

A New Basis for Spreadsheet Computing

INTERVAL SOLVER for Microsoft EXCEL®

Eero Hyvönen and Stefano De Pascale

■ There is a fundamental mismatch between the computational basis of spreadsheets and our knowledge of the real world. In spreadsheets, numeric data are represented as exact numbers and their mutual relations as functions, whose values (output) are computed from given argument values (input). However, in the real world, data are often inexact and uncertain in many ways, and the relationships, that is, constraints, between input and output are far more complicated. This article shows that interval constraint solving, an emerging AI-based technology, provides a more versatile and useful foundation for spreadsheets. The new computational basis is 100-percent downward compatible with the traditional spreadsheet paradigm. The idea has been successfully integrated with Microsoft EXCEL as the add-in INTERVAL SOLVER that seamlessly upgrades the arithmetic core of EXCEL into interval constraint solving. The product has been downloaded by thousands of end users all over the world and has been used in various applications in business computing, engineering, education, and science. There is an intriguing chance for a major breakthrough of the AI technology on the spreadsheet platform: Tens of millions of EXCEL users are making important decisions based on spreadsheet calculations.

The world is full of uncertainty and complexity. Everyday we are faced with questions such as, How can I live within the given budget? Is this technical design possible, given the inaccurate component data? Uncertain data and constraints are extensively used in decision making, but spreadsheets, one of the most commonly used decision-making aids of today, force us to use exact numbers for representing inexact data and to use only functions for constraints, thus distorting reality.

For example, consider the problem of computing the present value p of a future cash flow c that will be received after three years. If the annual future interest rates are r_1 , r_2 , and r_3 , then p can be computed by using the (discounting) formula:

$$p = c / ((1 + r_1/100) \cdot (1 + r_2/100) \cdot (1 + r_3/100)) \quad (1)$$

The problem is that future interest rates are volatile and that the value of c can be uncertain too. The value of p is then uncertain as well. The question is, How do we represent uncertain numeric values and how do we compute them?

Another major limitation of spreadsheets is that the relationships between cell values can only be expressed with functions evaluating output cell values from given input cell values. In the real world, things are more complicated. For example, consider the following formula for computing the y -coordinate of a projectile trajectory as a function of the x -coordinate, firing angle a , and initial velocity v .

$$y = x \cdot \tan(a) + 1/2 \cdot 9.81 \cdot a^2 / (v^2 \cdot \cos(a)^2) \quad (2)$$

Assume that the target is on a 120-meter-high hill (y) at a distance of 3200 meters (x). The initial speed (v) of the projectile is between 1250 meters a second and 1300 meters a second. The task is to find out what are the possible angles (a) between 0 and 90 degrees for hitting the target. The formula and the given data clearly provide the answer, but it is not clear how to back solve a from the function.

In a more general setting, the application problem might consist of a set of functions, equations, and inequalities, and the task is to solve any subset of variables involved, not only one variable. For example, what are the solutions to equations 3?

$$\begin{aligned} \sin(x_1) + \cos(x_2) &= \ln(x_3) \\ \cos(x_1) + 2 \cdot \ln(x_2) &= -\sin(x_3) + 3 \\ 3 \cdot \ln(x_1) &= \sin(x_2) - \cos(x_3) + 2 \end{aligned} \quad (3)$$

Traditional numeric techniques might find a solution to this problem but not necessarily. The success depends on the equations and the initial guess values used as the starting point for the iteration. At best, one solution is found for one starting point. The average spreadsheet user is not interested in such hidden technical details but simply wants to find a solution or

	A	B	C	D	E	F
1						
2		PRESENT VALUE (P)	R1 (%)	R2 (%)	R3 (%)	CASH FLOW (C)
3		[28.6, 32.8]	[4, 5]	[5, 7]	[6, 8.75]	[35, 38]
4						

Figure 1. Evaluating an EXCEL Formula with Interval Arguments.

all solutions to his/her problem by just pushing a button—or (s)he wants to be sure that the problem has no solutions at all!

These examples indicate that the following two key concepts of the current numeric spreadsheet paradigm are not flexible enough: (1) cell value and (2) formula.

Cell value: Only exact numbers can be used as values and can be evaluated by formulas. Given the widespread and diverse use of spreadsheets, a simple way for representing uncertainty is needed.

Formula: Only functions can be used as formulas that explicate the relations between cell values. Means for representing arbitrary constraints between variables involved are needed. Especially, equations and inequalities used everywhere in business, engineering, and science should be available.

The case study of this article shows that the idea of interval constraint solving developed in the fields of AI and interval analysis provides a new practical way to overcome these limitations. By generalizing the two core concepts of the spreadsheet paradigm—value and formula—a new basis for the very idea of using spreadsheets can be laid.

The new vision has been materialized as the commercial deployed add-in product INTERVAL SOLVER for Microsoft EXCEL, the result of some 16 person-years of research and development in Finland and Japan.

From the mathematical viewpoint, the solving power of the new technology is greater than with any traditional noninterval technique: All solutions (within a given precision level) to equations and other constraints can be found if enough time and memory are available. For example, INTERVAL SOLVER can actually prove that equation set 3 has exactly five different solutions. It suffices to push a button.

