Data Uncertainty in Constraint Programming: A Non-Probabilistic Approach

Neil Yorke-Smith* and **Carmen Gervet**
IC–Parc, Imperial College
London, SW7 2AZ, U.K.
{nys,cg6}@icparc.ic.ac.uk

Abstract

The constraint programming paradigm has proved to have the flexibility and efficiency necessary to treat well-defined large-scale optimisation (LSCO) problems. Many real world problems, however, are ill-defined, incomplete, or have uncertain data. Research on ill-defined LSCO problems has centred on modelling the uncertainties by approximating the state of the real world, with no guarantee as a result that the actual problem is being solved. We focus here on ill-defined data, motivated by problems from energy trading and computer network optimisation, where no probability distribution is known or can be usefully obtained. We suggest a non-probabilistic *certainty closure* approach to model the data uncertainty, discuss the formalism and semantics required to build a constraint solving system based on such a computation domain, and give examples in the case of linear systems.

Introduction

Uncertainty can arise from many sources. In the real world, we find dynamic environments, over-constrained problems, and partial or incomplete data. The difficulties of a large-scale combinatorial optimisation (LSCO) problem are correspondingly greater when the problem is ill-defined.

Four factors characterise an optimisation problem and uncertainty may feature in any of them: the input (or data), the constraints, the decision criteria, and consequently, the output (or solution). Some aspects have been examined in the literature:

- dynamic environments (new or changing constraints, anticipated changes): (Dechter & Dechter 1988; Wallace & Freuder 1999; Fowler & Brown 2000)
- over-constrained problems (hard and soft constraints, decision criteria): (Schiex, Fargier, & Verfaillie 1995; Dubois, Fargier, & Prade 1996)
- probabilistic models and data (incomplete or inconsistent information): (Taha 1997; Fargier & Lang 1993)

The constraint programming (CP) paradigm has been extended to address the first two of these areas, whereas the third — going back even to Dantzig (1955) — has been more traditionally the realm of operational research (OR). Interest in uncertainty within CP has been growing recently (Shazeer, Littman, & Keim 1999; Walsh 2000), but there has remained little work on data uncertainty.

Gervet *et. al.* (1999) describe a speculative constraint optimisation project in energy trading. Due to pending market deregulation and in conjunction with other ill-defined problem elements, there was no reasonable probability distribution that could be forecast for the data parameters. Instead, an iterative prototyping approach using simulation was employed.

In network management, by contrast, all the data may be available in theory but may be too costly and voluminous to collect in practice. Moreover, the data (such as traffic flow through a router) fluctuates constantly, and to sample at all points at one instant would be prohibitive (bear in mind clock drift between routers). Combined with the effects of dropped packets and other dynamic events, the data may be expected to be inconsistent by the time it is collected. It is of commercial interest to make decisions to optimise the network (Gervet & Rodošek 2000)[1], but we must do so based on incomplete and inconsistent information for which no stochastic description is apparent.

Ben-Haim (1999, Section 5) outlines the need for robust seismic reliability in the design of buildings in earthquake zones. Here, a stochastic definition of reliability, in terms of minimising probability of failure, is of little value, for the "information is much too scanty to verify a probabilistic model". Instead, convex modelling has been successfully applied to guarantee the integrity of the building.

These motivating problems from the real world exemplify that we face potentially the most difficult of the four types of data characterised in Figure 1: dynamic, non-deterministic data. Harder still, there is no stochastic characterisation available. It is clearly inadequate to address this situation in CP either by ignoring the uncertain data or by assuming the deterministic case.

The experience of traditional OR helps little, on the whole. The approach to uncertain data (Hoffman 2000; Liu 2000) has been largely to reduce to the deterministic

*Corresponding author

[1]This research has been applied to a commercial product, under development at IC–Parc and Parc Technologies: `www.parc-technologies.com`.

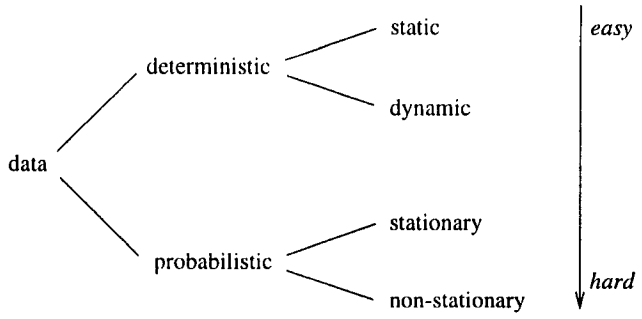

Figure 1: Classification of data

case and perform simulation, or to attach a probability distribution, perhaps derived by forecasting from past data, and consider risk and utility factors. In *robust optimisation*, for instance, variance measures or expected utilities are used; in *stochastic optimisation*, penalties based on expected feasibility of scenarios. In *fuzzy programming* (Sengupta, Pal, & Chakraborty 2001), constraints (and hence data) and goals are modelled as fuzzy sets, but now the set membership functions are assumed known.

