# Knowledge Management Applications at Pirelli Tires: the P-Race and P-Truck projects

***Stefania Bandini, Sara Manzoni***
Dipartimento di Informatica Sistemistica e Comunicazione,
Università di Milano-Bicocca
{bandini, manzoni}@disco.unimib.it

This work presents two experiences of application of the knowledge-based approach to Knowledge Management (KM). The described projects (P-Race and P-Truck) have maturated from the collaboration of the Department of Computer Science, Systems and Communication (DISCo) of the University of Milano-Bicocca (mainly the Knowledge Engineering Lab) and Pirelli Tires (in particular the Business Unit Truck and the Motorsports Department).

Different knowledge-based solutions have been designed and developed to support the representation and the development of core knowledge in the product innovation process and two different contexts have been investigated. The first experience, P-Race project, refers to the chemical formulation of rubber compounds of tire tread in order to take part (and win) in motor racing. Because of the different competences involved in the decision making process (the compound designer and the race engineer), multiple knowledge representations have been adopted, and integrated into a unique Case-Based Reasoning (CBR) computational framework. The case-based approach captures the episodic knowledge that characterizes most of reasoning activity of the race engineer. Moreover, a dedicated representation formalism called Abstract Compounds Machine (ACM) has been studied in order to allow the core knowledge about rubber compounds to be explicitly represented, computed and integrated in the CBR architecture. The most meaningful and innovative contribution of P-Race consists of a general case-based architecture where the adaptation step is performed by the ACM chemical formulation model.

P-Truck is an ongoing project that has born after the success of the P-Race project. Its aim is to develop a KM system to support the Business Unit Truck of Pirelli Tires in order to improve both design and manufacturing activities of truck tires. The design approach of the solution proposed to Pirelli Tires is, again, deeply knowledge-based. A knowledge acquisition campaign and direct observation of Business Unit working activities have been conducted. The result of this design phase is an integrated KM system in which different knowledge-based and service modules exploit the Pirelli organizational memory to support the main design and productive activities of the Business Unit Truck of Pirelli Tires. In particular a knowledge-based module, based on the previously introduced model for the design of rubber compounds (i.e. ACM model), supports compound designers in the definition of product specification. Moreover, rule based reasoning and CBR approaches are exploited to support the design of some tire production processes (e.g. mixing of rubber compound ingredients and tire vulcanization). Moreover, the CBR paradigm has been applied within the P-Truck project in order to support the tuning of production anomalies that may occur during the production process. This is a fundamental task within the manufacturing context and it is strictly related to the experience of people involved in manufacturing.

# The P-Race Project

### *The Design of Rubber Compounds for Motor Racing*

In motor racing the role of tires is crucial. Several competences are involved in the decision making process about the right rubber compound to be provided to each racing team, and the main ones are owned by car drivers, racing team components, tire designers, race engineers, and compound designer. The whole team acts according to the experience of its components on the field and their knowledge about a very complex decision making problem. In particular, race engineers and compound designers are usually the ones that collect knowledge and experience of the whole community that can be physically distributed and directly interact in order to provide a solution to a problem concerning the whole community. The experience of this team is strongly dependent on performance and results obtained in previous races on tracks they consider similar to the current one. Moreover, even if in a previous race their choice had led to success, the improvement of some factors (e.g. grip, warm-up, thermal and mechanical stability, resistance to wear) is anyway advisable, because of possible improvements in the products of competitor tire providers. Therefore, the choice of a tread for a given race depends mostly on the results of previously solved cases: the general problem solving mechanism used by race engineers and compound designers is strongly based on reasoning about past cases in order to solve a new case.

Among the parts that must be assembled in order to build a tire, tread is one of the most important. It is a chemical compound represented by a recipe that determines its major properties. The basic material composing tread obtained by the recipe is called in jargon *batch*. Tread batch comprises a set of ingredients: artificial or natural elastomers (rubber), active fillers (carbon black, silica), accelerants, oils, and some others. All these ingredients are essential for the acquisition of the desired chemical-physical properties determining tire performance. The knowledge about the recipes is the chemical formulation of rubber compounds, and is a large part of core knowledge of a tire company (while the knowledge about the structure of the tire is another large portion). Tread batch is a core product and any innovation of it represents an innovation involving the whole tire. Any innovation on tread batch influences the production process, both in the case of tires dedicated to car racing, and all the other products of a tire company (large-scale products, as tires for cars, trucks, motorbikes, and so on).

In this context, the involved knowledge that has to be acquire, structured, classified and represented regards, for instance, morphological features of the racing tracks, weather and track conditions, data about the type of the race, data concerning of the car team, the adopted recipes, time measurements for each test or race, comments about the race, and so on. A computer based tool for the support of this community-of-practice has to propose the reuse of previously adopted tires to analogue races and teams and to support the adaptation of previously adopted rubber compounds of tires.

