On Refining Ill-Defined Constraint Problems: A Case Study in Iterative Prototyping

Carmen Gervet, Yves Caseau, Denis Montaut

On Refining Ill-Defined Constraint Problems: A Case Study in Iterative Prototyping

Carmen Gervet
IC-Parc, Imperial College
London SW7 2AZ
c.gervet@icparc.ic.ac.uk

Yves Caseau
Bouygues, D.T.N.
1 av. E. Freyssinet
78061 St Quentin en Y.
caseau@dmi.ens.fr

Denis Montaut
Euro-Decision
9 rue de la porte de Buc
78000 Versailles
dcm@eurodecision.fr

Abstract

In this paper we describe an iterative development process to deal with speculative constraint optimization projects. Speculative projects are ill-defined in nature, as they are often new to the customer organization who wishes to anticipate market changes, but also because their main complexity lies in completing the problem definition. We consider one speculative application tackled in the `Chic-2` Esprit project by different partners: Risk management in the area of energy trading[1]. The ill-defined and incomplete components of the problem are: 1) model adequacy, 2) multi-criteria objective, and 3) forecast data. We show how the constraint technology can help refining the problem definition by building iterative models and methods to address the incomplete parts incrementally. Each partner applied his own technology using constraint programming, mathematical programming and local search techniques. Finally we discuss the new potentials for constraint technology to deal with stochastic problems.

1. INTRODUCTION

Constraint technology with various forms of hybridization has shown in the past decade its flexibility and efficiency to tackle well-defined projects, where the main objective is to model an existing problem and to build efficient solution methods. However, little has been published on using the technology to cope with *speculative* projects where an additional complexity lies in handling the ill-defined components of the problem definition. Speculative projects appear for instance when the customer organization wishes to anticipate market changes, and wonders to which extent the current approach should be revisited and how. In such cases, users expect our technology to help with identifying new issues based on their current experience and knowledge, and to complete the problem definition by providing models and solutions. The standard methodological process consisting of problem definition, design and solving is inadequate in this situation because the main goal is to refine the problem definition as opposed to solely solving it.

[1] `Chic-2` focuses on Creating Hybrid solutions for Industry and Commerce. Its activities concern the tackling of industrial Large Scale Combinatorial Optimization applications using hybrid algorithms, and the building of a methodology for LSCOs. See http://www.icparc.ic.ac.uk/chic2/.

This is quite a new challenge for constraint technology, and we propose to address it using an iterative prototyping approach.

In this paper we describe a speculative project tackled within the `Chic-2` Esprit project in the past two years: risk management in energy trading. Currently, electricity suppliers sign energy contracts with generator companies to satisfy best the end-consumers' demand and minimize the energy cost. With the coming deregulated market in England and Wales, suppliers foresee an increased degree of financial risks coming from an increased volatility in electricity prices (i.e. the forecast price data will be more uncertain). This will require suppliers to change fundamentally their current approach in order to assess and control the new risks associated with uncertain forecast data. Energy suppliers wish to anticipate the market changes by building portfolios of energy contracts that will maximize potential profits while minimizing the risks. The main technical issues amount to: 1) evaluating the consistency of the current constraint model, 2) deriving means to refine and possibly complete the problem definition. The incomplete components are the cost function which involves several criteria including a measurement of risk, and the forecast data for which no reasonable probability distribution can be secured a priori.

We present an *iterative prototyping* development process to evaluate the current model and refine the problem definition. Each successive prototype is developed with a specific purpose. A new prototype is built from the previous ones by adjusting the model of the problem and tackling a new ill-defined component. Starting from a simple model to formulate the core problem, the prototypes iterate towards more complex models that address the multi-criteria issue and the uncertainty in forecast data. For each technical issue, we propose modeling and visualization means to tackle them in the most realistic manner.

Each application addressed within the `Chic-2` Esprit project was tackled by several project partners for a fixed period (one to three months). The objective was to explore different modeling and algorithmic approaches based on a common problem definition. With respect to the Energy trading application, IC-Parc was the main partner and has been working with ICL, Euro-Decision, and Bouygues. In this paper we present each partner's solution to address the various ill-defined components. The main contributions arising from this case study are: 1) the active role of constraint technology in iterative prototyping to handle speculative projects, 2) the use of goal programming techniques from Operations research to deal with over constrained problems, 3) the role of visualization tools to deal with multi-criteria optimization, and 4) the introduction of simulation models in constraint programming to refine the problem definition and deal with uncertain forecast data.

This paper is structured in four stages that reflect both the different working partnerships with IC-Parc, and the iterative prototypes we have implemented. Section 2 presents the definition of the problem and the business case. Section 3 presents the first prototype to model and solve the core problem. Section 4 describes the second prototype to address the multi-criteria aspect. Section 5 presents a simulation model that deals with the stochastic nature of the problem. Finally we conclude with the lessons we have learnt form this application and draw some research perspectives.

2. PROBLEM DEFINITION

2.1 Business case

The electricity chain works as follows in England and Wales: power generators provide electricity to the national electricity Pool, and electricity suppliers purchase electricity from the Pool in order to satisfy the demand from their end-consumers (see Figure 1). As a trading practice, the price of electricity is set by the Pool and varies every half hour. The electricity price is calculated according to the maximum amount of electricity consumed by the end-consumers for every half hourly period. End-consumers' demand can also be unpredictable due to factors such as weather changes, economic growth, special events or simple bad forecasting. So pool prices are extremely volatile and unpredictable. Consequently, the risk of Electricity Suppliers not being able to make any profit (when buying and selling electricity) is very high.

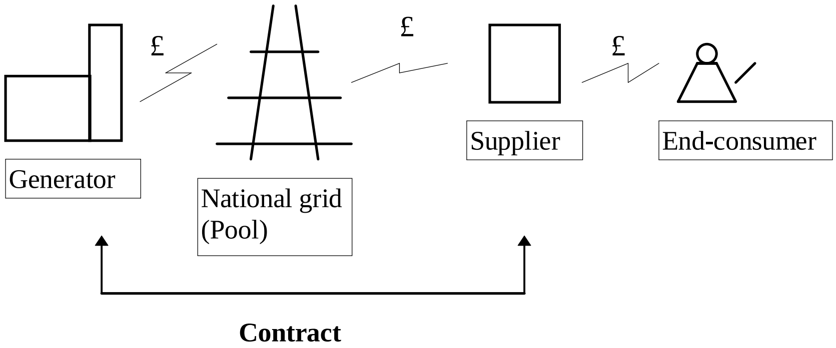

Figure 1: Electricity supply and trading chain

To protect themselves from the Pool price fluctuations and volatility, electricity Suppliers and Generators use a financial instrument, i.e. a contract which effectively fixes the price of electricity for the purchase of a pre-set volume of electricity over a given contracting period. Currently, Suppliers work on deriving optimal portfolios of contracts (between electricity suppliers and generators) that cover their forecast demand for a period of 12 rolling months such that their financial risks are minimized and their potential profits are maximized. The actual measure of the financial risks is based on two criteria illustrated below:

- The *demand exposure*. Sum of all distances between the demand forecast and the contracts' cover at any one time.
- The *expected energy cost*. Sum of all differences between the Pool price forecast and the total cost of the expected values of selected contracts.

