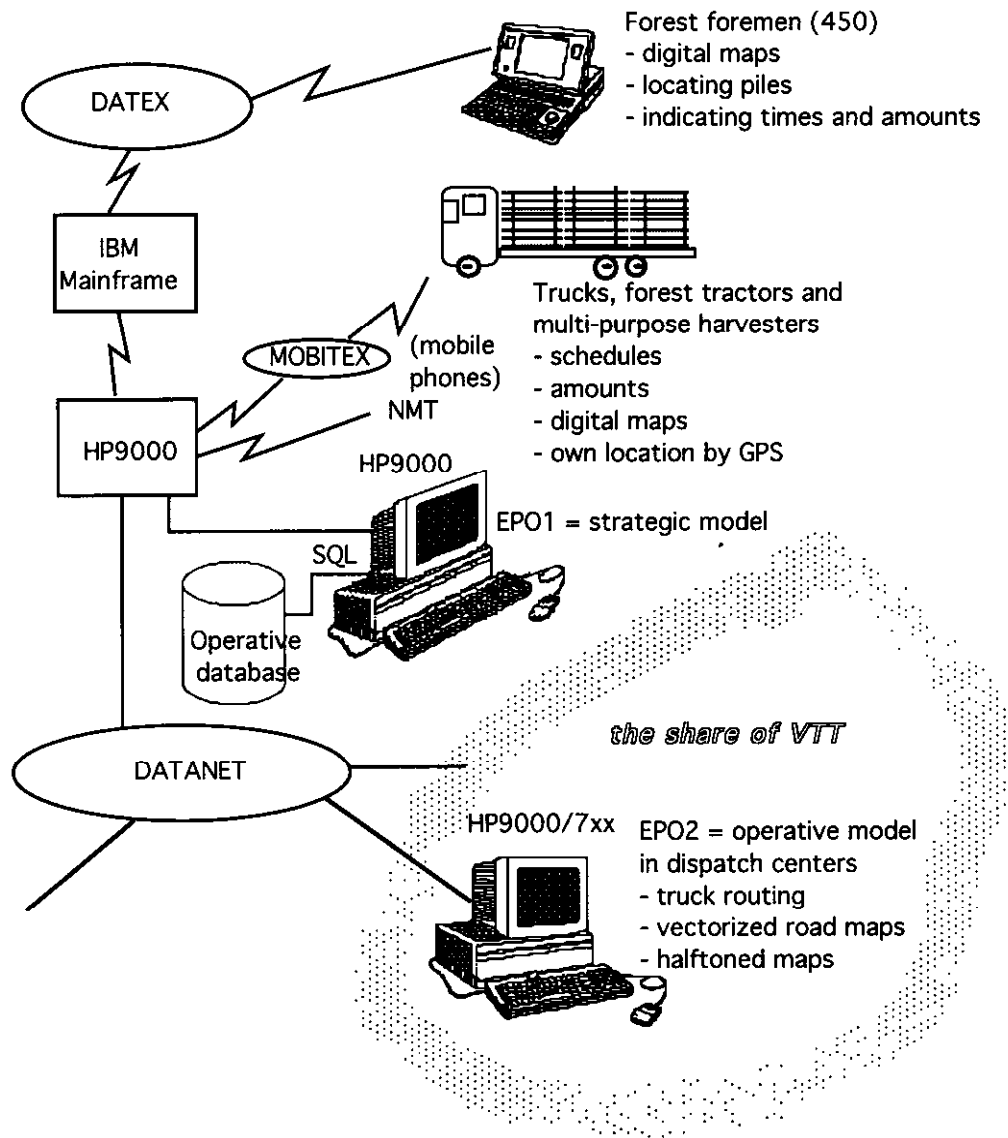


Figure 1. An overview of the EPO system

Even the strongest methods developed by Operations Research require far too much resources to be applicable to solve the whole problem. Our approach to tackle the problem has been combination of exact mathematical methods and heuristic, knowledge based methods. Exact methods are used whenever they are considered to be fast enough to solve the problem where they are applied. This kind of combination does not guarantee that the strictly optimal result is reached, but by any practical measures it produces a good result, anyway. In fact, due to the conflicting goals no unique optimal result does even exist. And for rare cases where the human dispatcher disagrees with the automatically produced results, he or she has also efficient interactive tools to modify the results manually or semi-automatically.

