Article reprinted from *OR/MS Today*, April 2005

ORMS TODAY

April 2005
Volume 32 • Number 2

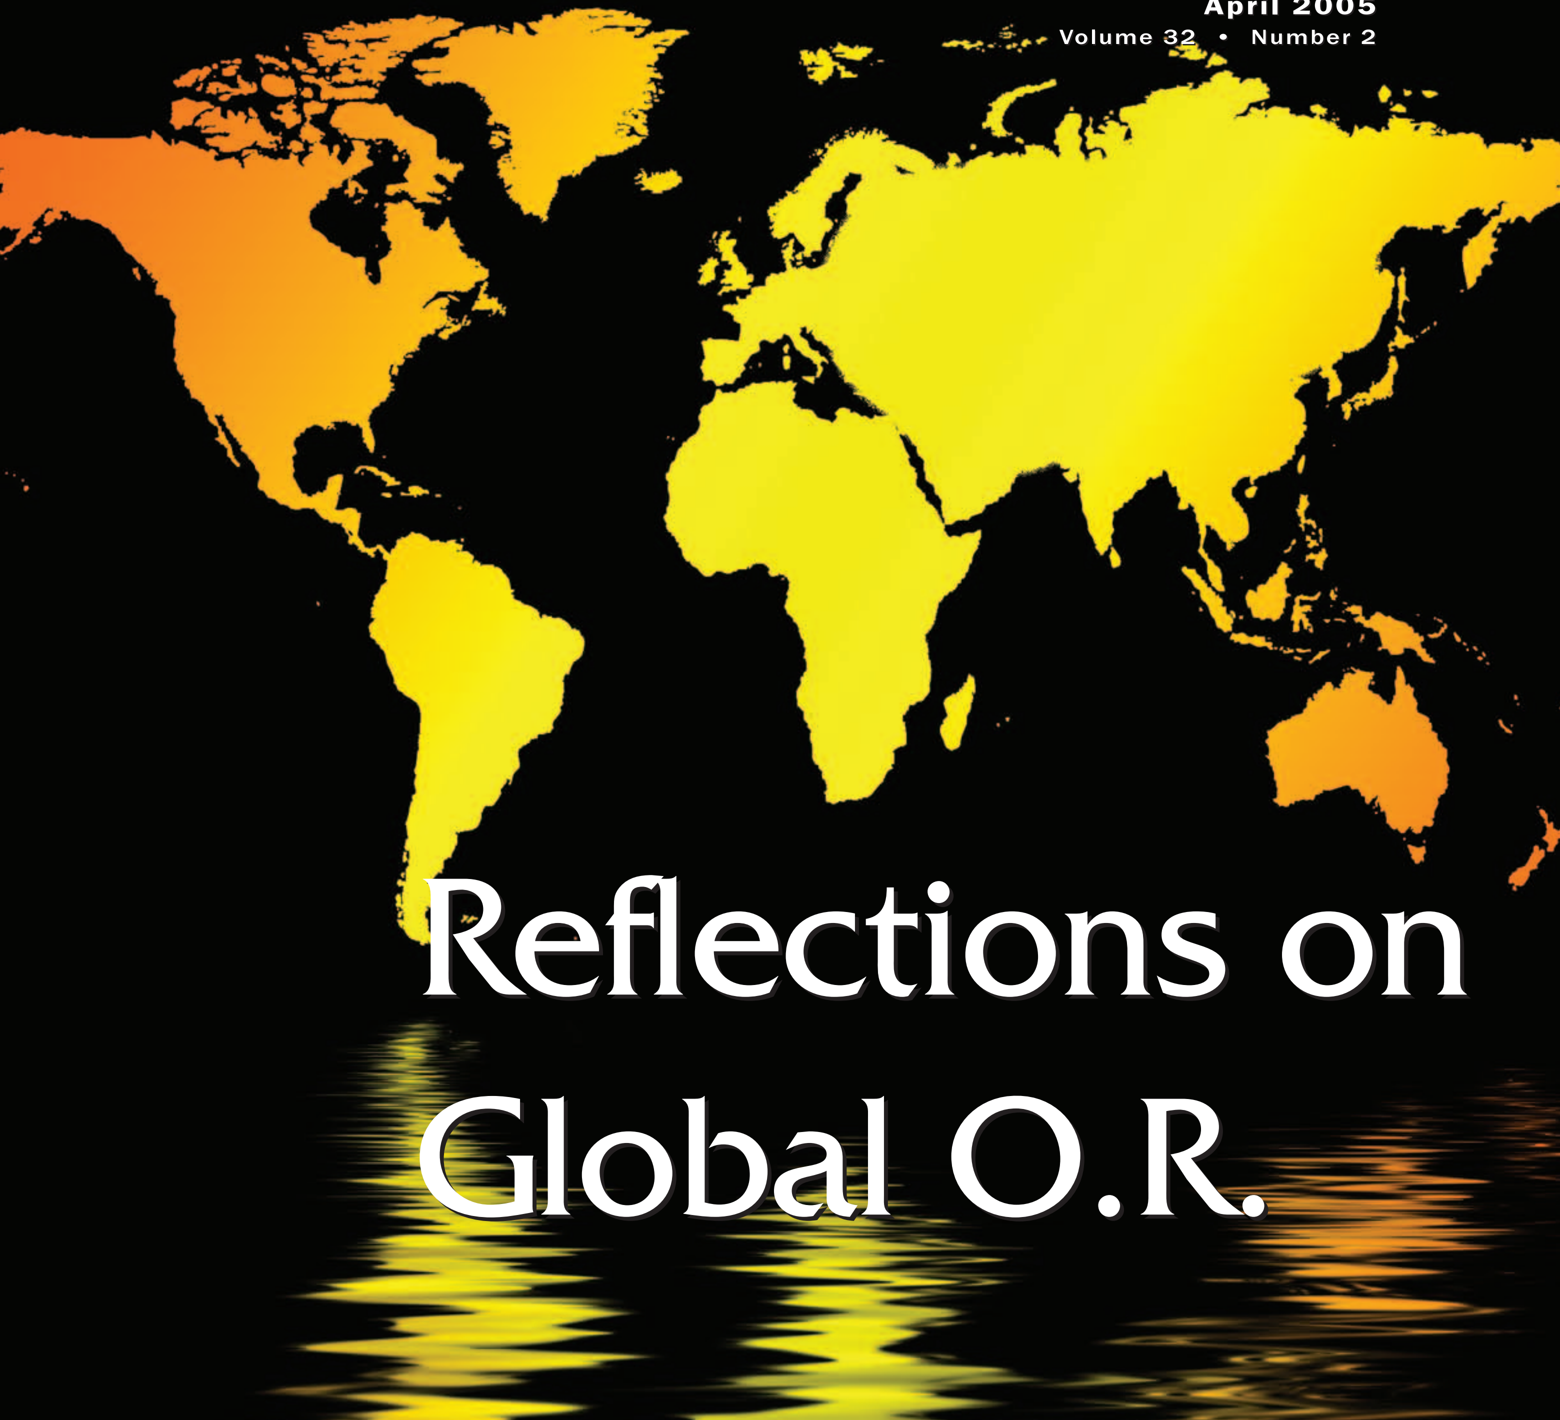

Reflections on Global O.R.