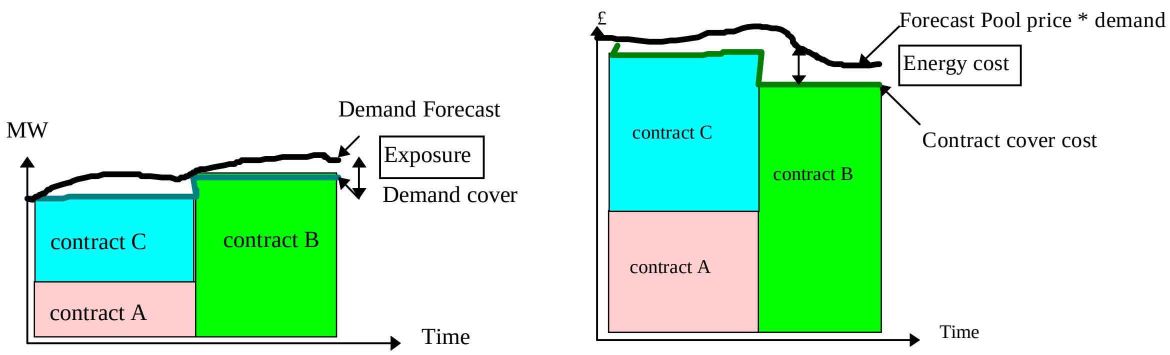

Figure 2: Demand exposure and expected energy cost

When the `Chic-2` project started, the deregulation of the energy trading market in England and Wales was planned, enforcing energy Suppliers to become energy Traders as well. Deregulation would have various consequences on the way current Suppliers work:

1. The forecast trends and seasonal patterns in energy costs would become more *volatile* as they would be subject to the financial market fluctuations as well as the actual energy demand market. The market becomes more fluid and players more adventurous. New risk factors take place and the selection of optimal contracts to hedge the risks is a much more difficult task.
2. In a non deregulated market, suppliers can secure contracts with energy generators for a long period (e.g. over 1 year) and are given time to select and negotiate the contracts most suited to their demand. After deregulation, traders will have to *react much faster* to the changes.
3. For the energy suppliers involved in the `Chic-2` project, the building of contract portfolios is being done by manual effort with the aid of spreadsheets. However, they have already reached the limit of spreadsheets tools. Deregulation requires an increase in number of contracts to be considered at any given time which means that a new model and better methods need to be found before the market is fully deregulated.

2.1.1 Problem specificity

There exist various software and stand-alone packages for risk management and portfolio optimization mostly based on the seminal work of Markowitz [13] that relate to decision theory and utility theory. However, risk management in the electricity trading market is a unique problem in many ways:

- Electricity is not a storable commodity, supply is always equal to demand.
- Electricity price set by the Pool change every half hour. The fluctuation trends and seasonal patterns do not match with any other commodity.
- The supply is continuous every hour, every day, every week and every year.

Our users, i.e. energy suppliers, did not find any software or stand-alone package addressing the energy trading problem, mainly because they are not looking for fully optimized solutions but for a Decision Support System (DSS) telling them “how to revise today’s best practice to cope with the coming deregulated market”. Consequently, the problem as specified by the users is *ill-defined* for different reasons: it is a new problem to them, the forecast data are uncertain, and the financial risks might involve new factors.

2.1.2 Expectations of the users in the decision support tool

For the near future, the main interest of the users lies in a bespoke DSS that considers their current practice and tunes or revises it towards a new model and solution to risk management. For instance, they would like to know to which extent the current rules of the company governing portfolio management are best viewed as soft or hard rules (constraints). In addition, the users expect to use the DSS to make informed judgment on which contracts to choose best depending on their risk adversity and confidence in data forecast. The criteria to be optimized for the portfolio selection should ideally derive a new measure of the financial risks.

2.2 Problem formulation

The initial user requirements are defined as follows. Given:

- Electricity demand and price forecast for every half hourly period of a rolling year (12 months starting from September)
- A set of contract profiles with a known shape specified by its baseload unit, duration (e.g. one month, six months, one year), and unit price.
- A set of company rules (constraints)

Minimize the demand exposure and expected energy cost (current financial risks) of the company by building a portfolio of contracts for the next 12 months such that the company rules are

satisfied. The unknowns are the selection (Yes/No) of contract profiles and the volumes attached to the contracts. A contract volume commits the baseload unit to an integer value.
The company rules are a set of constraints defined by:

- The total volume covered by each selected contract should not exceed 500GWh
- The total cost of each selected contract should not exceed 15 million pounds
- The volume of a contract is either nil or greater than 10MW
- The total cover of contracts should remain under the electricity demand curve or 10% above. Even though we have studied both cases, for reasons of space, we will present the results for the under cover case only. Results and comparisons between both cases can be found in [6].

Remark: A baseload unit gives the unitary shape of a contract, and is (most commonly) built on time periods divided into week days (WD) and week-end days (WE) where each day is itself divided into 4 hourly periods (i.e. 6 * 4). A four hourly period is called an EFA period. This terminology will be used throughout all the problem design and solving stages. For example a contract profile that covers 4 hourly periods 3, 4 and 5 during Monday to Friday is defined by the following loadshape, where *x* is the volume (decision variable) to instantiate:

WD345

EFA	WD1	WD2	WD 3	WD4	WD5	WD6	WE1	WE2	WE3	WE4	WE5	WE6
MW	0	0	x	x	x	0	0	0	0	0	0	0

2.3 Algebraic model

The algebraic model structures and formulates the user requirements around the following items: sets and indices, input data, constraints, decision variables and cost function [3]. It is presented below.

2.3.1 Sets and indices

Let $P = \{p_1, \ldots, p_n\}$ be a set of profiles (load shape, duration, strike price).
Let $C = \{c_1, \ldots, c_n\}$ be a set of contracts (profiles + fixed volume).
Let $M = \{m_1, \ldots, m_{12}\}$ be all the months in the rolling year.
Let $I = \{i_1, .., i_{12}\}$ be a set of 12 EFA time period indices, where the first six indices denote EFA periods for week days and the last six, EFA periods for week-end days.

2.3.2 Input data

- Relative to a profile `p`

`price(p)` : initial price of the profile.
`name(p)` : name of the profile.

- Relative to a contract `c`.

It contains the information inherited from the corresponding unique profile.
`duration(c)` : list of months covered by the contract.
`str_price(c)` : strike price of a contract. It is a piece-wise linear function defined by the initial profile price plus five pence per 50 MW.
`tot_cover(c)` : total number of base loads covered by the contract (computed using the EFA system scaling).
`shape(c,i)` : base load for the EFA period `i` in the WD/WE load shape ($i \in \{1, \ldots, 12\}$).

- Relative to demand forecast

The input data for demand are average data that transform the 365 x 48 matrix into a 12x12 matrix. This is the current users' practice that considers the average demand for a week per month and divides the 48 half hourly periods into 6 * 4 hourly periods. A distinction is made between week days and week-end days to match with the contract profiles.

`demAvr(m,i)` : average demand. It represents the average demand per day in a month *m*. A WD can be accessed using indices i in $\{1,\ldots,6\}$, and a WE using indices i in $\{7,\ldots,12\}$.

`demTot` : total demand over the whole financial year.

- Relative to pool price forecasts

`pppAvr(m,i)` : average Pool Purchase Price forecast data for each EFA period `i` ($i \in \{1,\ldots,12\}$) of month m.

`pspAvr(m,i)` : average Pool Selling Price forecast data for each EFA period `i` ($i \in \{1,\ldots,12\}$) of month m.