This article first explains why and when interval constraint solving and INTERVAL SOLVER are of use to a spreadsheet user. The architecture of the software is then presented, and the application of AI techniques is discussed. In conclusion, the significance of interval constraint technology in the development of the

current spreadsheet paradigm is discussed.

INTERVAL SOLVER for Microsoft EXCEL

INTERVAL SOLVER is an add-in that virtually extends the mathematical basis of EXCEL into interval constraint solving. From the user's viewpoint, he/she can make better use of imprecise real-world data and constraints and solve new kinds of problems that could not previously be addressed with spreadsheets. With INTERVAL SOLVER, one can bound worst and best cases satisfying the spreadsheet formulas, solve back argument intervals from given goals, solve equations and other constraints needed in the application, and find the best solution to a problem.

Bounding Worst and Best Cases

Intervals are perhaps the simplest way of representing uncertain numeric data. An interval $[min, max]$ is a continuum of values between the bounds *min* and *max*. For example, the interest rate of the next year can be estimated by interval $[4.0, 5.0]\%$, meaning any value between 4 percent and 5 percent. By using interval analysis, safe bounds for function values with interval arguments can be computed.

Figure 1 depicts the situation of equation 1 on an EXCEL sheet with the INTERVAL SOLVER add-in loaded. The formula for *P* (seen in the formula bar) has been computed with the given uncertain interest rate and cash flow intervals. All the user has to do is to write the EXCEL function inside the $=I(expression)$ formula of INTERVAL SOLVER. Then, interval arguments can be used.

The function value was initially *free*; that is, its value was unknown, which is equal to interval $(-\infty, \infty)$. INTERVAL SOLVER narrowed this value to $[28.6, 32.8]$, the interval that is guaranteed to bound all possible values of the function down to user-given precision. The minimum represents the global minimum and the maximum the global maximum of the function within the argument interval limits.

		=I(F3/((1+C3/100)*(1+D3/100)*(1+E3/100)))				
	A	B	C	D	E	F
1						
2		PRESENT VALUE (P)	R1 (%)	R2 (%)	R3 (%)	CASH FLOW (C)
3		29	[4, 5]	[5.7, 7]	[7.42, 8.75]	[35, 35.4]
4						

Figure 2. Interval Goal Seeking in INTERVAL SOLVER.

Notice that a number x is actually a collapsed interval $[x, x]$ having the same lower and upper bound, meaning that interval computations generalize traditional spreadsheet computations with exact values.

There are alternative approaches for representing uncertainty of numeric values too. A simple way is to enumerate scenarios. For example, in figure 1, one could compute the formula with, for example, all combinations of minimum and maximum values of the arguments. This simple approach is feasible only if the number of scenarios is small. In figure 1, there would be $2^4 = 16$ scenarios if only bounds were possibilities. Another problem is that usually it is difficult to say with what argument values the formula evaluates its global minimum or maximum that is often of greatest interest to the user. In the interval approach, all infinitely many scenarios are bounded within a single interval, and the actual global minimum and maximum can always be found.

A more sophisticated approach for representing numeric uncertainty is to use probability distributions as function arguments and then evaluate the functions using Monte Carlo simulation. This approach is widely used, and there are several software add-in packages available for spreadsheets, such as CRYSTAL BALL and @RISK.^{1,2} In the probabilistic view, an interval can be seen as a distribution whose form is completely unknown and whose definite integral over the interval is one. All variables are statistically independent from each other.

For the average spreadsheet user, probabilistic modeling might be too difficult to use. Intervals provide a simpler low-end approach for representing numeric uncertainty and, thus, have a better chance of being adopted by the spreadsheet users. Furthermore, intervals can be used for constraint solving, as is seen later. This is the key contribution of INTERVAL SOLVER.

Solving Back Argument Intervals

Assume that cell A1 contains the formula $= A2 + A3$. Given argument values A2 and A3, the value of A1 is computed. The computational

model of spreadsheets is a classic example of forward propagation.

However, in many problems, the goal is known, and the task is to back-solve argument values that lead to feasible solutions. Constraint propagation is a handy classical AI technique for solving such problems. For example, if A1 and one of the arguments, say A2, are known in this example, then the remaining argument, A3, can be computed ($A3 = A1 - A2$). This value can then be propagated further to formulas in which variable A3 is used, and so on. Constraint propagation makes it possible to evaluate formulas backwards or symmetrically, not only forward from known argument values to the function value.

The idea of constraint propagation is not new in spreadsheet computing. It was actually adopted by the early developers of the first major spreadsheet program, VISICALC, in the late 1970s. The best known result of this branch of development is TK!SOLVER,³ a tool for mathematical modeling.

In interval constraint solving, the classic numeric value propagation is generalized into a still more versatile computational model: Intervals are propagated instead of exact numbers. The idea is to narrow initial variable intervals by using local consistency filtering techniques developed originally for solving discrete constraint-satisfaction problems (CSPs). The result of the narrowing procedure is a set of intervals that definitely bound all exact solutions to the constraints, that is, the *solution set*.