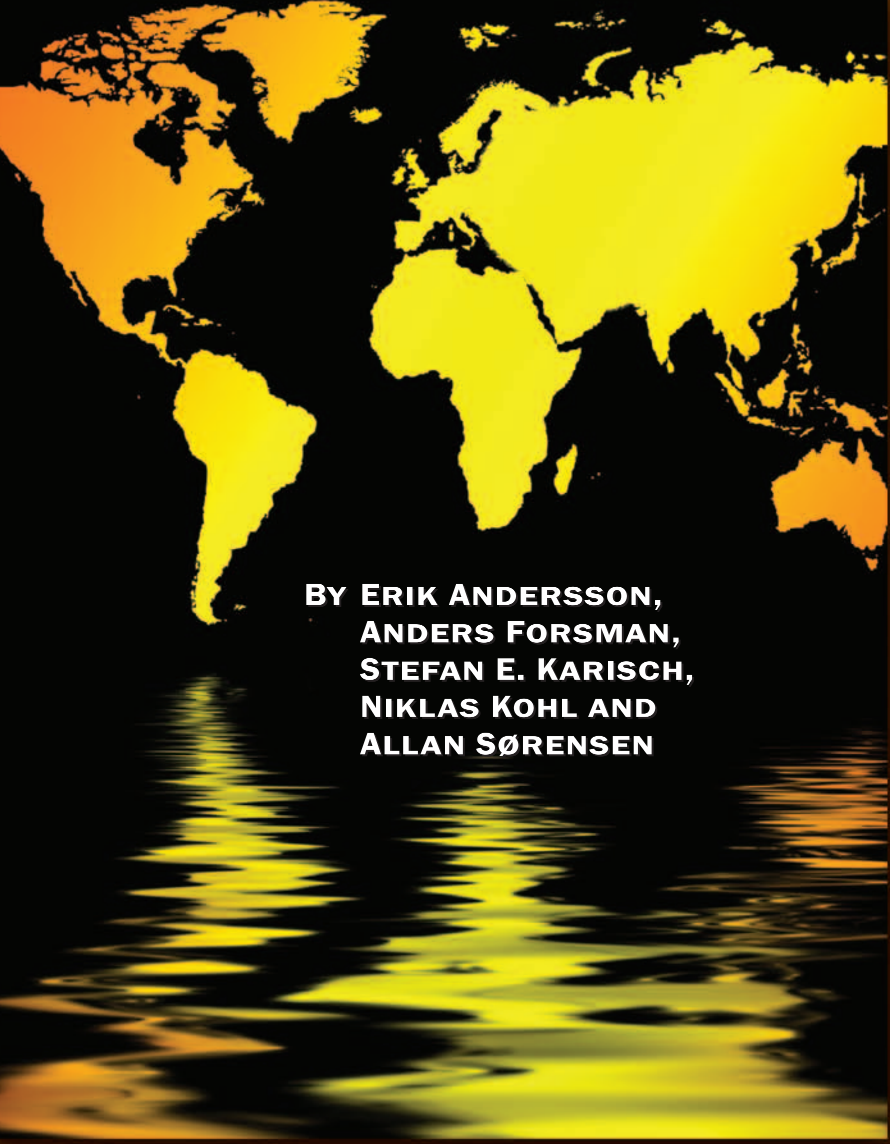

BY ERIK ANDERSSON,
ANDERS FORSMAN,
STEFAN E. KARISCH,
NIKLAS KOHL AND
ALLAN SØRENSEN

Marrying end-user modeling and large-scale optimization pays off at SAS and other carriers.

Problem Solving in Airline Operations

THE AIRLINE INDUSTRY EXPERIENCES VERY CHALLENGING TIMES, and many airlines need to undertake substantial changes to their business processes to get back to profitability. Because an airline's operations are generally considered a cost driver, carriers emphasize cost-effectiveness to make improvements to the bottom line. This push towards cost savings is supported by the use of optimization systems that improve the utilization of scarce and expensive resources such as aircraft, crew, gates, etc. To maximize the benefits of resource optimization, one needs to identify, model and solve the right operational problems. Once this is done, it is crucial to maintain flexibility and adjust to changes in the business environment.

The challenge of problem solving is to find the best possible solution, to the right problem, as fast as possible. These three dimensions are all essential to the success of optimization systems. Depending on the environment, some dimensions are emphasized more than others, and some might even be neglected. Often, the second dimension – the right problem – is neglected in scientific work related to real-world problems such as those arising in the airline industry or other complex transportation systems.

There are many planning and operations problems at airlines for which, due to their complexity, detailed and accurate modeling is required to obtain useful and efficient solutions. At the same time there is continuous change in an airline environment, e.g., modifications to flight schedules, changes to agreements and disruptions of plans. Other characteristics that make these airline optimization problems challenging and important are their size, the amount of data involved and their impact on the profitability of an airline.