- Relative to the constraints

`maxVol` = 500 GWh, maximal volume of cover allowed per contract.

`distDem` = 500 GWh, maximal difference between total contracts' cover and total demand.

`maxCost` = 15 M£, maximal gross value (volume times strike price) of a contract.

2.3.3 Decision variables

cvol(c) : volume of a contract **c**, initialized to the semi-continuous interval [0,10,..,500].

$$= \begin{cases} 0 \text{ if the contract is not selected} \\ \geq 10 \text{ otherwise } (\leq\ 500) \end{cases}$$

bool(c) : boolean variable indicating whether a contract **c** is selected or not.

2.3.4 Constraints

The constraints are the current rules of the company. They are hard constraints unless specified otherwise.

1) The volume of a contract is either 0 or greater than 10.
   $\forall c \in C: 10 \times \textbf{bool}(c) \leq \textbf{cvol}(c)$
   $\forall c \in C: 100000 \times \textbf{bool}(c) \geq \textbf{cvol}(c)$
2) The volume covered by each selected contract should not exceed 500 GWh.
   $\forall c \in C: \texttt{tot\_cover}(c) \times 4 \times \textbf{cvol}(c) \leq \texttt{maxVol}$
3) The gross value (strike price times volume covered) of each selected contract should not exceed 15 million pounds.
   $\forall c \in C: \texttt{tot\_cover}(c) \times 4 \times \textbf{cvol}(c) \times \texttt{str\_price}(c) \leq \texttt{maxCost}$
4) The company is under-covered against demand by no more than 500 GWh in the rolling 12 months.
   $$\texttt{demTot} - \texttt{distDem} \leq \sum_{c \in C} \sum_{w \in W | w \in duration(c)} \texttt{tot\_cover}(c) \times 4 \times \textbf{cvol}(c) \leq \texttt{demTot}$$
5) The company is left under-covered against demand by no more than a given percentage in the rolling 12 months (80-100%).
   $$0.8 \times \texttt{demTot} \leq \sum_{c \in C} \sum_{w \in W | w \in duration(c)} \texttt{tot\_cover}(c) \times 4 \times \textbf{cvol}(c) \leq \texttt{demTot}$$

6) The company should be covered by a given percentage of the 80%-100% of the demand forecast *for each EFA* of each averaged monthly week.

$$\forall m \in M, \forall i \in I,$$

$$0.8 \times \texttt{demAvr}(m,i) \leq \sum_{c \in C | m \in duration(c)} \texttt{shape}(c,i) \times \mathbf{cvol}(\mathrm{c}) \leq \texttt{demAvr}(m,i)$$

2.3.5 Objective function (criteria to optimize)

The current criteria to minimize are:

- The expected energy cost over the financial year.

$$f\mathrm{energy\_cost} = \sum_{\mathrm{m} \in M} \sum_{\mathrm{i} \in \mathrm{I}} \texttt{demAvr}(m,i) \times \texttt{pspAvr}(m,i) - \sum_{m \in M} \sum_{i \in I} \Big(\sum_{c \in C | m \in duration(c)} (4 \times \mathbf{cvol}(\mathrm{c}) \times \mathbf{bool}(c) \times \texttt{shape}(c,i)) \times (\texttt{pppAvr}(m,i) - \texttt{strike\_pr}(c) \Big)$$

- The demand exposure.

$$f\mathrm{demand} = \sum_{m \in M} \sum_{i \in I} \left| \texttt{demAvr}(m,i) - \sum_{c \in C | m \in time(c)} (\mathbf{cvol}(\mathrm{c}) \times \mathbf{bool}(c) \times 4 \times \texttt{shape}(c,i)) \right|$$

3. FIRST PROTOTYPE

Our first objective is to ensure that the constraint set is consistent and that the solution produced is realistic to the users. This stage has been addressed by Euro-Decision and IC-Parc with slightly different implementations of the algebraic model.

3.1 Deterministic model for the core problem

In order to focus on the feasibility check, we consider one main decision criteria to optimize: the demand exposure. The reason is that we do not know yet about any correlation between the demand exposure and the expected energy cost. Thus by working with one criterion we hope to make sure with the users that the optimized solution is accurate and realistic.

3.1.1 Assumptions made

The first prototype is built using the following assumptions:

- Consider a set of 20 different profiles (loadshape, duration, price)
- Consider all the constraints given by the user
- Consider one demand and price forecast at a time
- Minimize the demand exposure mainly

3.1.2 Model for the objective function

The demand exposure is a sum of distances between demand forecast and contracts' cover, i.e. a sum of absolute values[2]. An absolute value can be modeled either using two inequality constraints or by introducing a difference of two positive variables and minimizing their sum. We used the second alternative since it considers twice less constraints. Even though it introduces new variables, it maps better with the algorithm we used. This corresponds to the following formulation:

[2] In the under cover case the demand should always be greater than the contract cover thus the difference must be positive. But since we also considered the over case where the difference can be negative, we defined a generic model for the demand exposure.

$\forall m \in M, \forall i \in I,$

$$\mathbf{goalPlus}(m,i) - \mathbf{goalMinus}(m,i) = \mathtt{demAvr}(m,i) - \sum_{c \in C | m \in duration(c)} \mathtt{shape}(c,i) \times 4 \times \mathbf{cvol}(c) \times \mathbf{bool}(c)$$

The function to minimize is $\sum_{m \in M} \sum_{i \in I} \mathbf{goalPlus}(m,i) + \mathbf{goalMinus}(m,i)$.

3.1.3 Two different models

A study of the mathematical features of the algebraic model showed us that the problem formulation describes a pure integer linear programming model.

IC-Parc's model. IC-Parc first coded the algebraic model as a Constraint Satisfaction Problem (CSP) [12] using the ECL^{i}PS^e II platform [10]. The objective was to investigate how good CP can perform on this type of problems. The decision variables range over discrete domains and correspond to the contract selection variables (in [0,1]) and the contract volumes (in [0,10,..500]). The objective function is to minimize the demand exposure.

Euro-Decision's model. Euro-Decision coded the algebraic model as an Integer Linear Programming (ILP) model using the LP-Toolkit[3] system. The objective function is a weighted sum of demand exposure (high weight) and expected cost (small weight). Since a contract cost depends on its volume, this approach leads to a non-linear expected energy cost function. In order not to handle non-linear functions, Euro-Decision created a piece-wise linear function by generating an a priori set of contracts for each profile (10) with an associated volume and fixed cost.

3.2 Resolution algorithms

3.2.1 Using ECL^{i}PS^e

IC-Parc checked the consistency of the constraint set by applying constraint propagation and a simple labelling procedure on the contract selection first. *Infeasibility* of the constraint set was detected during the constraint propagation phase. It appeared clearly that some of the constraints 6) alone were inconsistent. We relaxed the lower bounds imposed by constraints 6) on the coverage of contracts, the problem became feasible. IC-Parc then searched for optimal integral solutions using branch and bound. This turned out to be very inefficient unless we devised powerful heuristics based on the expected value of the contracts and their coverage power. However, due to the linear nature of the model, IC-Parc decided at this point to use a linear solver. The model presented previously was then solved using the Mixed Integer Linear Programming library of ECL^{i}PS^e (CPLEX [4]). We obtained in 3-4 seconds an optimal solution to the relaxed problem on a Sparc Server 2 170 (two 167 Mhz CPU's) depending on the data set. The relaxed problem consists of allowing the contract's volumes to take real values. The computed solutions correspond to an MILP model where the boolean variables take discrete values and the contract volumes take real values.