Knowledge creation and sharing among race engineers and compound designers, together with incremental learning, are other main problems that the supporting tool has to address. The following section will introduce the P-Race system, a supporting tool developed for the support of the motor-sport department of Pirelli Tires in designing and proposing tire solutions to racing teams (see for more details on the domain and approach of the P-Race project (Bandini and Manzoni, 2000), (Bandini and Manzoni, 2002)).

### *The P-Race System*

Figure 1 shows the general architecture of the P-Race system. It can be divided into three main parts (A, B, and C, drawn by dotted lines in the figure).

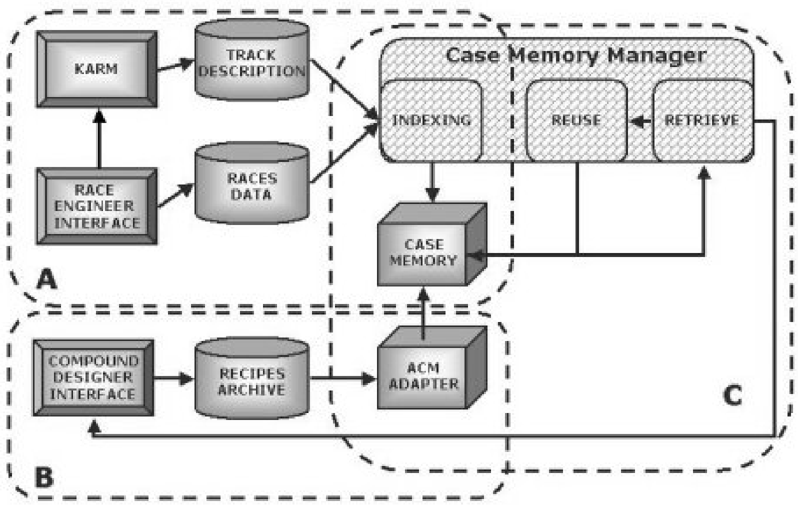


**Figure 1. The P-Race System**

Part A contains the main components dedicated to the race engineer:

- a user interface for system users;
- a database containing all meaningful data about past racing activity (dates, championships, cars, teams, drivers, trial and race times, coded recipes of the used tread batch, coded information about the tire structure, and so on);
- a dedicated Knowledge Acquisition and Representation Module (KARM) for the management of the main knowledge that is involved in the race engineer decision making process.

Part B is composed by components supporting the activity of the compound designer:

- a dedicated integration interface with the recipes database and other confidential data contained in the information system of the company;
- the Abstract Compound Model (ACM) module, that adapts retrieved solutions to the current problem;
- a user interface allowing to system users to access and use the ACM module.

Finally, the Part C constitutes the Case Based Reasoning core, and comprises:

- the Case Memory, where the pragmatic features of races are indexed and structured in cases;
- the Case Memory Manager indexing data from track descriptions and races database in form of cases, and evaluating the similarity between the current case and the stored ones.

The race engineer can interact with the system in three different ways, that is, he can input a track description, update other racing data, or activate the case based engine. The first two activities have the purpose to provide the system with all the information needed to support the race engineer's choice of the most suitable tire for a given race. The KARM has been designed in order to let users express their knowledge about tracks both in qualitative and quantitative ways, avoiding as much as possible subjective descriptions. In particular, the profile of tracks where races take place is acquired by both a dedicated user interface (see Figure 2) and a telemetric data acquisition module.