A large airline operation comprises several hundred aircraft, 10,000 or more air crew members, up to 100,000 flights per month and tens of millions of passengers per year. Obtaining continuously the best possible solution to the right problem is hence crucial for an airline to stay competitive and survive.

Figure 1 provides a simplified, high-level overview of the operations of an airline and gives a typical timeline for the different planning steps. First, the schedule (or timetable) is produced. Here the objective is to match marketing expectations with available fleets and constraints on the network. The second step in airline planning is the allocation of aircraft to the flights. This step involves determining the right type and size of aircraft for each flight leg in order to maximize the expected profit and constructing aircraft rotations that satisfy operational constraints and maintenance requirements. Once this is done, the crew needs on each flight are known and become input for the crew-scheduling problem.

Crew scheduling is traditionally done in two steps. First the crew-pairing problem is solved, where flight legs are sequenced into anonymous crew rotations (or trips) such that the crew requirements on each leg are satisfied, crew and other costs are minimized and contractual, and operational constraints are met. Then, in the crew-rostering problem, these trips and other activities are assigned to individuals, thereby building personal rosters.

On the day of operations, all resource areas in an airline (aircraft, crew, gates, etc.) need to be controlled, and in case of disruptions, plans need to be repaired. Operations research has been very successfully applied to problems arising in airline operations. (For more details, see the recent survey on applications of O.R. in the air transport industry by Barnhart et al. [1].)

Traditionally, the size and complexity of airline operations required the decomposition of the problem into more manageable sub-problems that are by themselves nontrivial planning or scheduling problems. A global modeling tool can provide at least consistency between these sub-problems and support their eventual integration to avoid sub-optimization.

Besides the complexity of the problems arising in airline operations, there is the additional challenge of not being able to model, for example, a 200-page specification of a particular operations problem reasonably accurately. First of all, such a specification usually does not exist. And even if it did, any specifications and the resulting models would have to be revised and thought through again, and new hypotheses tested and verified. Hence, it is essential in a business environment in general, and the airline industry in particular, to be able to support a process with fast iterations to get as close as possible to the real problem and then subsequently adapt to changes in the future.

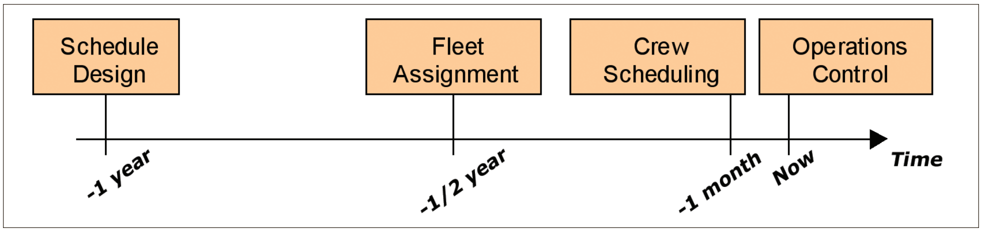

Figure 1: Major phases in airline operations

Modeling and rule tools provide transparency for the business users and allow them to change models and rules, thereby making optimization techniques more accessible for non-experts.

Modeling Systems and Rules Engines

THE MODEL IS the foundation of an optimization application, especially when used for solving real-life problems. A common approach in problem solving is the separation of the problem definition (model) from the problem solution to allow the user to focus on the modeling. General-purpose modeling systems have been continuously developed for large-scale optimization. Many different commercially available modeling languages and rule systems support this problem-solving paradigm of separating problem definition and solution. The NEOS Guide on Optimization Software lists at least 15 such modeling systems [2]. In comparison, the number of commercially available rule systems is smaller.

Besides following the problem-solving paradigm, modeling and rule tools also provide transparency for the business users and allow them to change models and rules, thereby making optimization techniques more accessible for non-experts. Using these tools, end-users can systematically explore, analyze and evaluate their models and the underlying optimization problems. Providers of modeling and rule systems have over the years also increased the usability of their systems, including application development tools and environments. These efforts have made business modeling less exclusive and more accessible to technically unskilled end-users who can now use optimization systems effectively. These strategic opportunities have already been described (see, for example, [3] and [4]).