The solver failed to find optimal integer solutions for contract volumes in a limited amount of time. However, users mentioned at this stage, that the integrality constraint over the contract volumes was not a hard one. In practice traders compute real values for the contract volumes and round them to get integer values which are easier to deal with their clients.

---

[3] Modelling language of Euro-Decision, see http://www.eurodecision.fr

3.2.2 Using LP-Toolkit

Euro-Decision used a Mixed Integer Programming solver (OSL) to solve the feasible model. The integral case was studied by considering a preprocessing model where 10 contracts were generated for each profile, and their respective strike price were fixed and totally ordered. In this model the energy cost is part of the objective function, so the solutions obtained differ from IC-Parc ones. The selection of contracts attached to a same profile, is such that one contract is selected only if the volume of the previous contract (in price ordering) has reached its maximum value (50MW). This is done automatically thanks to the global behaviour of linear optimization. Although this approach requires the handling of a predefined number of contracts per profile, it allows us to have an *exact* pre-computation of the expected value of each contract (used in the expected energy cost criterion).

On a SunSparc 10 workstation, solving the relaxed problem (volume variables take “real” values) took 2-4 seconds. An integer solution was available after about 3 minutes using branch and bound search (without proof of optimality). It was less than 0.015 % far from the optimum of the relaxed problem.

3.3 Result & Analysis

The initial problem definition had no feasible solution because the problem was over-constrained. Indeed, the lower bounds of the constraint set 6) required the contract cover to be at any time at most at 80 % of the demand curve. In fact, high peaks of the demand curve could never be 80% covered by the contracts available due to their very flat shapes (as opposed to the high fluctuations in the demand curve). The solution to this problem was either to relax the lower bound constraints (into soft constraints) or to allow for new profile shapes that would be more flexible. Both approaches have been implemented but we only presented the first one.

Constraint relaxation has been achieved by introducing *goal variables*. The modeling technique called “Goal programming”[2] comes from Operations research and aims at solving over-constrained problems as *unconstrained* ones where the objective is to try to satisfy a set of soft constraints as much as possible. Goal variables introduce a distance that expresses the degree of satisfiability of a constraint. Minimizing the distance aims at reducing the degree of unsatisfiability of a constraint. In our application, the model for the relaxed constraint set 6) corresponded to the modeling of the demand exposure, where we introduced a goal variable in each constraint, and minimized their sum as part of the objective function.

Convincing the users that their initial problem had no feasible solution was not an easy task. But they finally acknowledged that the lower bounds on contract coverage imposed by constraints 6) corresponded to an ideal case. In practice Suppliers allow themselves to break the rules. Figure 3 shows a solution produced with a goal programming model. The irregular curve represents the demand forecast, and the more regular curve represents the cover of the selected contracts. The days are numbered from October 1 1998 to September 30 1999. Clearly some peaks in winter (e.g. days 40 to 100) could not be covered at 80% by the given contract profiles. The discontinuity at day 151 illustrates the shift between winter and summer in terms of contract profiles and demand forecast.

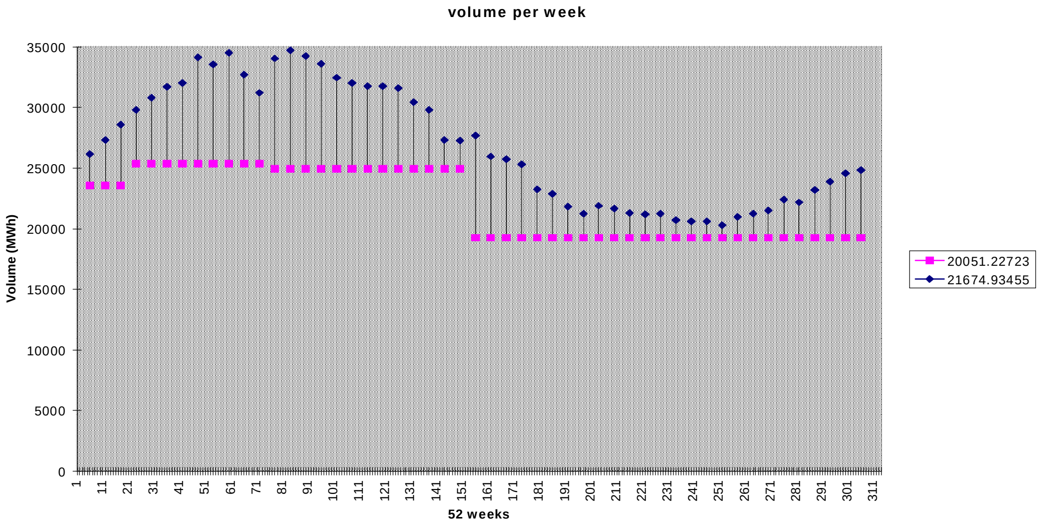

Figure 3: Solution with unconstrained lower bounds

4. SECOND PROTOTYPE

The objective of the first prototype was to investigate the feasibility and algorithmic issues. We defined and modeled the problem as a pure combinatorial search problem by making a set of assumptions that ignored the ill-defined aspects: the cost function was reduced to one main criterion (the demand exposure), and the data were considered known with certainty. The prototype gave us a lot of fruitful insight to the problem and showed that:

- The feasibility check helped the user seeing that a set of hard constraints needed to be soft constraints and that new profiles were required.
- The relaxed constraint set was easily solvable by mixed integer linear programming methods.

However, it was clear that the set of assumptions made was too restrictive with respect to the client's expectations. Also, it highlighted that the actual complexity of the problem lies in its *definition and modeling* rather than its solving. Consequently, the main direction for further work was to concentrate on the definition of the real problem and its realistic modeling. The main points to be addressed were the modeling of uncertainty relative to the input forecast data, and the modeling of a multi-criteria objective function. Both aspects are very complex issues in terms of modeling and solving (whether one uses mathematical programming or constraint programming). In the second prototype, IC-Parc decided to work on the multi-criteria component.

The simplest approach to model multi-criteria optimization problems is to define a unique cost function as a weighted sum of all the criteria. However, when the users do not know explicitly the correlation between the different criteria it is very difficult to weigh the criteria in a realistic manner. So IC-Parc decided to build a deterministic multi-criteria model by means of an interactive visual system. The objectives were 1) to study the correlation between two well-defined criteria: the demand exposure as a risk measurement and the expected profit, and 2) to visualize the solutions when tuning the cost function. This required an extension of the algebraic model. This section describes the demonstrator built by IC-Parc to address these objectives.

4.1 Multi-criteria model

Given:

- Electricity demand and price forecasts 12x12 average matrices
- A set of 20 contract profiles (duration, unit price, and load shape)
- A set of hard and soft constraints that set bounds on the local and global contract coverage,

Find the contract volumes that:

- Minimize the demand exposure,
- Maximize the expected profit (sum of expected value of each contract multiplied by the volume of cover).

The expected value of one contract is defined by the difference between the pool price forecast for each EFA period covered by the contract and the contract strike price. The algebraic formulation of expected profit per contract is:

$$\forall c \in C,$$

$$f_{profit}(c) = \mathbf{cvol(c)} \times \left(\frac{\sum_{m,i} (\texttt{pppAvr}(m,i) \times \texttt{shape}(c,i))}{\sum_{m,i} \texttt{shape}(c,i)} - \texttt{str\_price}(c) \right)$$

The total expected profit is the sum of each expected profit: $\sum_{c \in C} f_{profit}(c)$

The model considers alternatively one of the two criteria in the objective function and constrains the other one by setting bounds on its possible value.