All of these approaches rely on *a priori* knowledge or empirical estimation of some kind of probability distribution for the data, or of some other stochastic description of it. Further, rationality of the decision maker is assumed. Not in every situation, as we have seen, are these requirements meaningful or even feasible.

The need for a robust, non-probabilistic alternative has been recognised in OR (Ben-Tal & Nemirovski 1999; Ben-Haim 1999). As Hoffman (2000) writes, "We hope that future research will also address [in addition to incorporating risk] the issue of how to incorporate ... data for which even the mean value is not known and for which one only has range estimates of its value."

In the following sections, we shall speak of uncertainty to mean data uncertainty. We introduce the concept of the certainty closure and apply the methodology to a case study arising from our motivational problems. We then place the work in a wider context and anticipate future steps.

The Certainty Closure

We give a definition of the certainty closure and introduce the semantic and algorithmic concerns that follow for data uncertainty in constraint programming. The idea is to enclose the uncertainty and so be able to solve an ill-defined problem without approximation via some mathematical model. Thus we guarantee that the actual problem is being solved and that the solutions found are robust.

Data uncertainty is reflected as uncertainty in the coefficients of the constraints. We suppose here that all the constraints are hard and static and that the problem is not over-constrained; such issues have been considered elsewhere (for example, (Schiex, Fargier, & Verfaillie 1995)). An introduction to CP may be found in (Marriott & Stuckey 1998).

Definition 1 *An* uncertain constraint satisfaction problem $\langle \mathcal{V}, \mathcal{D}, \mathcal{C} \rangle$ *is a standard CSP in which the coefficients in the constraints* $\mathcal{C}$ *may range, not necessarily independently, over an uncertainty set* $\mathcal{U}$.

It is unknown which values in the uncertainty set the data might take, but it is necessary in every case that no solution to the resulting hard constraints be excluded. This motivates the definition of the certainty closure:

Definition 2 *The* certainty closure *of an uncertain constraint satisfaction problem* $\langle \mathcal{V}, \mathcal{D}, \mathcal{C} \rangle$ *is the CSP* $\langle \mathcal{V}, \mathcal{D}, \mathcal{C}' \rangle$, *where* $\mathcal{C}'$ *is derived from* $\mathcal{C}$ *as follows. For each constraint* $C \in \mathcal{C}$, *let* $C^{(i)}$ *be the* $0 < k < \infty$ *possible constraints resulting from* C *as the coefficients in* C *vary over the uncertainty set of their domains. Let* $C' = \bigvee_{i=1}^{k} C^{(i)}$ *be the disjunction of the* $C^{(i)}$. *Then let* $\mathcal{C}'$ *be the conjunction of the constraints* C', *without loss of generality with repeated and trivially redundant constraints removed.*

That is, the constraints in $\mathcal{C}'$ are the generalisation of those in $\mathcal{C}$ to hold under all possible realisations of the data. We write ρ to denote the certainty closure mapping:

$$\langle \mathcal{V}, \mathcal{D}, \mathcal{C} \rangle \xrightarrow{\rho} \langle \mathcal{V}, \mathcal{D}, \mathcal{C}' \rangle$$

We give below an extended example in the form of a case study. The idea can be seen in the transformation of the simple uncertain constraint $X \geq [20, 30]$ under ρ to $X \geq 20$ (where $[\underline{X}, \overline{X}]$ denotes a real interval). Satisfaction of the latter constraint guarantees satisfaction of the former, no matter what the true value of the data.

It might be objected that this is the worst case and over-emphasises pessimistic combinations of data values. In reply, we assume (motivated as discussed) only that all elements of $\mathcal{U}$ are possible; we have no information as to likelihood. Further, random values within $\mathcal{U}$ can be almost as bad as the most difficult values (Ben-Tal & Nemirovski 2000).

The certainty closure allows robust inference on the original problem in the sense that the domains resulting for the variables are *sure*: valid whatever the realisation of the data within the uncertainty set. The new CSP will be satisfiable provided the original is satisfiable for some element of $\mathcal{U}$: the system $X \geq [20, 30] \wedge X \leq [20, 25]$ need not be satisfiable, for example, whereas its certainty closure will be.

Note we do not assume that $\mathcal{U}$ is the direct product of the domains of the coefficients of $\mathcal{C}$: the values for the data may be dependent. This is in contrast to probabilistic models, where independence of data parameters is nearly always required.

The distinction must be made between the *natural* domains of the variables $\mathcal{V}$ and the constraint or *calculus* domain which is reached by applying the certainty closure. The former is the domain with which the user interacts and over which constraints are applied; the latter is how we represent the inferences and calculate with surety, and is hidden from the user. The different levels are portrayed in Figure 2.