For example, reconsider figure 1. The problem now is to determine the needed cash flow C and interest rates that would match a desired present value P. In figure 2, the user has set present value goal $P = 29$ in the situation of figure 1. In ordinary EXCEL, one cannot assign a formula and a value simultaneously to a cell, but with INTERVAL SOLVER, this can be done by double-clicking the cell. In response, EXCEL has refined two interest rates and the cash flow accordingly. Modified values are shown in bold for the user's convenience. INTERVAL SOLVER has bounded all possible scenarios that might lead to the goal within the given initial intervals—a

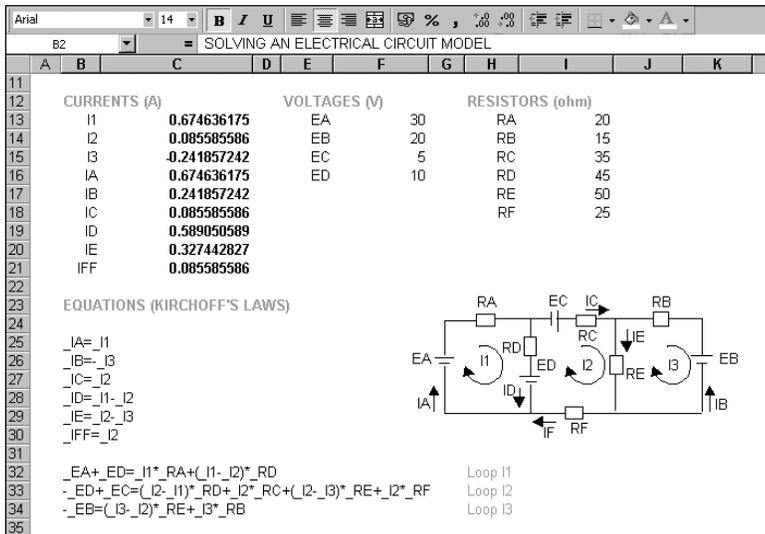


Figure 3. Solving the Nine Unknown Currents of an Electrical Circuit.

Cells have been given mnemonic names using EXCEL's Name command on the Insert menu.

useful piece of information for the decision maker.

INTERVAL SOLVER is capable of determining the global value interval of function formulas, but in back-solving, only locally consistent (Hyvönen 1992) bounds can be obtained in general. A problem of locally consistent bounds is that they can have excess width in certain situations; that is, local narrowing does not necessarily result in the narrowest intervals bounding all solutions. Various consistency criteria can be used for filtering, such as arc consistency, box consistency, 2B consistency, and 3B(w) consistency (Lebbah and Lhomme 1998), but the problem of excess width remains in the general case. In spite of this limitation, narrowed bounds provide insight to the user regarding the safe space of possibilities available. If needed, the following methods can be used for verifying the feasibility of an arbitrary point within a cell interval.

Solving Equations, Inequalities, and Other Constraints

The earlier goal-seeking example illustrated the idea that a spreadsheet formula is, from the mathematical viewpoint, actually a constraint equation. It tells how the function value and its arguments relate to each other, that is, what value combinations are mathematically possible. In the same spirit, the whole spreadsheet can be interpreted as a set of equations, inequalities, and other constraints whose variables have initial interval ranges. A spreadsheet thus formulates an interval CSP (ICSP). The natural task then is either to bound all solu-

tions of the constraint system or find its individual solutions.

This interpretation extends the usability of spreadsheet computing tremendously. Equations and inequalities can now be used on the sheet in addition to the traditional functions. In INTERVAL SOLVER, the expression inside the = $I(expression)$ formula cannot only be a function but also an equation or an inequality, such as

$$= I(A1^2*B2 = SIN(C1) + 3)$$

$$= I(LN(A1)^B1 > TAN(C1) C1^3)$$

One can also mix logical and numeric constraints, for example,

$$= I(IMPLIES(A1 > B1^2, AND(D2 = 0, A1 < SIN(B1))))$$

Solutions to the equations and logical constraints can be generated automatically by pushing a button. The spreadsheet has become an expert for equation solving and Boolean logic.

For example, figure 3 depicts a typical problem encountered in electrical designing. The circuit to be analyzed consists of batteries and resistors whose voltages and resistances are known. The task is to solve the nine currents C13:C21 shown on the sheet, which can be done in a standard way using Kirchoff's laws (equations) that relate resistances, voltages, and currents with each other. The equations are written and shown in cells B25:B30 (for the currents) and B32:B34 (for the voltages).

Because an equation does not have a numeric value, the equation itself is shown as the value of the corresponding = $I(equation)$ formula. Initially, values for the voltages and resistances were given, and values for currents were unknown; that is, they have very large interval values. The unique solution is found immediately and is shown in the figure.

Interval constraint-solving techniques differ from other numeric techniques in one important way: Possible solutions are never accidentally lost. As a result, all solutions can always be found if enough time and memory are available. Furthermore, if a situation is found infeasible, then the problem has no solutions for sure. This guarantee holds even when rounding errors are present. In interval computations, outward-rounding interval arithmetic is used, and imprecise floating-point numbers are represented by tiny safe intervals bounding the actual value.

For the spreadsheet user, this theoretical robustness is of great importance. Traditional numeric methods cannot, in general, guarantee that a solution will be found even if there is one. Convergence of iteration depends, for example, on the gradients of the equations and initial guess values for the variables. It is not

feasible to assume that a non-expert spreadsheet user understands the restrictions, conditions, and limitations related to traditional numeric equation-solving techniques. In interval solving, all solutions can, in principle, always be found.