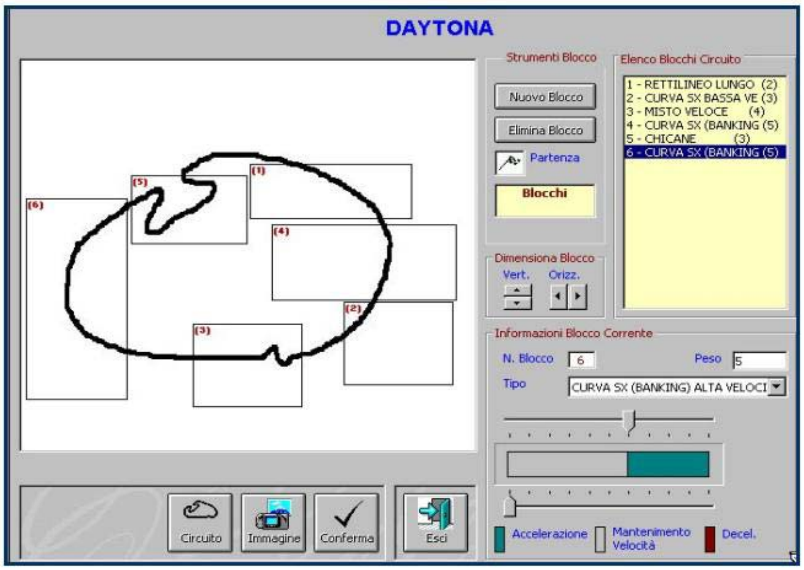


**Figure 2. The Race Engineer Interface for the description of tracks' profile**

The qualitative description provided by system users about weather and track condition forecasting is acquired by another dedicated user interface. This knowledge is acquired and represented by KARM, and the P-Race system exploit it to suggest the reuse of previously experienced tire solutions under similar weather, track and racing conditions (see (Bandini and Manzoni, 2001a) for more details about the fuzzy knowledge representation and similarity computation performed by KARM and Case Memory Manager).

Each track is described through the graphical user interface of the KARM as a set of *blocks*. Each block is characterized by a type (stretch, bend, chicane, and so on) and by the level of stress tires have to withstand (computed by the system, for instance, according to features of the road ground). The resulting description of tracks allows capturing race engineer experience and knowledge in terms of crucial information about tracks. The main ones are, for example, the features of track bends according to their severity and the required tire performance, the characteristics of the track surface, the thermal variation from a straight stretch to a bend and vice versa. This type of representation allows the system to compare tracks according to their morphological profiles in terms of race heuristics (e. g., initial and final speed in a bend, gear used in a given part of the track, weight supported by each wheel, and so on).

Another dedicated interface allows the race engineer to update the database containing data about the racing activity (dates, kind of championship, car, team, drivers, trial times, warm-up times, race times, coded recipe of the used tread batch, coded information about the tire structure, tire performances, and so on). This type of activity is usually done directly on the field, that is, at the circuit during the competition.

The system uses these data about tracks and races to support the race engineer in the solution of new problems. The reasoning process starts with the representation of the current problem as a new case to be solved. A case represents a set of chronometrical measurements, concerning a race or a trial, relevant for the performance or the technical solution adopted. As in any CBR system, the three major parts of a case are problem/situation description, solution, and outcome. In P-Race, the description of the current problem contains both qualitative (weather forecast and track conditions) and quantitative information (date, time and location of the event) used by the system to retrieve from the case memory the most similar cases. The solution for a case describes the coded recipe of the batch used in that case, while the outcome represents the resulting state in terms of performances obtained when the solution was applied. Starting from the description of the current problem, the system examines the case memory containing past problems already solved, and proposes a list of solutions (the most similar cases) to the race engineer.

The main task of the retrieval algorithm is to apply a function giving a measure of similarity among cases. In the P-Race system, the similarity function has been defined as the weighted sum of differences between attributes, some of which are the result of a fuzzy interpretation of the user input. Case retrieval is based on knowledge about tracks, weather conditions, and type of track surface. The list of solutions proposed by the system could, at this point, include a feasible solution for the problem at hand that could be directly applied. Otherwise, an adaptation process has to modify one of the solutions proposed. The system also supports users in feasibility evaluation, reporting in a structured way the outcomes of proposed cases, including all documents associated to each case (race engineers' and drivers' comments after a race or a test on track; quality-values vectors stating results in terms of performances obtained applying the solution; and so on). Thus, the system offers the view of all the current conditions the user needs to make his decision on which performance must be reached. At this point, if some modification to the basic recipe is needed, in order to improve or achieve some desired performance, the adaptation process (i.e. product innovation) is invoked. Adaptation could be necessary, for instance, when the proposed solution contains ingredients no longer available for tire production or when the past use of the solution had led top undesired outcomes.

A dedicated module of the P-Race system has been developed in order to support the compound designer in product innovation (i.e., the chemical formulation of a batch). It provides access to the information contained in a database of recipes and other confidential data about raw materials. The formulation of a new compound is guided by the requests of the race engineer that asks for the improvement of some performance of an existing batch.

Starting from the solution and the outcome of a case retrieved by the P-Race module dedicated to race engineer, the compound designer examines the recipe and the race conditions in order to fulfill the race engineer's requests. The decision process of the compound designer can be divided in three separate stages:

- Batch evaluation: the expert is usually able to assess the performance of a compound from its parameters. The evaluation of the compound designer is usually different from the one given by the race engineer: the former judges according to his theoretical knowledge about materials, while the latter examines the results of the races.
- Definition of the objective: starting from to the results of the analysis of the previous point, and to the information about the race context, the expert decides which property of the batch has to be changed in order to obtain the desired performance. The properties mostly involved are grip, thermal stability, mechanical stability, and warm-up. At the end of this stage, the compound designer has a set of possible options leading to the needed improvement of the performance.
- Choice of the ingredient: finally, the compound designer describes the batch recipes contained in the archive of the company as lists of ingredient together with their respective quantities. Then, according to the chemical and physical properties of the raw materials, chooses an ingredient, and decides whether (and how) its quantity has to be changed in quantity or if the ingredient must be substituted by another one.