The solution of an uncertain CSP is derived from its certainty closure such that all possible solutions to the original CSP are contained within. This means, in general, that the

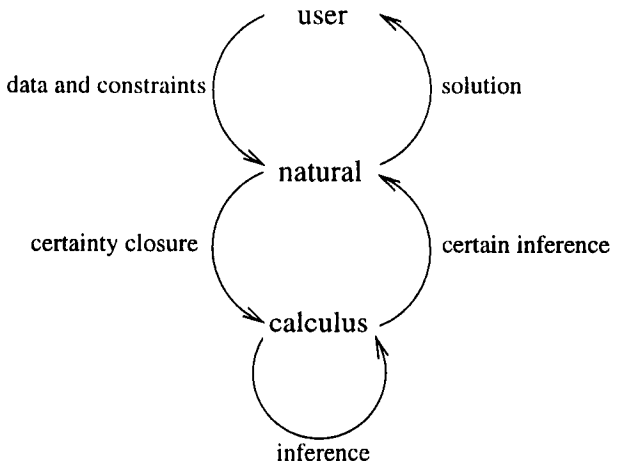

Figure 2: Natural and calculus domains

decision variables will take a set of possible values rather than be instantiated to one value — but the domains that result are sure and correct.

In enhancing the CP paradigm to cope with uncertainty, we do not at the calculus level assume the rationality of the decision maker. Rather than proposing a single instantiation, precluding all others, we return all possible solutions. This contrasts with the fundamental assumption of rationality often made in probabilistic approaches.

The type of constraints found in the energy trading and network optimisation problems of the Introduction are principally linear. Motivated by this, consider the case of data uncertainty in a linear problem, formalised as an interval linear system (ILS), as defined below. We follow the notation of Neumaier (1990), except to use bold font to denote an interval quantity.

Definition 3 *Let $\mathcal{V}$ be a set of n variables over $\mathbb{R}$, and $\mathcal{C}$ be a set of m linear constraints (equalities or inequalities) for $\mathcal{V}$. Writing each constraint in normal form, let $A \in \mathbb{R}^{m \times n}$ be the matrix of left-hand sides, and $b \in \mathbb{R}^m$ be the vector of right-hand sides. Let R be the list of m relations (or constraint symbols), one for each constraint; $R_i \in \{<, \leq, =, \geq, >\}$, $\forall i = 1, \ldots, m$.*

Then an interval linear system *induced by $\mathcal{C}$ on $\mathcal{V}$ is a tuple $\langle \mathbf{A}, R, \mathbf{b} \rangle$, where $\mathbf{A} \in \mathbb{IR}^{m \times n}$ is an interval matrix $[\underline{A}, \overline{A}]$ with $A \in \mathbf{A}$, and $\mathbf{b} \in \mathbb{IR}^m$ is an interval vector $[\underline{b}, \overline{b}]$ with $b \in \mathbf{b}$.*

For example, the ILS given by

$$\mathbf{A} = \begin{pmatrix} [-2,2] & [1,2] \\ [-2,-1] & -1 \\ 6 & [\frac{5}{2},3] \end{pmatrix} \text{ and } \mathbf{b} = \begin{pmatrix} [3,4] \\ [-5,5] \\ [4,15] \end{pmatrix}$$

and $R = (\leq, =, =)^\top$ has the *solution set* $\Sigma(\mathbf{A}, \mathbf{b})$, shown in Figure 3, which is given by the rays $2X + Y = 3$, $6X + 3Y = 15$, $2X - Y = -5$ and $12X + 3Y = 8$, and the points:

$$(X, Y) \in \Big\{ (0,4), (\tfrac{1}{3}, \tfrac{14}{3}), (\tfrac{5}{3}, \tfrac{10}{3}), (\tfrac{5}{2}, 0), (\tfrac{23}{6}, -\tfrac{38}{3}), (\tfrac{2}{3}, 0) \Big\}.$$

Hence the solution set is non-convex and the *interval hull* $\square\Sigma(\mathbf{A}, \mathbf{b})$, the smallest hyperbox enclosing the solution set, is unbounded.

The above formulation is very general. Following (Ben-Tal & Nemirovski 1999), we do not impose what form the uncertainty set might take. Two common choices, from the field of robust computation, are *intervals* and *ellipsoids*.

An interval (strictly, a closed non-empty bounded interval) given by lower and upper bounds is the simplest description of an uncertain value, and the properties and uses of interval computation are well-known (Alefeld & Herzberger 1983). Ellipsoids, while more complicated than intervals, arise naturally in problems in engineering. An ellipsoid can capture dependency information between variables and normally-distributed measurement error, and ellipsoidal computation is a little easier than interval computation due to the compact parametric representation of an ellipsoid (Kreinovich *et al.* 1998).