For example, figure 4 depicts the problem of finding the solutions to a difficult nonlinear set of 12 kinematics equations. After evaluation, the View Solutions dialog box of INTERVAL SOLVER has popped up, and all 16 solutions to the equations can be viewed on the sheet. The user can be sure that this equation system has precisely these 16 solutions within the precision criteria used.

Finding the Best Solution to a Problem

In figures 1 and 2, intervals were used for bounding the solution set. After the system has narrowed the intervals, the user can constrain the problem further by inserting new constraints or by modifying the intervals. For example, in figure 2, the target present value P , was modified. After any modification, INTERVAL SOLVER might be able to narrow related intervals further. The user and INTERVAL SOLVER can work together in a mixed-initiative mode, and the problem can be solved in a top-down fashion by refining stepwise constraints for the solutions. This approach is not possible in traditional spreadsheet computing.

This approach can be used for finding the best solution to the problem at hand, that is, for solving optimization problems. The user sets desired goal values, INTERVAL SOLVER narrows related cell values, the user modifies them again according to his preferences, and so on. If the situation is found at some point to be infeasible, special relaxation (Hyvönen 1991) commands of INTERVAL SOLVER can be applied to enlarge intervals and make the bounds feasible again.

INTERVAL SOLVER also contains a tool for solving traditional optimization problems directly with the help of the SOLVER add-in that comes with Microsoft EXCEL. An interval CSP in INTERVAL SOLVER consists of the interval bounds set for the cell values, their value types (real or integer), equations, inequalities, and logical constraints written on the sheet. These constructs can be transformed into a classical EXCEL SOLVER model and be solved using EXCEL SOLVER. In this way, individual solutions can be found by which a given target function (cost function) gets its minimum, maximum, or a preset specific value, given a set of constraints. By bounding solutions first with the interval model, the initial guess values can be selected within a reasonable range, and the optimization problem

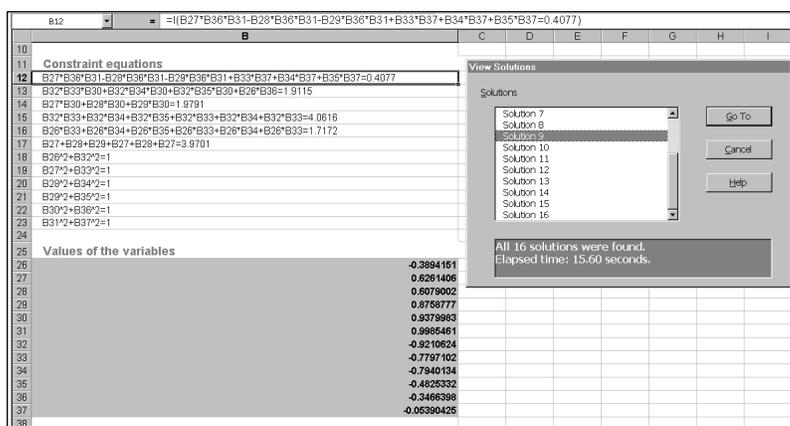


Figure 4. Finding All 16 Solutions to Kinematics Describing a Robot Arm.

The ninth solution is viewed.

can be solved more easily. INTERVAL SOLVER provides a natural way for expressing optimization problems. Any variable involved, not only the target cell value, can be optimized dynamically based on the interval model.

However, because EXCEL SOLVER is a classical optimization tool, interval techniques are not used, which means that the solution found might be only a local optimum and that only at most one solution corresponding to the set of initial guess values can be found. In general, there are no guarantees that a solution will be found even if there is one. Constrained interval optimization (Hansen 1992) provides a remedy to this problem and is planned to be made available in future releases of INTERVAL SOLVER as an alternative optimization tool.

Uses of AI Technology

The mathematical basis of INTERVAL SOLVER lays in the three InC++ interval libraries for C++, developed originally at VTT Technical Research Centre of Finland (Hyvönen and De Pascale 1995): (1) LIA INC++, (2) GIA INC++, and (3) ICE INC+ (table 1) (Delisoft 1998b).

Among these libraries, GIA and ICE are interesting from the AI viewpoint. Reconsider figure 1. The problem of determining the actual global minimum and maximum of P is easy in this case because the function happens to be monotonic. However, in the general nonmonotonic case (for example, equation 2), the problem is very difficult both from the algorithmic and computational viewpoints. The global min-max is then not obtained by a combination of argument interval limits.

There is only one class of numeric techniques that is guaranteed to always find the global minimum-maximum—global interval optimization techniques (Hansen 1992). These

```

#include <ice.h>
// Include ICE InC++ library header
main () {
    // Construct the equation set object I (of class Ice)
    Ice I;
    I.SetDefaultUnknown ("[-1e8, 1e8]");
    // Default interval bounds for variable values
    I.InsertConstraint ("x1^2+x2^2+x3^2-1=0");
    I.InsertConstraint ("x1^2+x2^2+x3^2-2*x1=0");
    I.InsertConstraint ("x1^2+x2^2-1=0");

    // Solve (evaluate) equation constraints
    I.SetPropagationMode(PMGlobal);
    // Set mode for finding individual solutions
    I.SetMaxSolutionNumber(1000);
    // Search for up to 1000 solutions
    ICEConsistencyType c;
    // Consistency type after evaluation
    I.Evaluate(&c);
    // Analyze and display the result
    if (c!=CTInfeasible)
        I.DisplayGlobalSolutions();
    return 0;
};

```

Figure 5. An Example of Programming with ICE INC++.