### Solving the dispatching problem

The EPO2 software, i.e. the dispatching subsystem of EPO, is the core of the whole system. Input for this subsystem consists mainly of following data:

- digitized **road network** with properties attached to each road segment
- halftoned **road map** for the graphical interactive user interface
- locations, categories, amounts, timings and possible advance reservations of **wood batches**
- needs and schedules of **factories** and other unloading sites
- available **trucks** with their properties and original locations
- sets of **parameters** determining e.g. preferences between conflicting goals

The main output consists of the schedules for each truck, indicating how much and which kind of logs the truck is expected to load at each batch that belongs to its route. Also exact driving directions could be provided, but the drivers prefer to use their own judgement to determine the optimal route between each consecutive pair of batches, instead of receiving too detailed instructions from a computer.

Practically everything in the dispatching subsystem is realized as objects, e.g. road segments, wood batches, trucks, factories, and problem solvers. The implementation language has been C++ and the system runs on HP9000/7xx workstations under the HP-UX version of the UNIX operating system.

To get an idea of the objects let's consider a road segment object. The basic information consists of its own position and length, but it also knows which other road segments are connected to it, and which is the typical *driving speed on it*. Also other aspects that may affect its usage may be indicated: its hilliness, possible narrow or shallow points, bridges with maximum loads etc. A truck object can inquire all these aspects when it determines whether it should choose a certain road segment for its

route. Of course, a truck object knows e.g. its own height and weight, that are needed when considering whether the driving on a road segment is possible.

The road network covered by the subsystem consists of almost 100000 road segments, a road segment being the part of a road that lies between two consecutive junctions. The most important property of a road segment, excluding its location, is the road class where it belongs to. This information is the main parameter when the driving speed on the segment is estimated. The digitized road network and their classification is provided by Karttakeskus (Map Center), a Finnish company specialized in map production. Also high quality halftoned digital maps, being exact copies of typical road maps used by drivers, are produced by the same company.


All map material is further processed in the preparatory phases of the EPO system to guarantee extremely fast processing both for calculations and for interactive viewing operations performed by dispatchers. For instance, the actual road network is reduced by effective mathematical methods to a "virtual" road network consisting only of such road segments that are included in at least one optimal route from some batch or factory to another. All other segments are hidden during the actual computations.

The information concerning wood batches, factories and trucks is collected from a relational database located in a mainframe computer, using predetermined SQL queries, before the actual dispatching process is started. The information is transformed and stored in an internal object database in the workstation of the dispatcher, allowing extremely flexible and fast processing of the required data during the dispatching process.

The actual dispatching is performed in three phases. In the first phase it is determined which batches are reserved to each factory. The second phase produces automatically *generated suggestions for the routes, loading tasks and schedules for the trucks*. The third phase is a semi-manual postprocessing phase where possible conflicts between the automatically produced plans and the human view of the dispatchers and the forest foremen are resolved.

To determine which batches are reserved to each factory or other unloading site, an "attraction model" is used. A virtual wave front for each wood category leaves every factory through all possible roads. The speed of the wave front is the larger the more the factory needs wood in the present category. When the wave front finds a wood batch, it reserves it to its own factory. The propagation of a wave front is stopped immediately when it hits another wave front looking for the same category. An example of a result of this phase is presented in Fig. 2.

The second phase in the dispatching process is the most demanding one. Due to the diversity of the methods used in this phase it is impossible to describe its operational principles briefly. It starts by noticing where the trucks are currently, and which batches should be transported to the factories to fulfil their requirements for amounts in each wood category and schedules. One of the subproblems in this phase is the well known travelling salesman problem,