A natural extension to an interval is a *multi-interval*, a finite union of intervals. This allows us to specify not only the bounds on $\mathbf{X}$, but also which values in $[\underline{X}, \overline{X}]$ are impossible. Unfortunately, because convexity is lost, multi-intervals make computation too difficult (Kreinovich *et al.* 1998, Chapter 24).

We will make use of intervals for three reasons. First, they describe well the uncertainty in the motivational LSCO problems, which have uncoupled data known only to be between lower and upper bounds. Second, libraries for interval computation are available in CLP systems. Third, intervals give ease of intuition.

In the next section, intervals are used to give a demonstration of certainty closure applied to a tractable subclass of uncertain linear systems.

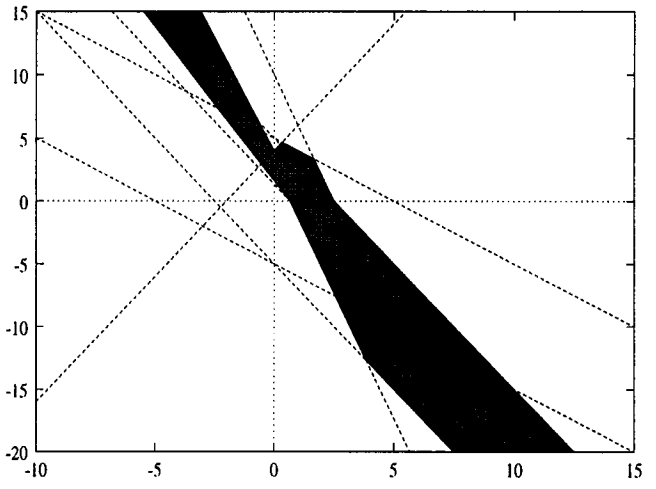

Figure 3: Solution set

Case Study: Positive Orthant Interval Linear Systems

Linear constraints, and systems which can be readily linearised, are sufficient to model many (although far from all) important real-world problems. A case in point is that of network optimisation. The difficulty here is not in the constraints, which largely form a linear flow model, but in the data, which is incomplete. The uncertain model variables are non-negative reals which may be assumed independent, thus motivating the sequel.

Definition 4 *A* positive orthant interval linear system *is an interval linear system in which the natural domain of each variable, dom$(V), V \in \mathcal{V}$, is non-negative. Thus, the solution set lies within the positive orthant of* $\mathbb{R}^n$.

An example is the system shown earlier in Figure 3, if we impose $X, Y \geq 0$. Such a system $\langle \mathbf{A}, R, \mathbf{b} \rangle$ is tractable, unlike the general case, because the solution set in the positive orthant, denoted $\Sigma(\mathbf{A}, \mathbf{b})_+ = \Sigma(\mathbf{A}, \mathbf{b}) \cap \mathbb{R}^n_+$, is convex.[2]

However, to avoid an exponential growth in the number of faces of the closure to a single constraint, we impose a restriction (★) on equality constraints:

$$\forall i = 1, \ldots, m, \quad \text{if } R_i = (=) \text{ and } 0 \in \mathbf{A}_{i,n}, \text{then}$$
$$\neg\left(\hat{\mathbf{A}}_i > 0 \wedge \mathbf{b}_i > 0\right) \wedge \neg\left(\hat{\mathbf{A}}_i < 0 \wedge \mathbf{b}_i < 0\right)$$

where $\hat{\mathbf{A}}_i$ denotes the vector $(\mathbf{A}_{i,1}, \ldots, \mathbf{A}_{i,n-1})$ of length $n-1$. Note that the intervals throughout and the solution set itself may be unbounded.

The condition (★) is sufficient to admit that a pair of hyperplanes bound the space described in the positive orthant by an uncertain equality constraint. While not restrictive, the condition may be weaker than it appears, for in certain circumstances the matrix $\mathbf{A}$ can be written in a suitable form, by pre-conditioning or transformation. The unsymmetric nature of (★) is explained by its geometric interpretation, namely we have chosen one axis normal e_n as 'up', and our concepts of 'less than' and 'greater than' will be with respect to this direction.

The intuition behind the method is that each inequality gives rise to one halfspace (a line, in 2D), and the points which correspond to feasible solutions of all the constraints are those which lie within the intersection of these halfspaces. This intersection is guaranteed (in the positive orthant) to be convex, and (★) permits equalities to be rewritten as a pair of inequalities.

Our motivational problems have no uncertainty in the objective function, and to simplify the presentation we postpone a discussion of optimisation under uncertainty. However, observe that for an ILS with linear objective $\max \sum_i \mathbf{c}_i X_i$ (without loss of generality), uncertainty in the objective is easily removed by adding the additional constraint $\sum_i \mathbf{c}_i X_i \geq Z$ for an auxiliary variable Z and optimising $\max Z$. This reduces the problem to the main case.