P-Race supports the activities described above with a dedicated adapter module based on the ACM model and called ACM Adapter. The ACM Adapter activates the integration interface with the recipes archive of the company in order to provide to the ACM component the decoded recipe expressed in terms of quantity of ingredients. In other words, it modifies the recipe of the proposed batch in order to improve the performance observed in the outcomes of the past case, or to obtain new performances in relation with the description of the new case.

Compound revision follows the application of four sets of rules:

1. *Description Rules*, describing a product already developed as a recipe according to the ACM model, that is, as a vector of quantities of ingredients.
2. *Performance-Properties Rules*, defining which changes are needed in the properties of the recipe, in order to obtain a given change in performance.
3. *Ingredients-Properties Rules*, defining which attributes of the ingredients of a recipe are involved in the modification of the properties of the recipe.
4. *Formulation Rules*, generating a revised recipe starting from another recipe. Three types of formulation rules have been defined: Substitution, Increase in quantity or Decrease in quantity.

Compounds adaptation follows the application of ACM rules. The knowledge base has been partitioned into *knowledge sources* corresponding to ACM rules. As shown in Figure 3, knowledge sources activation starts from the application of Description Rules that split the coded batch representing the solution for the retrieved case into the quantities of its ingredients. This step invokes the integration interface to the Pirelli Tires Archive. Then, the Performance-Properties Rules knowledge source is activated, in order to determine the needed properties of the product starting from the required performance. From the global properties of the product, the Ingredients-Properties Rules knowledge source finds out which ingredients have to be considered in order to obtain a variation of the properties satisfying the required performance. Finally, the Formulation Rules knowledge source formulates the modified recipe applying Substitution, Increase in quantity or Decrease in quantity Rules. Interested readers can found in (Bandini and Manzoni, 2001b) more details about the Abstract Compound Machine model and its use within the P-Race project.

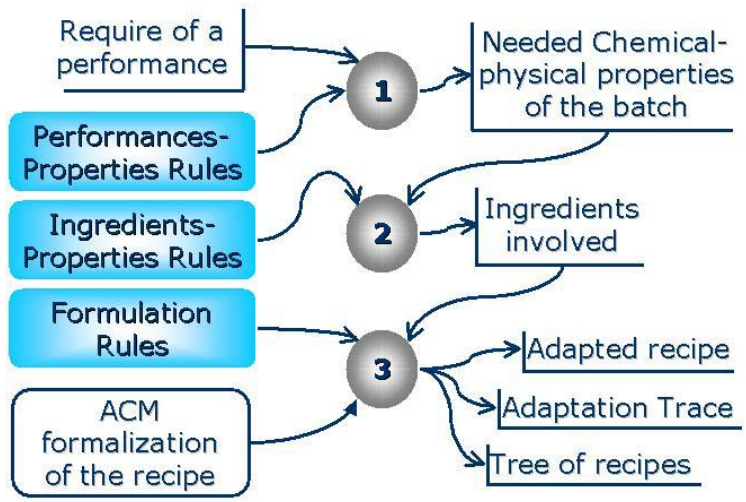


**Figure 3.** Product Innovation occurs according the application of ACM Rules

## The P-Truck Project

P-Truck is an ongoing project that has born after the success of the P-Race project. Its aim is to develop a KM system to support the Business Unit Truck of Pirelli Tires in order to improve both design and manufacturing activities of truck tires.

### *Truck Tire Manufacturing*

The life-cycle of a truck tire can be divided into the following main phases:

- *Design of rubber compounds*: a rubber compound is a blend of different ingredients, both natural (e.g. natural rubber, resins) and synthetic (e.g. carbon blacks, oils). This design phase has to decide the blend composition, identifying a set of ingredients and their amount, in order to achieve the performances that are required for the blend and for the tire (e.g. tensile strength, resistance to fatigue).
- *Mixing*: the ingredients must be suitably mixed in order to obtain a homogeneous blend with the required viscosity (again, related to rubber compound and tire properties).
- *Semi-manufactured production*: reinforcements are added to rubber compounds, producing the different parts that will compose the tire.