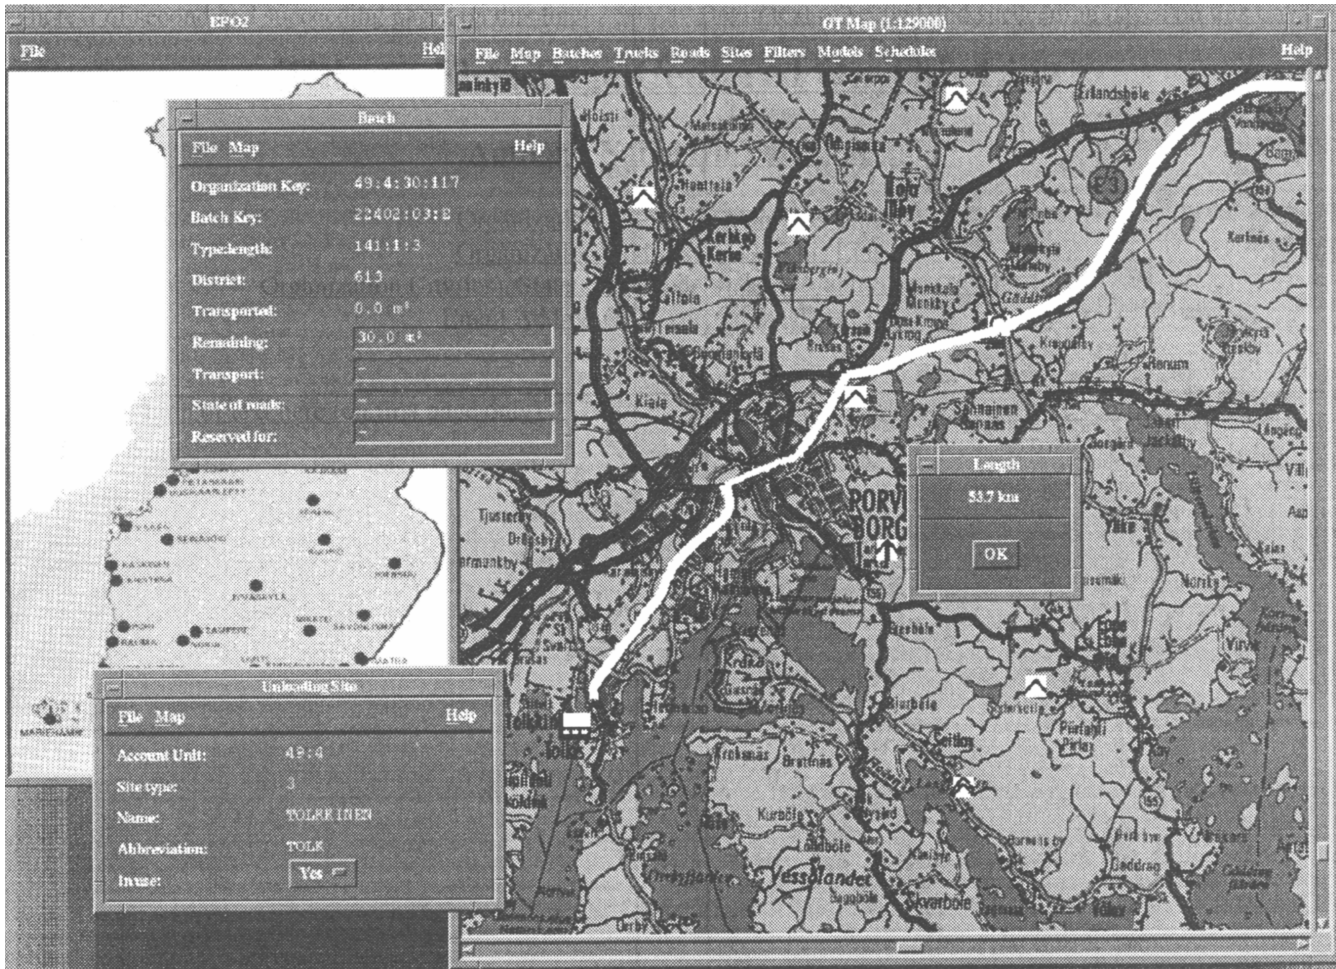
**Figure 2.** Visualization of attraction regions. Wood batches are indicated by circles, unloading sites by squares. Each batch has the same color as the factory where it is heading to.

finding the optimal route for a truck during its trip once the involved batches are determined. Heuristic AI related methods are heavily used due to complexity of the problems. Also many kinds of special knowledge is used in this phase. For example, some drivers are known to be more efficient than others in general, and the system can also be told that a driver is very familiar with some region and thus more preferable than other drivers. Weather conditions may affect the schedules, too. The results of this phase depend heavily on the values and preferences determined by settings of various parameters.

The third phase, the postprocessing, is required for various reasons. One of them is the fact that heuristic methods do not necessarily always produce an at least near optimal solution. An experienced dispatcher with an efficient visual interface can easily spot possible tasks

suggested by the computer that do not make sense. The possibility for manual postprocessing relieves problems caused by the fact that it is obviously impossible to program into the software all human aspects that should be considered in some special cases. Also possible errors in the initial data may be detected in this phase. A foreman may notice that he or she has been too optimistic in estimating how fast a wood batch is ready for loading, etc. Manual changes may also be required because of broken trucks, jammed roads or changes in weather conditions. The dispatcher may also change parameters to reflect the changed situation and run the automatic dispatching phase again, possibly with a subset of the original data.