[2]That is, the solution set is the intersection of halfspaces specified by the constraints, intersected with the positive orthant; the set need not otherwise lie entirely within the positive orthant and thus need not be convex (contrast (Aberth 1997)).

We will transform the system $\langle \mathbf{A}, R, \mathbf{b} \rangle$ to its certainty closure, the system $\langle \mathbf{A}', R', \mathbf{b}' \rangle$. Figure 4 gives the algorithm in outline; the details are discussed below.

INPUT:	ILS $\langle \mathbf{A}, R, \mathbf{b} \rangle$ as part of a CSP $\langle \mathcal{V}, \mathcal{D}, \mathcal{C} \rangle$
OUTPUT:	Bounds for the variables in $\mathcal{V}$
METHOD:	1. Check for suitable form
	2. Rewrite equalities
	3. Transform to certainty closure
	4. Find convex hull
	5. Derive interval hull

Figure 4: Overview of algorithm

Steps 1 and 2: The required form is to have linear constraints with interval uncertainty, and to obey (★). Each equality constraint can be replaced by a pair of inequalities, taking care to choose the relational operators correctly depending on the gradient with respect to e_n.

Step 3: The transformation ρ is given by replacing the matrix $\mathbf{A}$ and vector $\mathbf{b}$ by non-interval versions as follows.

Proposition 5 *The certainty closure of a positive orthant interval linear system* $L = \langle \mathbf{A}, R, \mathbf{b} \rangle$ *satisfying* (★) *is the numeric linear inequality system* $A'x \leq b'$, *where* $A' \in \mathbb{R}^{m \times n}, x \in \mathbb{R}^n$, *and* $b' \in \mathbb{R}^m$ *are given by the following.*

If $0 \in \mathbf{A}_{i,n}$, *then* $((A')_i, b'_i)$ *is*

$$\begin{cases} \left(\underline{\hat{\mathbf{A}}_i}, \overline{\mathbf{A}_{i,n}}, \underline{\mathbf{b}_i}\right) & \text{if } \{<, \leq\} \in R_i, \\ -\left(\overline{\hat{\mathbf{A}}_i}, \underline{\mathbf{A}_{i,n}}, \underline{\mathbf{b}_i}\right) & \text{if } \{>, \geq\} \in R_i, \end{cases}$$

while if $0 \notin \mathbf{A}_{i,n}$, *then* $((A')_i, b'_i)$ *is*

$$\begin{cases} \left(\underline{(\mathbf{A})_i}, \overline{\mathbf{b}_i}\right) & \text{if } \{<, \leq\} \in R_i \text{ and } \underline{\mathbf{A}_{i,n}} \geq 0, \\ \left(\overline{(\mathbf{A})_i}, \underline{\mathbf{b}_i}\right) & \text{if } \{<, \leq\} \in R_i \text{ and } \overline{\mathbf{A}_{i,n}} < 0, \\ -\left(\overline{(\mathbf{A})_i}, \underline{\mathbf{b}_i}\right) & \text{if } \{>, \geq\} \in R_i \text{ and } \underline{\mathbf{A}_{i,n}} \geq 0, \\ -\left(\underline{(\mathbf{A})_i}, \overline{\mathbf{b}_i}\right) & \text{if } \{>, \geq\} \in R_i \text{ and } \overline{\mathbf{A}_{i,n}} < 0. \end{cases}$$

Proof Omitted.

Consider again the example of the previous section. From the initial system

$$\mathbf{A} = \begin{pmatrix} [-2,2] & [1,2] \\ [-2,-1] & -1 \\ 6 & [\frac{5}{2},3] \end{pmatrix} \text{ and } \mathbf{b} = \begin{pmatrix} [3,4] \\ [-5,5] \\ [4,15] \end{pmatrix}$$

and $R = (\leq, =, =)^T$, the transformation to the certainty closure yields

$$A' = \begin{pmatrix} -2 & 1 \\ -2 & -1 \\ 1 & 1 \\ -6 & -3 \\ 6 & \frac{3}{2} \end{pmatrix} \text{ and } b' = \begin{pmatrix} 4 \\ 5 \\ 5 \\ -4 \\ 15 \end{pmatrix}.$$

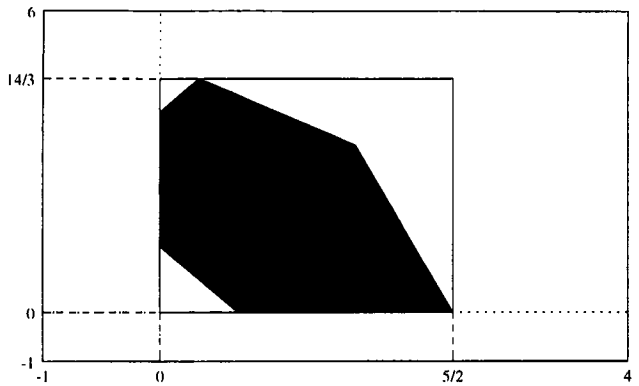

Figure 5: Tight bounds

Hence the positive orthant interval hull is

$$\square\Sigma\,(\mathbf{A},\mathbf{b})_+ = \left([0,\tfrac{5}{2}],[0,\tfrac{14}{3}]\right),$$

and the bounds derived can been seen in Figure 5.