- *Assembly*: semi-manufactured parts are assembled into a semi-finished product (i.e. *green-tire*, in the tire jargon).
- *Vulcanization*: the green tire is processed in order to give it the required thermal-mechanical features.

The above-summarized life cycle is sufficiently general to characterize all tire production realities, where different specific way to perform it can be possible. In the specific case of the Business Unit Truck of Pirelli Tires, a knowledge acquisition campaign has revealed that each of the above-summarized phases is accomplished by a specific Community of Practice (CoP). One of the main aims of the P-Truck Project is, thus, to support the CoPs that are involved in the Design of rubber compounds (i.e. compound designers), Mixing (i.e. Mixing technologists) and Vulcanization (i.e. Vulcanization technologists). Specific KBS solutions have been designed to support these COPs, and each one has been designed according to the specific design style of the COP.

Knowledge acquisition with Pirelli designers' COPs revealed that a truck tire can be considered as a *chemical device*, made up of both chemical components and other elements. Different from chemical commodities, whose life cycle is focused on the optimization of manufacturing process, the one of a chemical device is centered on the product innovation, in order to meet the requirements of evolving markets they are devoted to. In the case of truck tires, it is necessary to optimize a lot of performances (e.g. tensile strength, resistance to fatigue), and the importance of them varying according to the target market (e.g. South America, Europe, Asia).

The life cycle of a product in manufacturing companies consists of several phases that can be divided into two main steps: design and production. The design step is usually triggered by the need of innovation defined by marketing strategies oriented to answer requests emerging from the market or from new needs induced in it. The goal of the design step is to define product and production process specifications. Product and process specifications define the constraints the production process has to respect in order to make products pass the quality control and be successfully distributed on the market. The P-Truck project focuses in particular in the support of those people that are involved in tire design and innovation.

However, in the manufacturing reality another step has often to be considered in the design-production cycle (i.e. production tuning). This need is due to the possibility to encounter unpredictable problems during the production step. Like any sort of specifications, product and process specifications cannot be exhaustive of all the possible details characterizing the instantiations of the production process. In fact, design is usually a process located where the core design competencies reside, while production is increasingly distributed in different plants to take advantage of local facilities or to make production scheduling more flexible to the market needs. Each instantiation of the production process happens in a different contingent situation where, for instance, raw materials can show even slightly different properties. Different production context might lead to variable or unexpected results on the final product and to the consequent modification of the product structure or composition.

The tuning of product and process specifications in order to adapt them to take care of structural properties of the local plant, of unanticipated events or non-uniformity in raw materials may be necessary in order to guarantee the necessary product quality (e.g. uniformity of production results). Thus, besides standard process design, process tuning has sometimes to be performed during manufacturing. Although these critical situations (also referred as, *production anomalies*) are not too frequent, they are extremely crucial since the time available to solve them and to continue the production under the new circumstances is very short. The search for a quick solution justifies to record anomalies connected to problems that have already been experienced and solved in the past.

Knowledge engineering with Pirelli technicians involved in product and process tuning (i.e. people working in the production line, typically experienced workers who are however very busy or involved in off-line production activities, like product testing) revealed that the nature of their problem solving activity is mainly episodic, that is based on previous similar situations: they often reason about past cases in order to solve the current one. Thus, Case Based Reasoning has been chosen as a suitable approach in order to captures the episodic knowledge characterizing most of the reasoning activity of people involved in production tuning, and to support the dynamical process of experience growth and knowledge creation through incremental learning. A dedicated Case Based Organizational Memory (CB-OM) has been designed for the Business Unit Truck of Pirelli Tyres. Interested readers can refer to (Manzoni and Mereghetti, 2002) for more details about the content and structure of the Pirelli CB-OM.

### *An Overview of the P-Truck System*

The architecture of P-Truck has been designed in order to have a *centralized knowledge repository* and a *distributed problem solving strategy*. The designed architecture provides a lot of benefits, in particular, from the knowledge maintenance and sharing viewpoints.

Figure 4 shows the high level architecture of the P-Truck system that can be divided into two main parts: a first part is devoted to capture, represent and store expert knowledge (i.e. KEPT, Knowledge Elicitation module of P-Truck), and another one is devoted to process this knowledge (i.e. KPM, Knowledge Processing Modules).