All solutions to a set of three equations in three variables are solved and displayed.

LIA INC++	Overloads C++ arithmetic into extended interval arithmetic.
GIA InC++	Evaluates the global value range of a function with interval arguments.
ICE INC+	Is an interval constraint-solving library based on LIA and GIA INC++.

Table 1. The Three InC++ Libraries for C++.

algorithms vary in detail, but the underlying idea in all of them is to perform an exhaustive branch-and-prune search in which the initial argument intervals are split into tighter and tighter subintervals until precision conditions for the solution are satisfied. The search is accelerated by various pruning heuristics based on the best min-max candidate found thus far, the mathematical properties of the function (such as first and second gradients), and special narrowing operators for the arguments such as the interval Newton operator). An indication

of the implementational complexities involved is that the GIA INC++ library used in INTERVAL SOLVER consists of over 50,000 lines of C++ code.

The key technology underlying INTERVAL SOLVER is interval constraint satisfaction (Hyvönen 1989; Cleary 1987; Davis 1987), developed in the fields of AI, (constraint) logic programming, and interval analysis.⁴ The ICSP corresponding to a sheet is represented as a C++ object of class Ice included in the ICE INC++ library. This class has a simple dynamic string-based interface by which constraints can be inserted and removed from the ICSP, interval domains set for the variables, precision criteria set for solutions, and so on. For an example of the programming interface, the C++ code in figure 5 shows how to solve three equations in three variables with ICE INC++.

User operations on a sheet, such as inserting a formula and setting a cell value, are mapped into sequences of member function calls of the underlying Ice object. This mapping is written in the macro language of EXCEL, VISUAL BASIC. One hundred seventeen different member functions of ICE INC++ library are used for the interface.

The ICE INC++ library does all mathematical constraint solving regarding formulas written inside the =I(expression). EXCEL itself maintains the algebraic formulas on the sheet mutually consistent, which makes INTERVAL SOLVER invisible to the user and the integration seamless. For example, changing cell names in formulas when copying, pasting, or moving cells is automatic as usual. User-defined cell names, as well as the different alias function names used in the various country versions of EXCEL, are available with INTERVAL SOLVER too. Figure 6 illustrates the general integration architecture.

The ICE INC++ library uses a large variety of interval constraint-solving techniques. The constraint set is manipulated and simplified by algebraic manipulation routines to make it easier to solve numerically. In numeric evaluation, the tolerance propagation approach (Hyvönen 1992) is enhanced with global interval optimization algorithms (Hansen 1992), narrow operators (Van Hentenryck, McAllester, and Kapur 1997; Van Hentenryck, Michel, and Deville 1997; Benhamou, McAllester, and Van Hentenryck 1994), conditioning matrixes (Kearfott 1996), and structure-sharing techniques (Hyvönen and De Pascale 1996b).

Interval constraint solving was originally proposed as a new computational basis for spreadsheet programs in Hyvönen (1991) and Hyvönen and De Pascale (1996a).

Application Use and Payoff

Interval constraint-solving technology has recently gained more and more attention not only in AI research but in business as well. There has been a lot of development activity in the logic programming community, resulting in several interval constraint extensions of Prolog, such as BNR Prolog and Prolog IA.^{5,6} The first stand-alone mathematical solvers based on new scheme, such as NUMERICA (VAN HENTENRYCK, MICHEL, AND DEVILLE 1997) and UNICALC,⁷ have been introduced in the market. These systems and tools are intended mainly for expert use and programmers.

In contrast, INTERVAL SOLVER (Delisoft 1998a) targets (also) non-expert users of spreadsheets. With the help of the new computational basis, the usage of spreadsheets has been extended to new classes of applications that deal with uncertain data or involve problem solving under user-given constraints, a typical situation in business planning, technical design, science, and many other fields.

INTERVAL SOLVER is a generic tool for end user spreadsheets. It virtually generalizes the arithmetical basis of EXCEL, and the range of applications is therefore as wide as that of spreadsheets. Most users indicate business computations as their main area of interest. Typical applications include cash flow analysis, budgeting, and risk analysis involving uncertain future data. In engineering applications, INTERVAL SOLVER provides the user with a simple tool for performing design and other calculations involving, for example, component data with tolerances, and for solving equations. The payoff comes from getting more realistic and better results to support decision making. Non-expert users can now solve—with little training in their customary spreadsheet environment—new classes of mathematically complicated problems.

Together with word processing, database, and internet tools, spreadsheets have been one of the most influential software applications of information technology. According to Gartner Group, some 30 million spreadsheet programs were shipped in 1997, most of which (90 percent) were copies of EXCEL, and the market is growing rapidly. There is the intriguing possibility that interval constraint technology will eventually lead to a paradigm shift in using spreadsheets.