By this means, from a constraint problem with interval uncertain data, we obtain a certain constraint problem. By inference on the transformed problem (that is, the certainty closure), sure domains can be found for the variables — domains which include the values that arise under every data realisation — and for an ILS, we give the inference explicitly.

Steps 4 and 5: The certainty closure of an ILS is a halfspace description of a polytope. The action of the constraints on the variable domains is equivalent to the projection of the convex hull of the polytope onto each axis (Oettli 1965). Geometrically, we seek bounds on the dual representation of the polytope, that is, the bounds of its vertex form. In practice, we can use linear programming to find the interval hull directly.

For each variable X_j, $j = 1,\ldots,n$, solve two linear programs which differ only in their objectives, $\min X_j$ and $\max X_j$, subject to the constraints given by the certainty closure. With $2n$ applications of simplex we have tight, sure bounds:

Proposition 6 *Let $\rho(L)$ be the certainty closure of L defined above. The bounds obtained on the domains $\mathcal{D}$ by the method given are the tightest certain bounds possible, and can be computed in expected time $\mathcal{O}(mn^3)$.*

Proof The transformation from L to $\rho(L)$ can be done in $\mathcal{O}(2mn)$ operations, and its correctness follows from Proposition 5. The constraints in $\rho(L)$ are linear and the solution space of their conjunction intersected with the positive orthant is a convex polytope, possibly unbounded (Aberth 1997). This polytope is the convex hull $\Sigma\,(\mathbf{A},\mathbf{b})_+$ by construction, and projecting onto each normal e_i immediately gives the interval hull. The bounds so obtained are tight to the solution set since the maximum and minimum possible value of each variable are found by the simplex iterations. The projection can be done with $2n$ iterations, at expected cost $\mathcal{O}(3mn^2)$ each. ■

This provides bounds for the natural domains of the variables. It may well be that not every value in the product space of the variable domains is feasible; other constraint inference techniques would be needed to determine this, since no further tightening will be obtained by reasoning solely on their bounds.

The certainty closure for interval linear systems, and the hull inference described above, have been implemented on the ECL^{i}PS^e CLP system (IC- 2001) using the interval computation library ic and the commercial LP solver XPRESS-MP. For problems up to several hundred variables and constraints, even a straightforward implementation gives the bounds in a few seconds.

We are investigating the application of the certainty closure to the network optimisation problem introduced earlier. Initial findings show promise. The current approach of data correction to the uncertain LSCO has been complemented in three areas:

- proving the correctness or otherwise of current results
- identifying deep inconsistencies, beyond the point of data correction, in the raw data
- understanding the source of bottlenecks in the network

Thus insight is gained of the relationship between the network topology and traffic flow, leading to robust quantitative results and improved global understanding.

Related Work

A problem with uncertain data is, in general, much harder than the same problem with certain data, as we might expect. With a linear system, for instance, there exists a polynomial time algorithm for certain data, but even checking a possible enclosure for interval uncertain data is NP-hard (Kreinovich *et al.* 1998, Chapters 11–12), and all methods for the general problem are exponential in the worst case.

We have considered interval linear systems from a CP viewpoint, as a case study of a non-probabilistic methodology for uncertainty. Problems similar to ILS have been examined from the viewpoints of mathematical programming and linear algebra.

Data uncertainty in linear programming has been addressed by Ben-Tal and Nemirovski (1999). They define a 'robust counterpart' (RC) to an LP problem, and seek the solution with the best objective that satisfies all realisations of the constraints. There are close parallels between the ideas of the robust counterpart and of the certainty closure of a linear system. For ellipsoidal uncertainty with boundedness assumptions, they use an interior point method to solve the RC problem in polynomial time. The uncertainty in the constraints must be independent.

The most general method for 'interval linear programming' is that of Chinneck and Ramadan (2000). Using first a transformation, they find the best and worst optimum for a linear system in which intervals may appear as coefficients in the constraints or the objective. Where the variables are unrestricted in sign or where equality constraints occur, they must resort to an exponential enumeration of cases.

Various authors have proposed analytic solutions to restricted cases of the ILS problem (see (Neumaier 1990; Ning & Kearfott 1997) among others). The iterative methods that usually result require the matrix $\mathbf{A}$ to be square and have certain properties, and the solution set to be bounded.