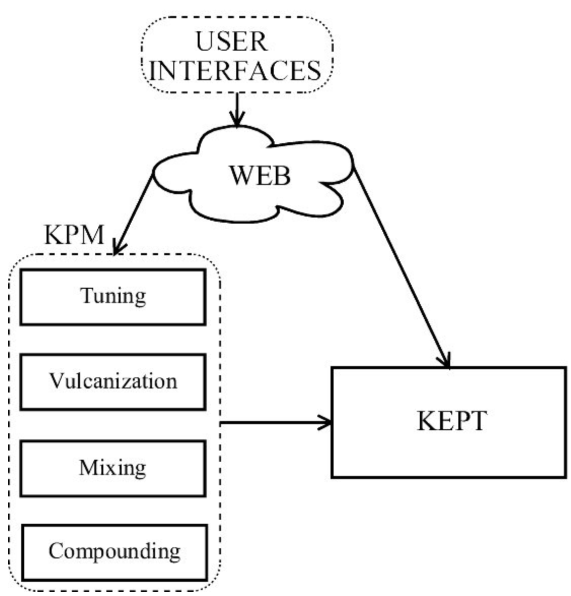


**Figure 4. The P-Truck System**

In particular, the KPM is made up of four Knowledge Based Systems:

- Compounding, a rule-based system that supports compound designers;
- Mixing, a rule-based system that supports Mixing technologists;
- Vulcanization, a case-based system that supports Vulcanization technologists;
- Tuning, a case-based system that support those experts whose aim is to handle possible anomalies that may occur during the production phases.

From the knowledge maintenance point of view, the development of KEPT has allowed to provide experts with knowledge repository and a flexible tool to manage it (i.e. visualize, update and retrieve). It is evident that a single knowledge base that collects data and knowledge coming from different sources is preferable from this point of view. This solution allows avoiding problems that must be considered dealing with distributed knowledge bases (e.g. data inconsistency). A centralized knowledge base facilitates all activities that require accessing it and, at the same time, it is an advantage also from the knowledge sharing point of view. In fact, this solution allows providing different KPM components with a specific view on the knowledge base that they can access to perform their tasks. Finally, new Knowledge Based Systems that exploit the P-Truck knowledge base content can be simply added and provided with an appropriate view on the knowledge base content.

Another important feature of the KEPT knowledge repository concerns its internal partition into two main parts. Each partition serves a different type of knowledge-based KPM. In particular, a first partition represents knowledge that is exploited by rule-based modules, and the other one constitutes the case memory of case-based modules. Within the P-Truck Project particular emphasis has been dedicated to the definition of the most suitable KBS approaches to support the different CoPs involved in the tire production process. A rule-based approach has been adopted for dealing with the decision making process of compound designers and mixing technologists, while the tuning of production process and the support of vulcanization technologists have been managed through a CBR approach. Different choices have been motivated by the different nature of decision-making processes that had to be supported by the KPM modules. At one hand, the design of rubber compounds is based on well-known chemical relationships between ingredients, blend properties and tire performances. Those relationships can be adequately captured, represented and stored into a rule base. Analogously, the design of mixing processes follows well-known and experienced rules that bind compound ingredients and machinery features. On the other hand, there are no explicit relationships between production anomalies and their solutions. Production experts, due to scheduling and costs requirements, often reason by analogy, and apply to anomalous production situations solutions that have already been experienced in the past.

The P-Truck system has been designed and developed according to a *Component Based* approach. P-Truck modules have been designed in order to satisfy all the identified user and non-user requirements (i.e. integration of heterogeneous and interacting components, independent development of components, extendible architecture, separation between business and presentation logic, portability, integration into the existing information system and Web orientation). Figure 5 shows two examples of the web based user interface of the P-Truck system.