Development and Deployment

In the late 1980s and early 1990s, research on using interval arithmetic as the basis for inter-

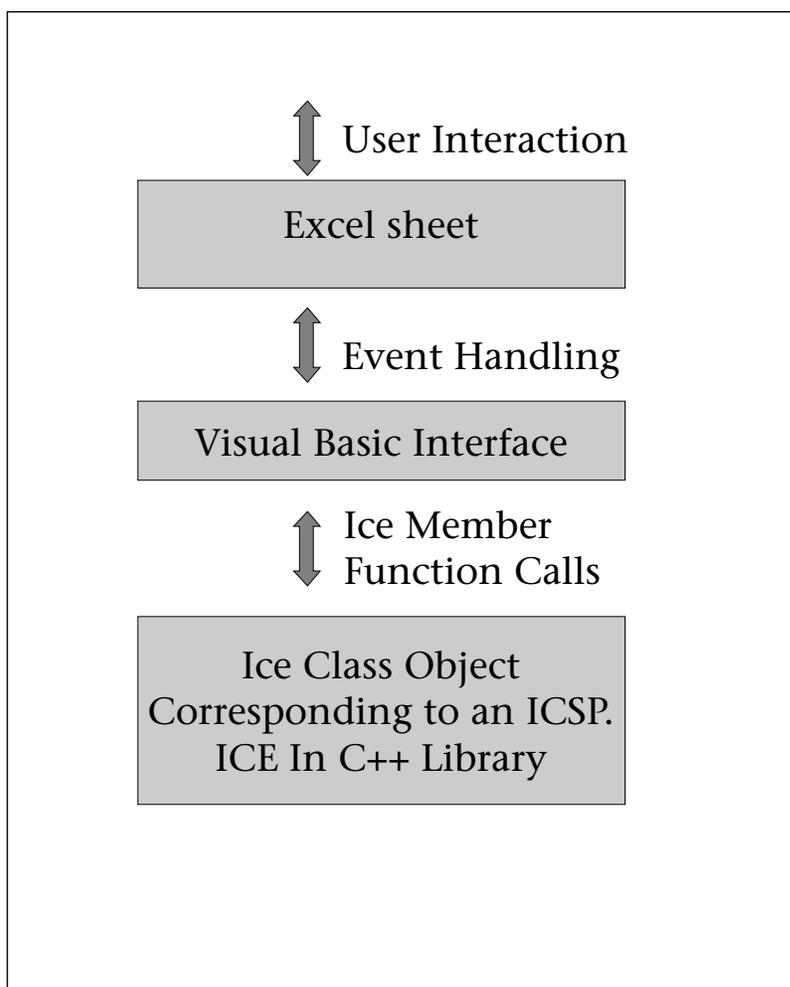


Figure 6. The Interaction Model of INTERVAL SOLVER.

The VISUAL BASIC interface layer catches user interactions (inserting data, moving cells, and so on) and commands. The ICE object, that is, the ICSP corresponding to the sheet, is updated accordingly or solved depending on the user interaction.

val constraint satisfaction was carried out at VTT Technical Research Centre in Finland and in a joint project with Electrotechnical Laboratory, Japan. A result of this work was an interval constraint solver and an interval spreadsheet demonstration system implemented in Lisp. Based on the first results, it was decided to implement the technology for industrial applications in C++ and apply it to a major commercial spreadsheet program, EXCEL.

It turned out, however, that the 16-bit address space provided by EXCEL at that time was too small for handling problems of reasonable size. Also, EXCEL's macro language turned out to be too limited for a commercial-level implementation of the new interval vision. With the new 32-bit WINDOWS versions and the new VISUAL BASIC macro language, the situation changed rapidly in 1996. The first implementa-

tion of interval constraints for EXCEL called RANGE SOLVER was exhibited by VTT at the CeBIT 96 fair in Hannover, Germany. INTERVAL SOLVER is its direct descendant, commercialized by Delisoft Ltd, a spin-off of VTT.

The first version of INTERVAL SOLVER 97 was finished during autumn 1997 and was released internationally in April 1998 at COMDEX Japan, Tokyo. The software evaluation kit has been available through various internet sites, including Ziff-Davis Libraries, Download.com, and Microsoft Office Update and also through some representatives in various countries (United States, Canada, Australia, Hong Kong, Korea, and so on). There have been several thousand downloaders all around the world.

The software consists of more than 100,000 lines of code, 80 percent of which is mathematical routines in C++. Several people were involved with the research on interval computations, but the actual code of INTERVAL SOLVER, as well as the manual, setup program and electronic tutorial, were written by the authors of this article.

The development, and especially the commercialization phases, of INTERVAL SOLVER were far more demanding than was initially expected. To meet the high efficiency requirements of spreadsheet users, the software had to be geared and tuned very carefully. Tiny modifications in the algorithms easily resulted in order-of-magnitude differences in performance. The computational efficiency of interval constraint-solving techniques is sensitive not only to the algebraic form of the ICSP but also to the initial interval values used. Various heuristics can be used to speed up convergence, but there is no single optimal strategy that always works well.

