

Equilibrium Theory and Constraint Networks*

Francesco Ricci

Istituto per la Ricerca Scientifica e Tecnologica
38050 Povo (TN)
Italy
tel. 39-461-814444
e.mail: ricci%irst@uunet.uu.net

Abstract

This paper presents a new way to map a Constraint Satisfaction Problem (CSP) onto a non-cooperative game. Constraint Satisfaction Problems arise in many areas of Artificial Intelligence and have usually been tackled with backtracking based algorithms. The relevance of equilibrium theory in games with respect to the problem of finding a solution or a partial solution for a CSP is shown. The concepts of Nash equilibrium, admissible equilibrium and perfect equilibrium are examined. We prove that the solutions of a CSP are equilibrium points in pure strategy, both in the sense of Nash and in the sense of Selten. Finally, it is shown a CSP solution can be reached through a step by step evolutionary process, in which variables update the probability of assuming a value, considering what the other variables will do. Some computer experiments are reported in the paper.

1 Introduction

Artificial intelligence (AI) has started to look very closely at game theory in the last years. Many typical artificial intelligence problems seem to have a better representation in terms of agents, strategies, cooperation. Furthermore this approach enables an increase in efficiency by harnessing multiple reasoners to solve problems in parallel. A new sub-area of AI called Distributed Artificial Intelligence (DAI) has arisen and many DAI researchers have exploited game theoretic methodologies. The first aim of this discipline is to solve a collection of real world problems, such as distributed and parallel computing, computer supported cooperative work, computer-aided design and manufacturing, and many others, enabling many agents to work together on different aspects of the problem. It is out of the scope of this paper to comprehensively report on the full range of works in this area. The reader may find a good collection of up to date papers in [Gasser and Huhns, 1989].

The main concern of this paper is to show how game theory may be applied to a particular, well defined class of AI problems, namely Constraint Satisfaction Problems (CSP). In a constraint

*This work was supported in part by the Italian National Research Council (CNR) under the project "CNR-PNF Sistemi Informatici e Calcolo Parallelo. This work has been presented in two preliminary versions at the Constraint Directed Reasoning Workshop, AAAI-90, Boston 1990, and International Conference on Game theory, Florence, 1991.

I would like to thank Giorgio Pederzoli for his help in the earlier research. Special thanks to Paolo Avesani and Alfonso Gerevini for comments on a draft of this paper.

satisfaction problem a set of variables and a set of constraints among them is given, and the problem is solved when an assignment to all the variables is found that satisfies all the constraints. The classical approach consists of a solver that searches in the space of all the assignments for a correct one. We shall not enumerate the pros and cons of such an approach, as there is an extensive and rapidly growing literature that may be consulted (see [Nadel, 1989] and the references quoted). Our approach is completely different from the previous one and arises from the desire to solve CSP in a distributed environment. Often it happens that the variables in a CSP may be associated in turn to a set of agents, representing some attribute of the agent. For example, in a resource allocation problem, which can be formally described as a CSP, a variable usually represents the quantity of resource units allocated to the user associated to that variable. Humans usually tackle these problems engaging complex processes of negotiation between individual hypotheses, which may or not lead to an optimum [Sathi and Fox, 1989]. Conversely the classical CSP approach would generate an intelligent problem solver, which can control every parameter, generate every assignment, and ultimately come up with a solution. We have tried to build a model which is strongly inspired by human problem solving even if we should argue that rather we are proposing normative rules to CSP solving.

In our proposal a constraint satisfaction problem (CSP) C with n variables is mapped - in a quasi canonical way - onto an n player non-cooperative game $G(C)$, where every player of $G(C)$ corresponds to a variable of C , the pure strategies set of every player is identified with the set of values of the corresponding variable, and the utility function of a player rates the degree of satisfaction of an assignment with respect to the variable identified with that player. We show the relevance of the equilibrium theory in games with respect to the problem of finding a solution or a partial solution for the CSP the game came from. The concepts of Nash equilibrium, admissible equilibrium and perfect equilibrium (Selten) are examined. We prove that the solutions of the CSP are equilibrium points in pure strategy, both in the sense of Nash and in the sense of Selten.

Finally we examine the “tracing procedure”, proposed by J. Harsanyi and R. Selten [Harsanyi and Selten, 1988]. That procedure overcomes the naive bayesian approach to the prediction problem and it defines for every non-cooperative game a unique equilibrium point called the outcome of the game. The work of Harsanyi and Selten has been our starting point in defining an evolutionary process that shares some features with the tracing procedure. The solution or outcome of the CSP/Game is reached through a step by step process, in which the variables/agents update the probability to make a choice (assuming a value), reflecting on what the others will likely do. Whereas Harsanyi Selten tracing procedure is not elementary computable, since it requires the computation of the Nash equilibrium points of a generic family of non cooperative normal form games, the procedure we propose can be effectively computed and the computation may proceed in parallel on the set of variables.