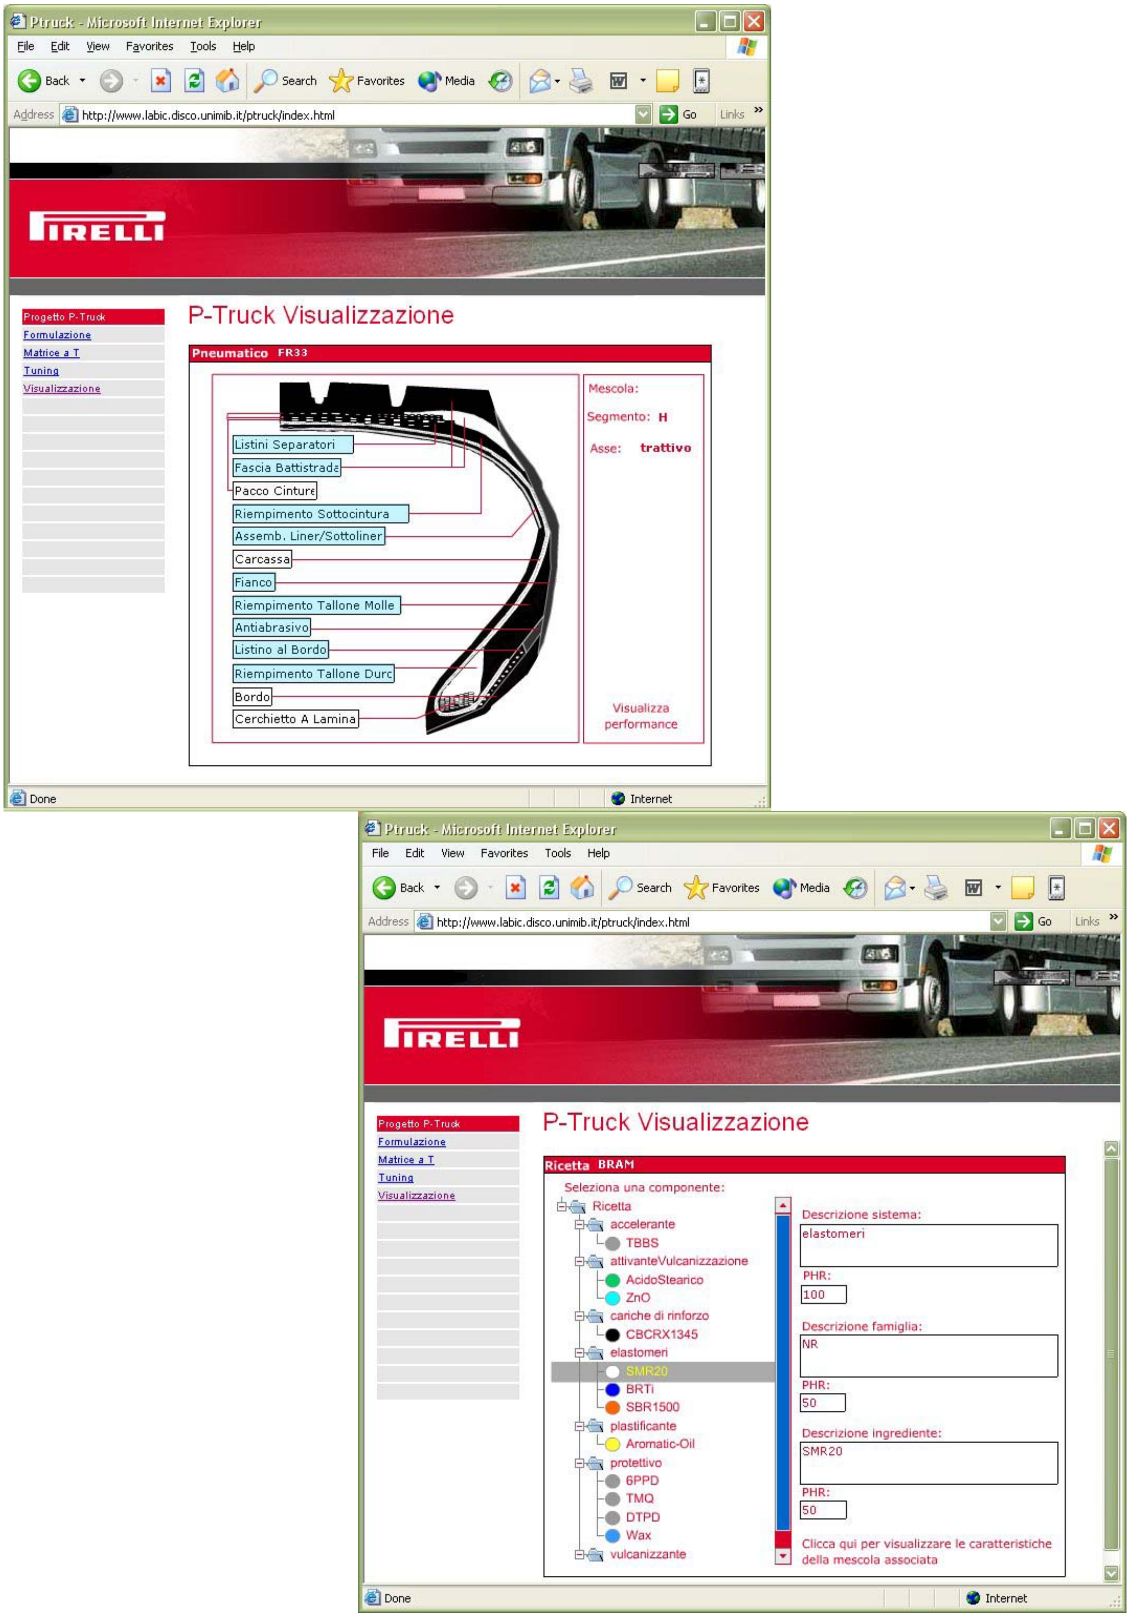


**Figure 5.** P-Truck system interface: tire components' and batch recipe view.

System requirements lead to the choice of the Java 2 Enterprise Edition (J2EE) platform as technological solution to develop the P-Truck system. Currently the J2EE technology supports component-based development (Enterprise Java Beans, EJB), integration between systems (it complies with CORBA and Remote Method Invocation, it supports XML, and so on), database accessibility and web-based user interfaces development. Moreover, the designed solution allows the integration of several knowledge-based modules based on different approaches (e.g. rule based and Case Based Reasoning) into an already existing and used company information system (made of a huge number of databases and applications).