Besides the technical difficulties in implementing and tuning the mathematical constraint engine, lots of difficulties were encountered with the interface to EXCEL. A key problem was the enormous versatility of ways in which the user can interact with EXCEL and potentially confuse the system by making the sheet and the underlying ICSP model mutually incoherent. Most operations in EXCEL, such as inserting a formula, can be made in several alternative ways. All of them have to be caught. An additional practical problem was that a new fundamentally different version of the macro language provided by Microsoft for EXCEL 97 was released in the middle of the development process, causing redesign needs for the interface. Fortunately, the new version was more versatile from the INTERVAL SOLVER viewpoint. Last but not least, several deficiencies were encountered in different EXCEL versions. They

had to be circumscribed using special programming tricks.

Conclusions

Interval constraint solving and INTERVAL SOLVER provide a more versatile basis for spreadsheet computing. Two key concepts of spreadsheets have been generalized: First, the idea of cell value is generalized from exact numbers to intervals. An exact value is a special case of the notion of interval. Second, the idea of (function) formula is generalized into equations, inequalities, and logical constraints. A function is a special case of an equation.

From the computational viewpoint, the idea of forward propagation of exact values is generalized into interval constraint propagation. Again, forward propagation is a special case of the new model, interval constraint propagation.

Interval constraint solving is a conceptually simple, robust scheme for representing and solving difficult mathematical problems under uncertainty. Solutions are never lost such as when using traditional numeric techniques. A price to be paid for the robustness and the ease of use is increased computational complexity. However, results indicate that in many cases, interval constraint-solving methods can successfully compete with, or even outperform, the best traditional numeric techniques (Van Hentenryck, Michel, and Deville 1997).

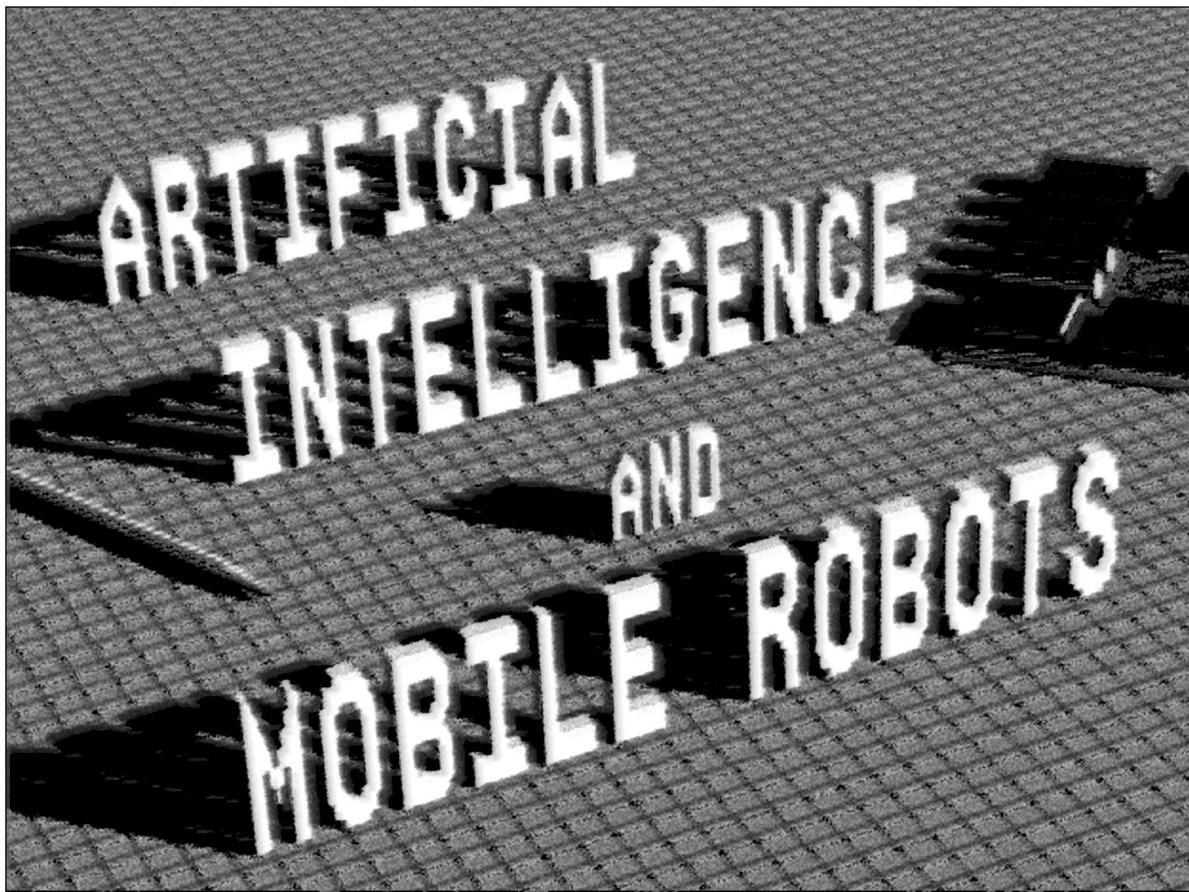
Spreadsheet programs are among the most widely used applications in information technology. However, after the pioneering days of VISICALC in 1979, their underlying computational idea has not changed much. INTERVAL SOLVER demonstrates that AI techniques can make a substantial contribution in the development of the spreadsheet paradigm.