A different approach, first suggested by Oettli (1965) in an early paper, is to find some characterisation of the solution set as a polytope, then apply the simplex algorithm $2n$ times to find the extreme points with respect to each axis normal. This is the method we use above to perform inference. Much later, Aberth (1997) extended this technique for solutions spanning multiple orthants, although his method therefore can be exponential in n.

Beaumont (1998) has suggested an adapted simplex method to reduce the complexity. Given an initial enclosure of the solution set, his method produces a refined enclosure, but the tightest bounds are not necessarily reached. It would be worthwhile to integrate some of Beaumont's technique, such as the improved starting point for the simplex invocations, into our inference solver for linear constraints.

Chiu and Lee (1995) were concerned to improve the handling of linear equalities with interval variables (interval data they did not consider). By making incremental preconditioned interval Gauss–Seidel iteration (Neumaier 1990), and integrating it with a solver for non-linear CLP($\mathcal{R}$) constraints, they performed the domain reasoning. Their method cannot handle inequalities.

The inference scheme given contrasted with methods developed for interval constraint logic programming (see, for example, (Benhamou 1995)). In the latter, the domains of the variables are treated as (usually real) intervals, and interval-based local consistency operators applied. Since the certainty closure yields a certain constraint problem, if intervals should be used as the uncertainty set then it is possible to exploit interval CLP methods, such as bounds propagation (Marriott & Stuckey 1998, Chapter 3), for the inference algorithm. However, the tight solution set is not guaranteed and, further, interval narrowing methods perform best in answering global questions arising from non-linear constraint systems (Chiu & Lee 1994; Benhamou 2001).

Perspectives and Future Directions

Work on data uncertainty in constraint programming is fairly new, despite the reality of uncertainty in many real world situations. Existing CP models do not address this aspect of the problem when a probabilistic framework is not applicable. The aim of the certainty closure approach is to provide robust solutions to uncertain data in ill-defined LSCO problems, and our first results are promising in terms of robustness and efficiency.

The case study points us to three issues in integration with the paradigm. First, the method introduced must be made incremental. Second, it must be interfaced with solvers for other constraints. Third, there is the question of optimisation: how does uncertainty affect a *constraint optimisation problem* (COP), and how can uncertainty in an objective function be handled?

Incrementality of execution is important for solvers in CP (Marriott & Stuckey 1998, Chapter 10). Subsequent to the initial system, it is common to receive updated or new information regarding the uncertain data, but it is inefficient to recompute the entire inference on each occasion. For the case of interval linear systems, we can think in terms of the addition of a new halfspace equation, or correspondingly a potential 'slicing-off' of some part of the convex hull.

One of the strengths of CP is the ability to handle many types of constraints (non-linear, scheduling, disjunctive, and so on), and problems involving heterogeneous constraints. The second area to explore in future work is to consider a wider class of constraints in uncertain LSCO problems, and consider how the certainty closure should be formulated.

The third area is optimisation. If the decision criterion is maximisation of a function $f : \mathcal{V} \to \mathbb{R}$, say, then solving the certainty closure as a COP gives a hard upper bound. The problem is more difficult if f contains uncertainty; here the definition of ρ will need to be extended.

There has been little previous work on uncertain COPs. Sengupta, Pal and Chakraborty (2001) model vagueness and uncertainty in linear inequalities by using intervals, and present an interpretation of the constraints based on the preference of the decision maker. Using this total interval ordering, they find one problem instance which provides a satisfactory exact problem and optimise it by means of classical LP techniques. They have no guarantee of the efficacy of the solution, nor of the best or worst optimum.

Chinneck and Ramadan (2000), discussed earlier, find the best and worst optimum and the data values which lead to these. The best optimum supposes the most favourable values of the data; the worst, like the certainty closure, the least favourable. Once again, rationality of the decision maker is assumed.

Of these three issues, it appears that research may well be best focussed on applying the certainty closure approach to wider classes of constraints, so that the successes of CP in solving heterogeneous real world problems may be extended.

Acknowledgement. The authors thank Warwick Harvey for constructive suggestions and much helpful discussion. This work was partially supported by the EPSRC under grant GR/N64373/01.

References

Aberth, O. 1997. The solution of linear interval equations by a linear programming method. *Linear Algebra and its Applications* 259:271–279.

Alefeld, G., and Herzberger, J. 1983. *Introduction to Interval Computations*. New York: Academic Press.

Beaumont, O. 1998. Solving interval linear systems with linear programming techniques. *Linear Algebra and its Applications* 281:293–309.

Ben-Haim, Y. 1999. Set-models of information-gap uncertainty: Axioms and an inference scheme. *Journal of the Franklin Institute* 336:1093–1117.

Ben-Tal, A., and Nemirovski, A. 1999. Robust solutions of uncertain linear programs. *Operations Research Letters* 25:1–13.