In an airline environment, there is a need to combine large production optimization with the prototyping power of modeling systems. To our knowledge none of the commercial modeling or rule systems listed in the NEOS Guide on Optimization Software is used for production planning and scheduling at airlines or in other large transportation systems.

For more than 14 years, Carmen Systems has developed a combined modeling and rule system, Carmen Rave (Rave stands for "rule and value evaluator"), which is currently deployed by 20 airlines and three railway companies. The Carmen Rave language is a purpose-built rule and modeling language tailored for resource management and optimization problems in the transportation industry.

A combined rule and modeling tool allows the implementation of costs, definition of objective functions for optimization, expression quality aspects, etc. In other words, a combined tool allows a user to model all the characteristics of a particular application and describe them in rule code. This maps knowledge and expertise of the users into the rules engine and transforms an optimization system into an expert system. Preserving knowledge and expertise is crucial for consistency and continuity and can not be achieved in a single iteration. Such a task needs continuous effort over time and a modeling tool that allows continuous adaptation to reflect the real-world problem characteristics. Furthermore, users of optimization systems should be able to analyze their environment to increase business knowledge and expertise. The creation of many what-if scenarios and simulations is necessary in this context. By supporting the co-existence of multiple rule sets, creating new scenarios can be done easily and quickly.

CARMEN RAVE:
RULE MODELING EXAMPLE

The Carmen Rave language is a purpose-built modeling language tailored for resource management and optimization problems in the transportation industry. The language is a declarative programming language in which all objects belong to a level hierarchy that must contain a chain level (e.g. a crew roster) and an atomic level (e.g. a flight or so-called leg). The rule programmer can introduce intermediate levels (e.g. duties and trips).

For each level, attributes can be defined using other attributes and a number of built-in or user-defined functions and operators. Basic building blocks are the keywords that are predefined attributes of the leg or the chain level. There are also so-called aggregators (such as sum and max) and specifiers (such as first and next) to make it possible to construct complex expressions. The aggregators remove the need of complex recursive functions that potentially could cause infinite recursion problems. Below is a small example of a rule ensuring that the connection time between two legs or flights is at least 30 minutes. Attributes that are derived from keywords must start and end with "%", and, in the example, duty is a user-defined intermediate level.

```
%connection_time% =
next(leg(duty), departure) - arrival
;
rule connection_time_ok =
%connection_time% > 0:30;
end
```

As a consequence of the declarative nature of the language, the rule programmer never needs to consider when or on what level a rule needs to be checked. This is automatically derived from the rule description. Furthermore, the Rave compiler and runtime system has the important responsibility of communicating the model efficiently to the optimizers. This way, the rule programmer can focus on what he/she knows the best – keeping the model as close to reality as possible. The rule source code is compiled into C code and dynamically linked to the decision support or optimization system. This makes it possible for the user to maintain and modify the problem description without downtime of the application.

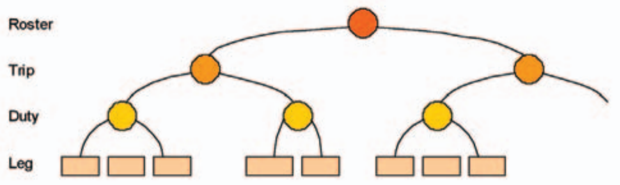

Figure 2: Example of level hierarchy.

Our business experience of selling optimization systems bundled with a combined rule and modeling system is that new customers will never compromise performance for flexibility. It is the way the optimization market works. If a major airline wants a new optimizer for their crew scheduling, it will invite several vendors and evaluate the quality of their solutions to a predefined benchmark problem. One percent difference in solution quality can mean $30 million of annual savings for the customer, so it is almost impossible for the buyer to justify anything other than the best quality solution, even if the difference is just a small fraction of a percentage point.