The resulting orders are transmitted to the computers in the trucks via mobile phones, and each driver has a clear schedule available while starting the daily driving.



**Figure 3.** A map interface for the EPO2 dispatching software. Note the mouse sensitive symbols for an unloading site and different types of wood batches (clicking of which generates textual pop-up windows with more detailed information), a shortest route indication, scrolling bars for map sliding, and the global map locator.  
Base map © Karttakeskus (Map Center, Finland).


### Running the dispatching subsystem

During each dispatching run, the scheduling is determined typically for a period of one week, but the length of this period can be freely changed via input parameters. Normally the dispatching starts during a night with automatically launched queries that collect the necessary data through data transmission network. All preparatory processing for the data is also made automatically before the actual dispatching process. The first suggestions are ready within a couple of hours after starting the queries, and the actual dispatching for 20 trucks for one week is done in a quarter of an hour. Thus the dispatcher has enough time to study several combinations of parameters with different weightings to see the effects in the results, if this is considered preferable.

There are many kinds of visual tools that support the dispatcher to judge and transform the results generated by the system. The main tool is a map interface where the target area can be easily selected, viewed in preferred

resolutions, and scrolled horizontally and vertically. The locations of the factories and wood batches are seen on the map, and there are versatile filtering possibilities to hide all but the interesting batches. The factories, batches and other icons seen on the screen are mouse sensitive. By clicking on them, detailed information on wood categories, amounts and other available information becomes visible. Of course, already the shape of a batch icon indicates rough category of the wood in the batch. The interface may also be used e.g. for such general operations as determining the optimal route between any two spots on the map, measured either in mileage or in driving time. A sample of a map interface is shown in Fig. 3.

The dispatcher is also able to generate summaries of the results produced by the system, such as the visualization of attraction regions, shown in Fig. 2. A very useful chart is e.g. a collected schedule for all trucks, as shown in Fig. 4. Similar charts can be generated also e.g. for showing the schedules of all loads arriving to each factory during the day.



**Figure 4.** Rough overview of the schedules for each truck. The circles indicate loadings, the numbers attached to them are foreman numbers. The line rising from each symbol of a factory indicates the originally expected timing for the arrival of the truck. For closer inspection, also amore detailed schedule for each separate truck is easily available.

### Deployment and conclusions

The development of the EPO system was started in summer 1991. Originally it was expected to be ready before the end of 1992, but this schedule showed to be far too optimistic. A major reason for this was the pioneering nature of the work. Many of the features that were required, e.g.

digitized maps and wireless communication systems were advertised to exist, but they were not tested in any such large and real system as EPO, and lots of improvements were necessary before all the operations became reliable enough. Also the multitude of all kinds of constraints that had to be considered in building the dispatching subsystem EPO2 was heavily underestimated.

Anyway, gradually the obstacles were overcome, and the approaches showed to be well chosen. Currently even the numerous people inside the involved organizations who were very sceptic at the beginning of the development project are enthusiastic to admit that the results are far better than what they were able to expect, if the system would ever be running.

The first trucks with routing generated by the EPO system started to drive in 1993. During the year 1995 the system will gradually cover all Enso's domestic wood procurement, once also the rest of the vehicles, dispatchers and foremen are equipped with the required hardware and software, and educated to use them.

Based on the achieved operational results, the direct annual savings for Enso caused by the EPO system are estimated to be at least several millions US Dollars. Exact amounts are impossible to determine, because the whole organization of Enso has been strongly changed, EPO being one of the driving forces in these changes. There are also numerous side effects where it is impossible to determine what is EPO's share. For example, much more exact information of available timber provides a much firmer basis in commercial negotiations than the previous rough estimates.

Currently the EPO system supports only trucks but later versions will include trains and waterways as well. Already now the time required for the dispatching has been reduced to a fraction of the time required by the first operational runs a year ago, by implementing many kinds of new insights. There is no doubt that this progress will continue. Also some completely new approaches will be experimented, for instance the possibilities to further improve the results of the automatic dispatching phase by using neural nets.