### *Concluding Remarks*

The process of acquiring and modelling core knowledge concerning a specific domain is a very important research topic. Many Knowledge Based Systems (KBS) have been developed to deal with several knowledge fields, but the phase of knowledge acquisition and representation is still the main problem of this type of tools.

Knowledge has been recently recognized as a very important asset for enterprises, needing knowledge engineering to:

- identify knowledge sources (people, data sets, texts, programs and so on);
- build an incremental knowledge base for the acquisition and representation of knowledge;
- share and reuse knowledge among different applications for various types of users (i.e., share existing knowledge sources and future ones).

Knowledge engineering methodologies, such as CommonKads and MIKE, have been proposed as standard and generalized solutions to satisfy enterprise needs. Another possible approach to knowledge engineering consists in the development of dedicated KBSs providing specific solutions to each problem; in this case, domain specific knowledge acquisition and representation tools should be adopted.

This type of approach has been adopted within both the presented project: P-Race and P-Truck and the KARM and KEPT tools have been, respectively, developed. The KARM module provides the P-Race system with a tool for the acquisition and representation of either qualitative or quantitative knowledge. In the former case, the block view technique allows race engineers to describe track profile features and, thus, to represent their core knowledge. Within the P-Truck project, KEPT has allowed to provide experts with a flexible tool to manage data and knowledge coming from different, possibly distributed, sources.

The integration of KARM and KEPT with knowledge processing modules (of P-Race and P-Truck systems, respectively) allows the acquired and represented knowledge to be used and exploited to support system users in KM. Thus, for instance, the experiential knowledge of race engineers can be profitably shared and used by all the members of the motor racing team. Similarly, the experience and core knowledge of compound, mixing and vulcanization designers can be exploited within the COP that owns it and can be shared with who perform product and process tuning.

Finally, once this core knowledge has been acquired and represented into the knowledge repository of a KM system, tools like KARM and KEPT can be exploited as training tools. For instance, in the case of P-Race, the block view description of a track eventually provided by beginners could be compared with the one supplied by P-Race Telemetry. Within the P-Truck project, that will last at the end of the following year, a dedicated module will be designed and develop specifically for this task.

## Some additional references on the presented case studies


Bandini, S. and Manzoni, S. (2000): *A Support System Based on CBR for the Design of Rubber Compounds in Motor Racing*. In Blanzieri, E., Portinale, L. (Eds.): Advances in Case-Based Reasoning, LNCS 1898, Springer 2000, ISBN 3-540-67933-2, pp. 348-357.

Bandini, S. and Manzoni, S. (2001a): *Application of Fuzzy Indexing and Retrieval in Case Based Reasoning for Design*. In Proceedings of the 2001 ACM Symposium on Applied Computing (SAC), Neuro-Fuzzy Applications Track, March 11-14, 2001, Las Vegas, NV, USA. ACM Press, New York, pp.462-466.

Bandini, S. and Manzoni, S. (2001b): *CBR Adaptation for Chemical Formulation*. In Aha, D.W., Watson, I., Yang, Q. (Eds.): Case-Based Reasoning Research and Development, LNCS 2080 Springer 2001, ISBN 3-540-42358-3, pp. 634-647.

Bandini, S. and Manzoni, S. (2002): *Modeling Core Knowledge and Practices in a Computational Approach to Innovation Process*. In Magnani, L. (Ed.), Model-Based Reasoning: Scientific Discovery, Technological Innovation, Values, World Scientific, Singapore.

Manzoni, S. and Mereghetti, P. (2002): *A Tree Structured Case Base for the System P--Truck Tuning*. In Proc. of UK CBR workshop at ES 2002, Cambridge, 10 december, 2002, University of Paisley, Glasgow, pp. 17-26.

S. Bandini, S. Manzoni, F. Sartori (2002): *Acquiring Knowledge and Numerical Data to support CBR Retrieval*, Proc. of the 13th International Conference on Knowledge Engineering and Knowledge Management, European Knowledge Engineering Workshop (EKAW02), 1-4 October 2002, Siguenza (Spain), LNAI 2473, pp. 8-13.