During the development of Carmen Rave, we had to overcome several technical difficulties. It was a challenging marriage between computer sciences and O.R. – every language feature had to be designed carefully to meet performance requirements. Techniques like dynamic sorting, caching, scope identification, partial evaluation and other program transformations were developed to this end. Planning problems and real-time problems require different techniques. A planning problem may take several hours of CPU time to solve, so it pays off to build large cache tables of precalculated rule values. Analysis of rule call patterns has been done carefully for many different use cases. There is no standard solution to the performance problem of rule evaluations, so the run-time system of Rave has evolved to a hybrid of many algorithms.

Practical Experience

Rave is used with all Carmen applications and also is available as a stand-alone product to integrate in third-party software. As indicated above, more than 20 transportation companies use the system as an integral part of Carmen's suite of resource optimization products. These companies include Air France/KLM, British Airways, Delta Air Lines, Lufthansa and Northwest Airlines – five of the 10 largest airlines worldwide – as well as Deutsche Bahn (German Railways), one of the largest passenger transportation companies in the world. The rule modeling language allows the users at these companies to interact with optimization systems by defining rules, objectives and targets rather than modifying flight, crew or passenger records.

One of Carmen's first clients was Scandinavian Airlines (SAS). In 1999, SAS decided to introduce a corporate rule system to be used within all crew management systems, regardless of vendor and technical platform. SAS had previously used Carmen Rave as a modeling and legality system with the Carmen crew-pairing system. The strategic objectives for using a corporate legality system were:

- centralized rule maintenance by users,
- automatic distribution of changes,
- consistent rules and interpretations at all times in all systems, and
- advanced simulation and test facilities.

Since the beginning of 2002, Rave controls all crew management systems at SAS, including the real-time tracking system and various optimization applications. SAS considers the ability to simulate and quickly respond to change as key to the survival of any airline. With a corporate rule system, SAS can constantly adjust to changes in the operation and can also perform simulations to analyze the consequences of potential modifications in the critical crew management systems.

The complexity of crew management at SAS is staggering. A total of 1,700 pilots and 4,800 flight attendants are scheduled with different crew management systems. Crewmembers are positioned in bases in three countries – Denmark, Norway and Sweden – and each country has its own set of national regulations. Crewmembers are also part of different unions that follow differing agreements. In addition, both pilots and flight attendants can work in several positions and belong to different qualification groups. The overall objective is to use all available crew resources to operate the domestic and international flight schedule of SAS as effectively and efficiently as possible. Last but not least, quality-of-life considerations for individual crewmembers are also a priority for SAS.

Besides meeting the strategic objectives listed above, SAS has also fulfilled the following business goals when implementing a corporate legality system:

- ***Cost reductions.*** For example, operation and maintenance costs in the crew legality and composition area have been reduced by more than 40 percent.
- ***Time to market changes and development.*** For example, new or changed agreements and rules are now implemented in hours or days instead of weeks or months, respectively.

- ***Minimization of risk.*** For example, independence of vendors and other key persons has been achieved due to transparency of rule code.

The use of the Rave modeling language at SAS is a good example of the application of the problem-solving paradigm in a complex airline environment. Being able to model problems in accurate detail is necessary to get the full benefits from optimization and decision-support systems. In general, airline clients estimate that the modeling capabilities provide them with additional 2 percent savings in crew costs on average, with some estimating savings of up to 5 percent. These savings are related only to the rules and modeling system and are achieved on top of the benefits delivered by the optimization systems. For an airline the size of SAS, 1 percent of crew cost savings corresponds to around $5 million a year. Additionally, SAS saves several million dollars in using a corporate legality system by meeting the business goals listed above.

A rule system has become a necessity for SAS to maintain the large number of labor rules and regulations that are very complex compared to industry standards. In the airline's ambition to automate and optimize the crew planning process, a global and efficient rule and modeling language has become essential. Therefore, modeling has become a strategic core competence within the crew management operations of SAS.