4.2 Algorithm

The users provided us with a list of preferences to guide the decision making process. Even though it was difficult to classify these preferences in a mathematical or hierarchical way it appeared that the selection of contracts with positive expected value was appropriate. So we considered this well-defined heuristic.

We developed a graphical interface in TCL-TK [14] to make the selection of criteria interactive, and to visualize the solution produced. Via the demonstrator interface, users could select an objective function among 1) maximize the expected profit, and 2) minimize the demand exposure. They could set dynamically a global constraint to restrict the bounds on the non-optimized decision criterion.

Since the model has an MILP structure similar to the one solved in the first prototype, the MILP library of $\mathrm{ECL^iPS^e}$ proved to be ideal. Some graphical solutions are illustrated in appendix A.

4.3 Results

The results obtained provided us with good material to analyze the relations between the two criteria and their influence on the solutions produced. When optimizing on each criterion independently we obtained very different results that can be summarized as follows:

- Minimizing the demand exposure corresponds to a pure two dimensional *bin packing* problem (minimize the uncovered space by putting as many objects as possible in the space).
- Maximizing the expected profit corresponds to a pure *knapsack* problem (maximize the total profit by selecting as many contracts as possible that fit into the two dimensional space).

The actual problem is neither of them, but a compromise between packing and knapsack. However, since only the decision maker knows which combination is best, we generated a set of solutions based on one criterion while constraining the other one (and vice versa). This led to the drawing of an *efficient frontier* (function of one criterion against the other one) depicted in figure 4. It is a powerful means for the decision maker to make informed judgments.

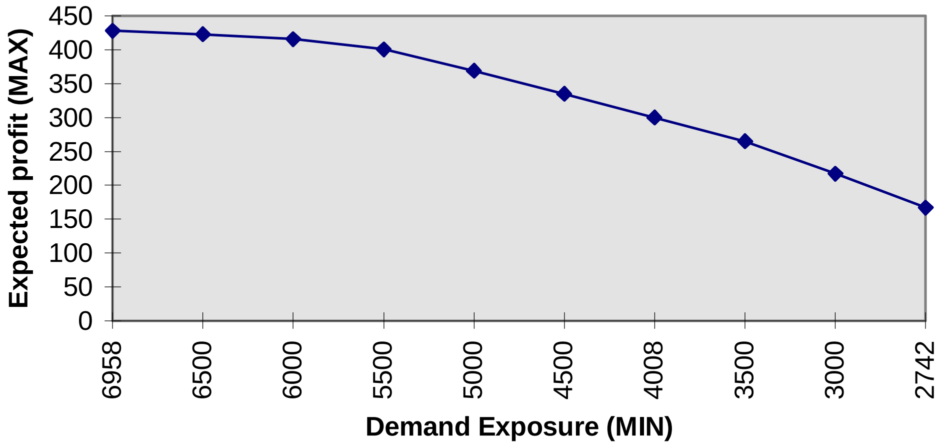

Figure 4: Efficient frontier

The figure shows that both criteria work against each other. The minimal exposure (2742 MW) corresponds to the worst profit, and respectively the optimal profit per unit (435 £/MW) corresponds to the worst exposure. An explanation is that contracts with good surface of cover specially in high peaks of demand in winter (which minimize exposure) have low or negative expected value and thus do not bring a priori any profit.

The demonstrator had many positive results. First it helped us understanding faster about the criteria correlation and their impact of the solutions produced in a visual manner. Second, it was a powerful communication tool between developers and users. It helped the users express and refine their requirements. Finally, it has been used to demonstrate Utility companies what could be done with the constraint technology with respect to "risk management decision support systems".

However, some additional requirements needed to be fulfilled before we reached a "good" model of the user's real problem. These requirements were:

- A revision of risk measurement. Demand exposure as defined did not represent a financial measure of risk since it was formulated as a quantitative measure of electricity cover.
- Dealing with uncertainty in the input data sets. Data were forecast data whose uncertainty needed to be considered and modeled.

5. THE STOCHASTIC NATURE OF THE PROBLEM

For specific "problem understanding" purposes, the first two prototypes considered a deterministic model of the energy trading problem, where we assumed that the input data were known with certainty. But the problem of "risk management in energy trading" is essentially a *stochastic* problem. The input data are forecast data whose reliability is uncertain. Making decisions under such conditions of uncertainty (data subject to changes) is core to this combinatorial problem.

Dealing with uncertainty is still a broad research topic in constraint programming and combinatorial optimization in general. There is no clear answer as to which model should be used to model the data and assess the risks in the decision making process. This section describes a risk model proposed by IC-Parc that attempts to measure and assess the risk in demand exposure, as well as a simulation model developed by Bouygues that evaluates the risk factor.

5.1 Risk model

The users say "*Risk is always related to the risk of not making profit which is affected by demand exposure and pool price exposure. The weight on demand and price risks varies depending on the time of the year and the time of the day. To a trader, the demand exposure is actually the exposure of availability. Indeed, the lack of electricity available at peak periods is the main cause for high price fluctuations. Thus, there is a high cost associated with buying electricity when there is no power set available at stand-by condition. Hence, the price will not be too high at off peak period even if demand is a lot higher than expected. This is because there is enough availability to supply the unexpected rise in demand*."

It is clear that assessing and measuring the risk is essential to this problem, as well as finding methods to control it. However, the risk factor is still ill-defined. At this stage of the project the deregulation of the electricity market has been postponed and the client has abandoned his search for a DSS due to company re-organization issues. IC-Parc and Bouygues were having no further advise from the client in deriving new models.

Based on the above trader's remark, IC-Parc defined a concept of risk exposure that expresses demand exposure in money terms. The idea consists of attaching a level of risk to each EFA of uncovered demand. The level is defined by introducing a cost factor based on a percentage of the forecast pool price. This allows profit and risk exposure to be measured in a common cost unit.

The input data model remains unchanged: electricity demand and pool price 12x12 matrices. In a pre-processing step, a risk factor is computed for each point of the 12x12 matrix by taking into account the nature of the peaks in demand, i.e. how high they are and when do they take place (which EFA period). So for each demand period of the 12x12 matrix is associated a function f_{perc} that returns a percentage of the pool forecast price. We considered three different percentages:

- 30 % (high risk) if exposure greater than 100 MW.
- 10 % (medium risk) if exposure greater than 50 MW (less than 100MW).
- 0% (low risk) if exposure nil.

The function f_{perc} becomes the indicator of risk adversity. When minimizing the risk one should aim at covering high peaks in demand in a better way than while minimizing demand exposure. This approach should ideally be an elegant compromise between the bin packing and the knapsack models highlighted previously. Testing the validity of the risk function together with the value of the percentages, was part of the collaboration between IC-Parc and Bouygues.

Thus, the new model is based on two criteria to optimize:

- **Minimize risk exposure.** Risk exposure is very close to the definition of demand exposure but contains the extra percentage factor. We have:

$$\forall m \in M, \forall i \in I,$$
$$f_{risk}(m,i) = \left| \texttt{demAvr}(m,i) - \sum_{c \in C | m \in duration(c)} \left(\texttt{shape}(c,i) \times \mathbf{cvol(c)}\right) \right| \times f_{perc}(m,i)$$

The risk exposure function to minimize is defined by $\sum_{m \in M} \sum_{i \in I} f_{risk}(m,i)$.