Acknowledgments

Thanks to Eero Peltola, Technology Development Centre of Finland, Sitra, and VTT Information Technology for fruitful cooperation.

Notes

1. CRYSTAL BALL product information available at www.decisioneering.com.
2. @RISK product information available at www.palisade.com.
3. TK!SOLVER product information available at www.uts.com.
4. Home page of interval computations research is cs.utep.edu/interval-comp/main.html.
5. BNR Prolog home page is www.als.com/als/clpbnr/clp_info.html.
6. Prolog IA home page is prologianet.univ-mrs.fr/Us.



Artificial Intelligence and Mobile Robots

Case Studies of Successful Robot Systems

Edited by David Kortenkamp, R. Peter Bonasso, and Robin Murphy

The mobile robot systems described in this book were selected from among the best available implementations by leading universities and research laboratories. These are robots that have left the lab and been tested in natural and unknown environments. They perform many different tasks, from giving tours to collecting trash. Many have distinguished themselves (usually with first or second-place finishes at various indoor and outdoor mobile robot competitions).

Each case study is self-contained and includes detailed descriptions of important algorithms, including pseudo-code. Thus this volume serves as a recipe book for the design of successful mobile robot applications. Common themes include navigation and mapping, computer vision, and architecture.

6 x 9, 400 pp., ISBN 0-262-61137-6

**To order, call 800-356-0343 (US and Canada) or (617) 625-8569.
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**

7. UNICALC home page is txt.rriai.org.ru/UniCalc/about.html.

References

- Benhamou, F.; McAllester, D.; Van Hentenryck, P. 1994. CLP(Intervals) Revisited. Paper presented at the International Symposium on Logic Programming (ILPS-94), 13–17 November, Ithaca, New York.
- Cleary, J. 1987. Logical Arithmetic. *Future Computing Systems* 2(2): 125–149.
- Davis, E. 1987. Constraint Propagation with Interval Labels. *Artificial Intelligence* 32(8): 99–118.
- Delissoft. 1998a. Interval Solver User's Guide. Delissoft Ltd., Helsinki, Finland.
- Delissoft. 1998b. User's Guides of InC++ Libraries. Delissoft Ltd., Helsinki, Finland.
- Hansen, E. 1992. *Global Optimization Using Interval Analysis*. New York: Marcel Dekker.
- Hyvönen, E. 1992. Constraint Reasoning Based on Interval Arithmetic: The Tolerance Propagation Approach. *Artificial Intelligence* 58(1–3): 71–112.
- Hyvönen, E. 1991. Interval Constraint Spreadsheets for Financial Planning. In Proceedings of the First International Conference on Artificial Intelligence Applications on Wall Street. Washington, D.C.: IEEE Computer Society Press.
- Hyvönen, E. 1989. Constraint Reasoning Based on Interval Arithmetic. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 1193–1198. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.
- Hyvönen, E., and De Pascale, S. 1996a. Interval Computations on the Spreadsheet. In *Applications of Interval Computations*, eds. B. Kearfott and V. Kreinovich, 169–210. New York: Kluwer.
- Hyvönen E., and De Pascale, S. 1996b. Shared Computations for Efficient Interval Function Evaluation. In *Scientific Computing, Computer Arithmetic and Validated Numerics*, eds. G. Alefeld and E. Frommer. Berlin: Akademie-Verlag.
- Hyvönen, E., and De Pascale, S. 1995. InC++ Library Family for Interval Computations. Paper presented at the International Workshop on on Applications of Interval Computations, 23–25 February, El Paso, Texas.
- Kearfott, B. 1996. *Rigorous Global Search: Continuous Problems*. New York: Kluwer.
- Kearfott, B., and Kreinovich, V., eds. 1996. *Applications of Interval Computations*. New York: Kluwer.
- Lebbah, Y., and Lhomme O. 1998. Acceleration Methods for Numeric CSPs. In Proceedings of the Fifteenth National Conference on Artificial Intelligence, 19–24. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Van Hentenryck, P.; MacAllister, D.; and Kapur, D. 1997. Solving Polynomial Systems Using a Branch-and-Prune Approach. *SIAM Journal of Numerical Analysis* 34(2): 797–827.
- Van Hentenryck, P.; Michel, L.; and Deville, Y. 1997. *NUMERICA. A Modeling Language for Global Optimization*. Cambridge, Mass.: MIT Press.



Eero Hyvönen received his doctor of technology from Helsinki University of Technology and is professor of computer science at the University of Helsinki. His research interests include interval constraint-satisfaction techniques and AI. In addition to research, he has been active in developing and commercializing software innovations in the industries. His e-mail address is eero.hyvonen@cs.helsinki.fi.



Stefano De Pascale received his M.Sc. in computer science from the University of Helsinki. He has worked at VTT Technical Research Centre of Finland as a research scientist and was cofounder of Delissoft Ltd. as the research and development manager. De Pascale's research area covers interval and constraint-solving techniques. He is also interested in software development methods and tools and has been working on both research and industrial development projects. He is currently an independent consultant at Stepax.com. His e-mail address is stefano.depascale@stepax.com.