Conclusion

End-user modeling and large-scale resource optimization can be successfully married for solving large and complex problems arising in airline operations. The combination of optimization and modeling power in the client's hands can be viewed as the real contribution of a modeling system. Being able to model problems in detail and accurately is necessary to get the full benefits from optimization and decision-support systems. There are a number of business problems SAS and other airlines solved that they would not have been able to solve without a modeling system. Through using a combined rules and modeling engine, operations and maintenance costs can be reduced, adjustment times to market changes shortened, and a higher level of quality and consistency achieved. The contributions of such systems to an airline's bottom line can be significant.

Most importantly, however, such a modeling tool gives the power of operations research to the users at large transportation companies – to people who were not initially O.R. specialists.

***Erik Andersson** (erik.andersson@carmensystems.com) is co-founder and CTO of Carmen Systems. **Anders Forsman** (anders.forsman@carmensystems.com) is product manager of Carmen Rave at Carmen Systems and is responsible for the development of Rave. **Stefan E. Karisch** (stefan.karisch@carmensystems.com) is vice president of Operations Research at Carmen Systems. **Niklas Kohl** (niko@dsb.dk) was a senior consultant at Carmen Consulting and is now manager of IT and Optimization in the planning department of the Danish State Railways (DSB). **Allan F. Sorensen** (allan.sorensen@sas.dk) is head of IT at Scandinavian Airlines Denmark A/S with more than 20 years experience within IT and crew management systems.*

REFERENCES

1. C. Barnhart, P. Belobaba and A.R. Odini, "Applications of Operations Research in the Air Transport Industry," *Transportation Science,* Vol. 37, No. 4, pgs. 368-391, November 2003.
2. NEOS Guide: Optimization Systems/Modeling Languages, http://www-fp.mcs.anl.gov/otc/Guide/SoftwareGuide/Categories/optsysmodlang.html
3. T.A. Grossman, "Spreadsheet Modeling is a Strategic Opportunity," *OR/MS Today,* October 2003.
4. R. Fourer, "Software for Optimization," *OR/MS Today,* December 1998.

CARMEN RAVE:
ARCHITECTURE AND DEPLOYMENT

Carmen Rave is embedded in an airline IT environment. The rule and modeling system interacts with other systems through a high-performance API or a message interface for plug-in transaction systems. This follows on from the pattern of separating business rules from standard systems. In airline operations, one distinguishes two application areas, namely planning and operations. Products in these two areas distinguish themselves mainly through the type of data updates and their response times. While planning products receive batch updates of data and have response times in the order of hours, operational products get real-time updates, and their response time is expected to be in the order of seconds.

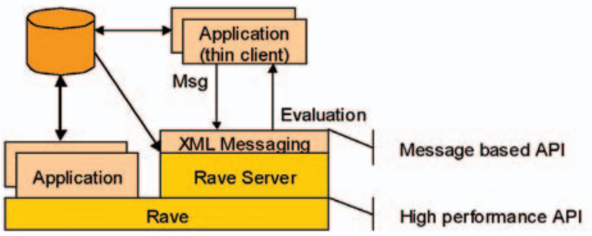

Figure 3: Directly linked Rave and the Rave Server architecture.

To meet these requirements, Rave can be employed in the following ways. In planning, Rave is usually directly linked to an application, thereby providing a high-performance API for large-scale optimization. Data that Rave needs for legality checks is accessed on demand directly from the application. For operations products, Rave is usually linked to thinner clients via messaging interfaces using a Rave Server as depicted in Figure 3. The thin clients only contain a small amount of planning data, while the bigger part of the data needed for legality checks is sent as references in the XML messages and loaded by the Rave Server.

An integrated development environment, the Rave IDE, is the interface to the rule developer. It is a graphical tool that provides one interface for centralized management of rules and objectives. The Rave IDE allows the user to edit, compile and navigate in the source code. The navigation capabilities are especially useful when working with the large amount of rules found in real-world applications.

The Rave IDE can also provide visualization of a rule or cost function evaluation. The evaluated values are displayed synchronized with the source code, which aids the rule developer in his understanding. It can be used as a debugging tool, as well.