- **Maximize expected profit**. The expected profit of one contract is the same as in the previous model $\sum_{c \in C} f_{profit}(c)$ such that:

  $\forall c \in C,$

$$f_{profit}(c) = \mathbf{cvol(c)} \times \left(\frac{\sum_{m,i} \left(\texttt{pppAvr}(m,i) \times \texttt{shape}(c,i)\right)}{\sum_{m,i} \texttt{shape}(c,i)} - \texttt{str\_price}(c) \right)$$

5.1.1 *IC-Parc' demonstrator*

IC-Parc extended the demonstrator presented in section 4 with the risk exposure criterion. Optimal solutions to the MILP model were found in 1sec. on a Sparc Server 2170. The multi-criteria optimization was done by considering the risk exposure as the cost function, and constraining the expected profit to range within restricted bounds. We ran a set of experiments on minimizing the risk exposure. They showed us that the risk exposure function we have created did not influence the solution differently than the initial demand exposure function (without a risk factor). One reason could be that the contract profiles provided as input data, were too regular in shape to fit high peaks in demand whatever the risk exposure was. Another reason could be the choice of the Pool price percentages attached to the risk factor. If this percentage is too low for high periods of demand then it does not push any contract to fulfill the gaps. We have tuned this percentage several times and it did not yield tremendous variations on the solution.

The demonstrator pointed out that the deregulation in the energy trading market would definitely require energy traders to consider new types of profiles. Current profiles seem too restrictive. In fact, an ICL consultant in the field confirmed this point and expressed the idea of introducing options and other derivatives in the portfolio management system. However, at this stage of the problem understanding, this would have extended quite considerably the fuzziness of the model and problem formulation.

5.1.2 *Bouygues' solver*

The main objective of Bouygues, as far as the deterministic risk model was concerned, was to investigate how good constraint programming with local search techniques could perform against the MILP solver. The search algorithm is a Least Discrepancy Search (LDS) [8], i.e. a branch-and-bound algorithm with a control over the depth of the search tree to be explored, extended with the greedy heuristic of best expected profit. The branching is done on the selection of contracts. The heuristic determines the best contract to pick. Then we create a choice point with two branches: "to pick or not to pick the selected contract". Note that the selection is *dynamic* because the expected profit of a contract depends on the potential volume (a priori unknown) that might be selected. This algorithm is incomplete because we do not try all possible volumes (from a predetermined interval): we either pick the maximal admissible volume (from its domain) or we decide not to pick any (see Figure 5).

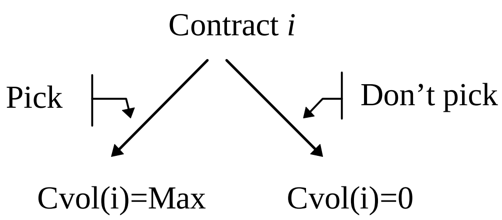

Figure 5: Branching scheme on the picking of contracts

There is no evaluation of a lower-bound function because the few simple estimates we tried did not yield interesting results (not enough pruning power). Propagation is very simple and limited to updating the maximal admissible volume for each contract (whether the contract is picked or not).

Because this algorithm uses no lower bounding, it is not adequate for solving a large problem. However, Bouygues noticed that the first branch at each node is almost always a good choice since it correspond to a " good " heuristic. When applying an LDS strategy we explore the sub-tree of the search space where right branches (contract refusals, on the right side of a graphical top-down representation) can be picked a limited number of times less than a value **K**. For **K** = 0, no contract is refused and the heuristic of best expected profit directs the selection. For small values of **K** we have a polynomial algorithm (the degree of which is **K**). In fact this approach amounts to fixing the minimum number of contracts to be picked. If **n** (15 in our case) is the initial number of possible contracts, LDS constrains us to have at least **n -K** contracts selected.

Bouygues' algorithm is approximate: each solution found at a node is feasible, but the optimal solution may be missed, either because the wrong contracts were selected (rare) or because the wrong volumes were assigned (frequent). We corrected this last problem using a local optimization method.

Local Optimization in *CLAIRE*[4]. The local optimization step is called at each node of the previous search tree. Its goal is to improve the cost function (expected profit minus expected risk) while keeping the previously selected values for each contract volume within the final bounds that were established during the search. Each contract volume ranges initially in an interval [**0 .. MaxVol**] which must contain the final value that will be eventually selected. A decision in the search tree has an impact on this interval in order to limit the domain to be explored in the local optimization. The goal is to try to increase the volume of contracts. This amounts either to increasing the current volume of a contract by an increment or to pushing a contract that was not yet selected (volume nil) in the portfolio. Namely, if we have assigned the volume of a contract **Cvol(i)** to **Vol** (in [**MinVol .. MaxVol**]) a local move from this node of the search tree corresponds to an alternative choice: a left choice moves the **MinVol** value (initially 0) to **Vol + increment,** and a right choice moves **MaxVol** to the current value of the volume **Cvol** (see Figure 6).

---

[4] See http://www.dmi.ens.fr/users/laburthe/claire.html

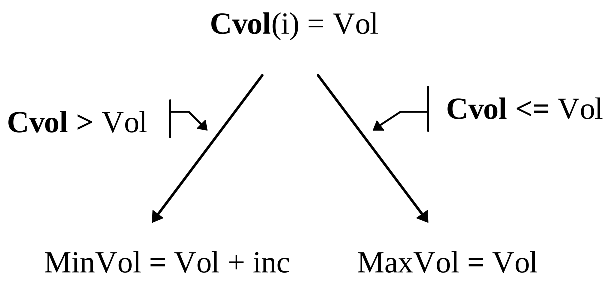

Figure 6: Local optimization step

Local moves are made using a simple push model. For any new contract we add a small quantity (using a fixed increment). This addition to the contract volume might create some over-cover for the global portfolio of contracts. We remove the excess using a greedy heuristic. The greedy heuristic is based on removing the most significant contributor to the over-cover, using at most the exact quantity that makes it the most significant contributor. A move is positive if the volume obtained after this substitution is better than the original. We search for all possible moves (for all contracts) and select the best one. We stop as soon as there is no strictly positive move, that is there is no improvement. To speed up the convergence process, we use a set of decreasing increments for pushing new contracts and we end the local optimization procedure only when no increment can produce an improvement.

As mentioned previously, the quantities that can be added or removed during a local move are bounded by the interval [**MinVol .. MaxVol**]. This ensures that the local search started at two different nodes will not produce the same solution.

5.1.3 Results

Bouygues obtained the optimal solution using simply the greedy algorithm and the local optimization step. This seems to be nothing more than a coincidence since we have observed a small but significant difference between LDS(3) and LDS(0) on randomly generated problems, where 3 and 0 represent different values of **K** (number of discrepancies allowed). Bouygues ran comparisons with IC-Parc's approach on 3 different problems only, but the estimate is that LDS search is definitely closer than 1% from optimal.

From a modeling accuracy point of view, the risk factor does not seem to influence the decision making process. A more insightful analysis of this point will result from the next model: the simulation model.

5.2 Simulation model

Bouygues designed a simulation model based on a set of simple probability distributions for the two sets of data parameters: the demand and the pool price. The (crude) model reflects our limited understanding of demand and pool price forecast. Although it reflects what was expressed by experts in the field, it should not be taken at face value but is rather an interesting indication of what can be accomplished through simulation. The job of collecting a realistic statistical model for demand and pool price forecast remains to be done.