Ben-Tal, A., and Nemirovski, A. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathmatical Programming Series A*. Published online, DOI 10.1007/s101070000163.

Benhamou, F. 1995. Interval constraint logic programming. In Podelski, A., ed., *Constraint Programming: Basics and Trends*, volume 910 of *LNCS*. Springer–Verlag. 1–21.

Benhamou, F. 2001. Interval constraints. In Floudas, C. A., and Pardalos, P. M., eds., *Encyclopedia of Optimization*. New York: Kluwer Academic Publishers.

Chinneck, J. W., and Ramadan, K. 2000. Linear programming with interval coefficients. *Journal of the Operational Research Society* 51(2):209–220.

Chiu, C. K., and Lee, J. H. M. 1994. Towards practical interval constraint solving in logic programming. In *Proc. of ILPS-94*, 109–123.

Chiu, C. K., and Lee, J. H. M. 1995. Interval linear constraint solving using the pre-conditioned interval Gauss–Seidel method. In *Proc. of ICLP'95*, 17–32. Kanagawa, Japan: MIT Press.

Dantzig, G. 1955. Linear programming under uncertainty. *Management Science* 1(3–4).

Dechter, R., and Dechter, A. 1988. Belief maintenance in dynamic constraint networks. In *Proc. of AAAI-88*, 37–42.

Dubois, D.; Fargier, H.; and Prade, H. 1996. Possibility theory in constraint satisfaction problems: Handling priority, preference and uncertainty. *Applied Intelligence* 6:287–309.

Fargier, H., and Lang, J. 1993. Uncertainty in constraint satisfaction problems: A probabilistic approach. In Clarke, M.; Kruse, R.; and Moral, S., eds., *Symbolic and Qualitative Approaches to Reasoning and Uncertainty (ECSQARU'93)*, volume 747 of *LNCS*, 97–104. Springer–Verlag.

Fowler, D. W., and Brown, K. N. 2000. Branching constraint satisfaction problems for solutions robust under likely changes. In *Proc. of CP2000*, 500–504.

Gervet, C., and Rodošek, R. 2000. RiskWise-2 problem definition. IC–Parc Internal Report.

Gervet, C.; Caseau, Y.; and Montaut, D. 1999. On refining ill-defined constraint problems: A case study in iterative prototyping. In *Proc. of PACLP'99*, 255–275.

Hoffman, K. L. 2000. Combinatorial optimization: Current successes and directions for the future. *Journal of Computational and Applied Mathematics* 124:341–360.

IC–Parc. 2001. *ECL^{i}PS^e User Manual Version 5.2*. `www.icparc.ic.ac.uk/eclipse`.

Kreinovich, V.; Lakeyev, A.; Rohn, J.; and Kahl, P. 1998. *Computational Complexity and Feasibility of Data Processing and Interval Computations*. Number 10 in Applied Optimization. Dordrecht: Kluwer Academic Publishers.

Liu, B. 2000. Uncertain programming: Optimization theory in uncertain environments. In *Proc. of the 9th IEEE Intl. Conf. on Fuzzy Systems*, volume 2, 941–944.

Marriott, K., and Stuckey, P. J. 1998. *Programming with Constraints: An Introduction*. Cambridge, MA: MIT Press.

Neumaier, A. 1990. *Interval Methods for Systems of Equations*. Number 37 in Encyclopedia of Mathematcs and Its Applications. Cambridge: Cambridge University Press.

Ning, S., and Kearfott, R. B. 1997. A comparison of some methods for solving linear interval equations. *SIAM Journal of Numerical Analysis* 34:1289–1305.

Oettli, W. 1965. On the solution set of a linear system with inaccurate coefficients. *Journal of SIAM: Series B, Numerical Analysis* 2(1):115–118.

Schiex, T.; Fargier, H.; and Verfaillie, G. 1995. Valued constraint satisfaction problems: Hard and easy problems. In *Proc. of IJCAI'95*, 631–637.

Sengupta, A.; Pal, T. K.; and Chakraborty, D. 2001. Interpretation of inequality constraints involving interval coefficients and a solution to interval linear programming. *Fuzzy Sets and Systems* 119(1):129–138.

Shazeer, N. M.; Littman, M. L.; and Keim, G. A. 1999. Constraint satisfaction with probabilistic preferences on variable values. In *Proc. of the Sixteenth National Conf. on Artificial Intelligence*.

Taha, H. 1997. *Operations Research: An Introduction*. Prentice Hall, 5th edition.

Wallace, R. J., and Freuder, E. C. 1999. Representing and coping with recurrent change in dynamic constraint satisfaction problems. In *CP'99 Workshop on Modelling and Solving Soft Constraints*.

Walsh, T. 2000. Stochastic constraint programming. Technical Report APES-23-2000, APES Research Group.