Short-term demand forecasting is usually very precise. The availability of a wealth of data has produced precise models. Long-term demand forecasting, on the other hand, is adversely affected by two things. First, electricity demand is strongly correlated with climate variations, that are hard to predict six months in advance. Second, there are a few special days in the year (such as the accident of Lady Diana) that are hard to predict and that create peaks which cause major troubles.

5.2.1 Description

Bouygues' model for demand is to add two corrective terms to the forecast values:

- A *random uniform distribution* (within +/- 8%) that is drawn for each month which reflect the climate variation.
- A *squared random distribution* (within +/- 16%) that reflect special occasions that could not be forecasted. A different value is drawn for each day.

Similarly, the simulation model for the pool price is defined as the sum of 1) original forecast, 2) climate variation (same term as for the demand model), and 3) a random uniform distribution (+/- 10%) that reflects the " noise " around the pool price. The market games that define the pool price make it extremely volatile and there is no well-established model for forecasting pool price since it is linked to demand fluctuations. The variation bounds that we used to instantiate the model reflect our current understanding of the model. However, because this model is so crude and our understanding so incomplete, we actually tested four scenarios, that correspond to the following parameters:

- **Scenario 0**: climate variation +/- 10%, special days +/- 25%, pool price variation +/- 15%
- **Scenario 1:** climate variation +/- 8%, special days +/- 16%, pool price variation +/- 10%
- **Scenario 2:** climate variation +/- 6%, special days +/- 9%, pool price variation +/- 8%
- **Scenario 3:** climate variation +/- 4%, special days +/- 4%, pool price variation +/- 4%

Scenario 0 introduces more fluctuations, while scenario 3 is the closest to the forecast demand and pool price. The principle of the simulation is simple. Each run of the simulation draws a random distribution for demand and pool price. The associated profit is then computed according to the formula. Bouygues performed 100 separate runs and computed the mean value and the standard deviation of the profit.

5.2.2 Results and lessons learnt

Table 1 presents the results using the four scenarios in the simulation model. For each scenario we give the mean value of the profit **p** and the standard deviation in the form **p ± x**. The risk factor indicates how much we considered the risk function.

Risk factor	Scenario 0	Scenario 1	Scenario 2	Scenario3
0.0	1.955 ± 0.541	2.146 ± 0.397	2.257 ± 0.292	2.32 ± 0.186
0.03	1.538 ± 0.379	1.803 ± 0.261	1.934 ± 0.173	2.00 ± 0.102
0.1	0.903 ± 0.306	1.348 ± 0.209	1.613 ± 0.132	1.755 ± 0.072

Table 1: Different Scenarios for the Simulation Model

The algorithm used was the greedy algorithm. A reading from left to right shows that the undesirable side-effects (high deviation) tend to disappear if the forecast is closer to the reality (scenario 3), as we might expect. Still it is interesting to see that there is not a single instance where the risk-factor in the cost function makes sense. The reduction in the standard deviation (minimized risk) is always obtained with a bigger reduction in the mean value of the profit (loss of profit).

The lessons we have learnt from this set of experiments may be summarized as follows:

- The risk-factor is not meaningful to the problem. Although we obtained a significant reduction of the standard deviation, thus minimized the risk, this is at such an expense in terms of average profit that it does not make any business sense. We tried different values of the risk factor but there was no good value. The situation depends obviously on the price distribution of the set of contracts. Moreover, different data set from the ones provided by the users could have yield different results.
- The energy trading problem is a stochastic problem. Trying to solve it while ignoring the stochastic nature makes sense only if there is a didactic purpose in using a deterministic model.
- Simulation helps to validate a model and its economic function. It provided us with a lot of insights about the business implications of the model.

6. CONCLUSIONS

In this paper we have presented a case study in risk management dealing with "how to tackle ill-defined applications using constraint technology". We have shown the use of iterative prototyping combined with technical means to address incomplete components of a problem definition. The approach has highlighted aspects we believe are essential and which are:

Credo 1: *Iterative prototyping* is an adequate development process that allows an incremental refinement of the model with the users. However, its use sets requirements on having high level modeling languages like ECL^{i}PS^e and CLAIRE, to enhance and tune the model in an efficient and flexible way.

Credo 2: The first model should ideally be simple enough (no uncertainties) to check the adequacy and validity of the core problem formulation. If it is infeasible, goal programming is an efficient modeling technique to deal with over-constrained problems. Users are often reluctant to admit that their model has no feasible solution.

Credo 3: Multi-criteria optimization appears in most real world applications and is often cumbersome to model. When the correlation between the different criteria is not well established, a weighted objective function is unrealistic. We believe that optimizing on one criterion while constraining the others is a practical approach to search for potential correlation. Also, the use of an interactive visual interface on top of a DSS is an excellent means to draw efficient frontiers and allow the users to make informed judgments.

Credo 4: Building simulation models is a practical and cost effective means to get more insight on ill-defined problems when searching for the adequate problem definition (e.g. cost function). Simulation models can be easily designed using constraint technology.

This speculative work led us to refine the definition of the problem with the users. However, in terms of modeling uncertain data and assessing the risks, more work is required. In this final section we present directions for future research work in decision making under uncertainty [5].

6.1 Risk and decision making

Risk management deals with *uncertainty* in the occurrence of an unforeseen event (a change in demand or price data), and the *impact* on the decision making process if the unforeseen event occurs [9]. Since the 50s, risk management problems have been tackled by mathematical programming tools based on utility theory, decision theory, and game theory (e.g. see [11][1]) but to our knowledge constraint programming has not yet been used nor contributed to this application domain.

From a "decision theory" point of view, decision making under uncertainty commonly means that no probability distributions of the data are available or can be secured a priori [15]. Our risk management problem exemplifies this situation. Demand and price data may experience seasonal variations and trends. They are non-stationary (time dependent). The risk lies in having a portfolio of contracts that is exposed to any random fluctuation. Designing an adequate data model is a complex task. It is tackled in mathematical programming using approximation models (figure 7).

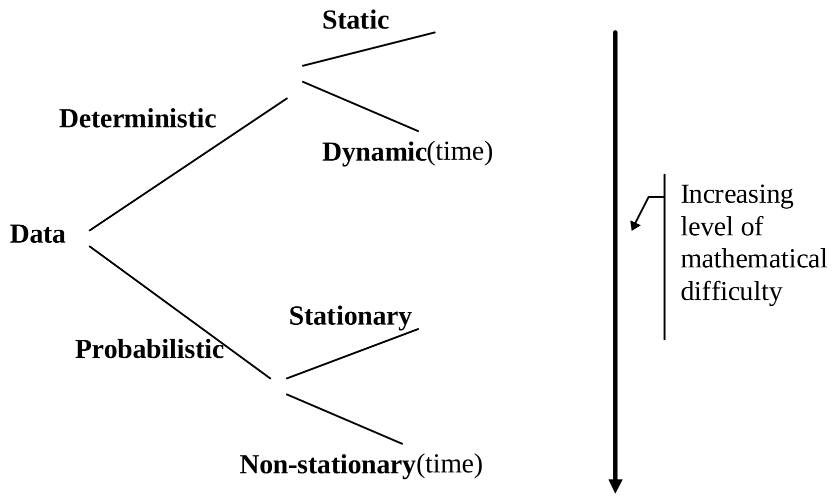

Figure 7: Approximation models for data representation

The data are either 1) reduced to a *certainty* case, and thus considered deterministic, 2) attached to randomly generated probability distributions, or 3) attached to well studied probability distribution generated by forecasting models. The choice is usually driven by the nature of the data[5]. Each type of approximation corresponds to a mathematical model for decision making under uncertainty: a simple deterministic model, random stochastic models, or stochastic models resulting from a forecasting model. Deterministic models average the data by their mean and reduce the problem to one under certainty. Stochastic models are based on generating probability distributions and studying the best and worst case scenarios. Forecasting models are designed to detect trends and seasonal variations in past data and are usually the basis for a future reasonable stochastic model.

The results we have presented in sections 3 and 4 correspond to the first approximation: decision making under certainty. Even though we introduced a decision criterion that measures risk as the demand exposure, our deterministic data representation did not reflect the uncertainty in trends. In section 5 we presented some experiments based on a simulation model. This requires further work among which an input from the client to adjust the model and introduce probability distributions. Concerning the use of forecasting models to identify trends and seasonal patterns in the data, our client explained that forecasting methods even though reliable for short term planning, are not reliable for long term planning (one year in our case). They are not worth considering for analyzing trends and fluctuation in electricity demand and price data.

6.2 Future work

One natural direction is to develop a better simulation model and to perform more complete experiments. We also need to find another cost function that represents a better model for risk.

[5] Several methods exist to characterize data; they are mainly based on the computation of their mean and variance (measure of the spread, degree of uncertainty)

This can be done in an empirical manner using a simulation tool or with a more analytical approach based on risk characterization, measurement and control [9]. Many different approaches can then be taken to build approximation models among which:

- Introducing stochastic models in constraint programming. We could generate probability distributions, and consider a well-defined cost function such as maximizing the difference between the mean and the standard deviation of the profit.
- Substituting the stochastic approach with a sampling strategy. We may suppose that we have a finite set of scenarios that are represented by pool price and demand forecast, and that our goal is also to minimize a combination of mean and worst case profit. Since it is clear that we have reached a good understanding of how to solve the basic problem, we are well equipped to attack a multi-sample version. Approaches based on case-based decision theory can then be considered (see [7]) and introduced in a constraint programming framework.

ACKNOWLEDGEMENTS

This work has been partially supported by the CEC in the Esprit project `Chic-2` 22165.

REFERENCES

[1] R.A. Breasley. *An introduction to Risk and Return from Common Stocks.* Blackwell, 1983.

[2] A. Charnes and W.W. Cooper. *Goal Programming and Multiple Objective Optimization.* European Journal of Operations Research, vol 1, 1977.

[3] `Chic-2` consortium: Bouygues, Euro-Decision, IC-Parc, ICL, NTUA, Renault, *`Chic-2` methodology user-guide*, technical report, Esprit project 22165, 1998.

[4] Using CPLEX Callable Library, version 3.0. CPLEX Optimization Inc., 1994.

[5] C. Gervet. Ongoing research in `Chic-2`: Decision making under uncertainty. In *CP'97 ERCIM/Compulog workshop on constraints*1997.

[6] C. Gervet, Y. Caseau, Denis Montaut et al. *Risk Management in Energy Trading,* `Chic-2` Final and public deliverable, 1998.

[7] I. Gilboa and D. Schmeidler, *Case-based Decision: An Extended Abstract,* (Invited talk) Proceedings of IJCAI'98.

[8] W. Harvey and M. Ginsberg. *Limited Discrepancy Search,* Proceedings of IJCAI'95.

[9] F.J.Heemstra and R.J. Kusters. Dealing with risk: a practical approach. In *Journal of Information Technology,* volume 11-4, pp 333-346. Chapman and Hall, 1996.

[10] IC-Parc. ECL^iPS^e version 4.0.7 user manual. Technical report, IC-Parc, 1998.

[11] F.H. Knight. *Risk, uncertainty and profit.* University of Chicago Press, 1921 (Edition 1971).

[12] A.K. Mackworth. *Consistency in Networks of Relations.* Artificial Intelligence 8, 1977.

[13] H.M. Markowitz. *Portfolio Selection, efficient Diversification of Investment.* Wiley, New York, 1959.

[14] J. K. Ousterhout, *TCL and the TK Toolkit,* Addison Wesley, 1994.

[15] Hamdy A. Taha. *Operations Research An introduction (fifth edition),* chapter 3-9, Prentice Hall International Editions, 1995.

Appendix A

The first screen presents a contract portfolio that maximizes the expected profit. The initial demand forecast is shown in the background as a chart model of the 12x12 matrix (12 EFA periods per month). The selected contracts are drawn on top of the demand forecast and listed in the small window (with their volume of unit cover). The contract names specify which time of the year the contract profiles correspond to (O for October, W for winter, S for summer, and A for annual), and which EFA period they cover.

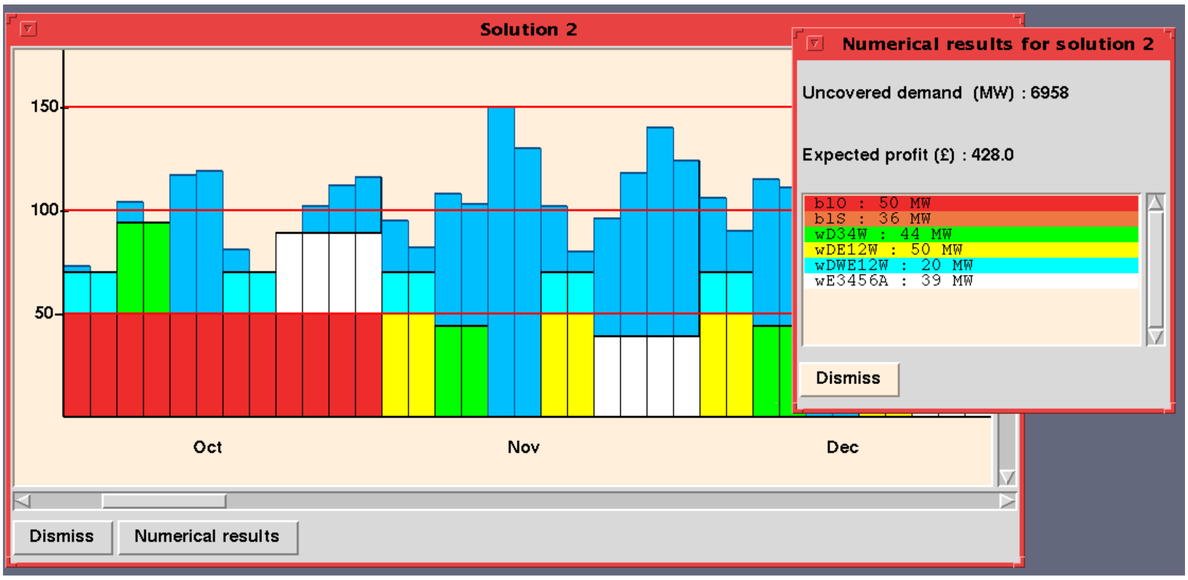

The second screen presents a portfolio that minimizes the demand exposure. A different set of contracts was produced with different associated volumes. One can see that it covers the demand forecast more closely (see EFA periods 5 and 6 in November). The amount of uncovered demand was reduced by more than half but the expected profit has dropped from 428£/MW to 251 £/MW.