The procedure we exhibit is to be considered only as an example of what can be done within a framework of non-cooperative bayesian reasoners that are only attracted toward more profitable strategies. Our a priori tenet is that the more rational a society of reasoners, the higher the total and medium payoff gained by the agents. Our research might be viewed as a first step toward an experimental confirmation of such a belief. It is quite surprising to note that “rational variables” thinking and acting as in the process shown in this paper, may achieve problem-solving without any form of cooperation and negotiation [Genesereth *et al.*, 1986], as our experimental results confirm.

Evolutionary processes have been studied by many authors both in artificial intelligence [Huberman and Hogg, 1988, Samuelson, 1988] and in more traditional contexts, such as economy [Brown, 1951] and biology [Smith, 1982]. Our contribution is more general in some aspects and more specific in others. All the previous papers other than [Huberman and Hogg, 1988] are confined to the two-player two-strategies case, while we consider the general case. Huberman and Hogg give a precise account of system dynamics considering noise and asynchronicity but with very simple payoff functions that by no means can model even a simple CSP. We have deliberately confined

ourselves to the synchronic and non stochastic case to study in a future paper what happens when one relaxes such requirements. Moreover those approaches are strongly influenced by population models, which in our computer experiments seem inadequate to CSP problems.

Finally, we note that the computer experiments seem encouraging. The technique works well both in determining partial solutions for CSPs that do not admit solutions and for transforming partial solutions into global solutions. We can imagine a possible application in all those domains where one is not only interested in getting a solution or a partial solution but one needs to rationally justify the single steps that have produced the solution.

2 Constraint Satisfaction Problems

Constraint Satisfaction Problems [Mackworth, 1987, Ricci, 1990] arise in many areas of AI including: vision [Waltz, 1975]; truth maintenance systems [de Kleer, 1989]; scheduling [Fox, 1987]; graphics [Borning *et al.*, 1987]; temporal reasoning [Dechter *et al.*, 1989]. We refer the interested reader to the quoted papers as it is impossible to give even a partial account of the matter here. We shall often refer in the following to the 8 Queens problem, which is a very simple instance of CSP. In this case 8 queens are to be placed on a chessboard in such a way that no queen attacks the others, that is there is no pair of queens placed on the same column, or on the same row or on the same diagonal.

Here we recall the basic definitions of constraint, constraint network and solution of a constraint network [Mackworth, 1977].

Let $\{X_1, \dots, X_n\}$ be a set of variables and suppose further that every X_i can take values in a set D_i , called the *domain* or the *label* of X_i . A k -ary *constraint* or *constraint-arc* on $\{X_1, \dots, X_n\}$ is a subset of $D_{i_1} \times \dots \times D_{i_k}$ where $i_j \neq i_l$ if $j \neq l$. If $r \subset D_{i_1} \times \dots \times D_{i_k}$ is a k -ary constraint we will say that r *connects* the variables $\{X_{i_1}, \dots, X_{i_k}\}$ and that $\{X_{i_1}, \dots, X_{i_k}\} = c(r)$ is the *connection* of r [Montanari and Rossi, 1989]. We shall also say that the variables X_i and X_j are connected if there exists a constraint r such that both X_i and X_j belong to $c(r)$. A k -constraint, from a "topological" point of view, is a labeled hyperarc, where the label contains the k -vectors allowed by the constraint, and the hyperarc expresses the ordering the nodes are connected in.

For example, in the 8 Queens problem, the column position of the queen placed on the i -th row can be represented by a variable X_i . The label of X_i is $D_i = \{1, \dots, 8\}$ and $X_i = j$ means that the queen on the i -th row is placed on the j th column. In particular the condition that two queens do not attack each other on a column is expressed by the binary constraint $r = \{(a_i, a_j) \in D_i \times D_j : a_i \neq a_j\}$. If $r \subset D_{i_1} \times \dots \times D_{i_k}$ is a k -ary constraint, let $\chi_r : D_{i_1} \times \dots \times D_{i_k} \rightarrow \{-1, 1\}$ be the (characteristic) function defined by the following condition:

$$\chi_r(d_1, \dots, d_n) = \begin{cases} 1 & \text{if } (d_1, \dots, d_n) \in r \\ 0 & \text{otherwise} \end{cases}$$

We shall also use the function $\chi_r^{a,b}$ which takes the value $a \geq 0$ on r and $b \leq 0$ outside r . Therefore we have $\chi_r = \chi_r^{1,0}$.

On a set of variables one can define more than one constraint, building what is called a constraint network. A *constraint network* or a *constraint satisfaction problem* (CSP) is a triple $C = \langle X, D, R \rangle$, where $X = \{X_1, \dots, X_n\}$ is a set of variables, $D = \{D_1, \dots, D_n\}$ is a set of labels and $R = \{r_1, \dots, r_m\}$ is a graded set of constraints, graded with respect to the arity of the constraints. We denote with $\alpha(r)$ the arity of r , and with $R_{[i]}$ the subset of R which contains all the relations that connect the i -th variable. Finally, to *solve* a CSP means to find an assignment of values to the variables in X such that all the relations in R are satisfied.

3 Noncooperative Game Theory

3.1 Definitions

In this section we shall review some basic definitions from noncooperative game theory in order to state the terminology.

From a formal viewpoint an n -person noncooperative game in normal form G is defined by giving:

1. n finite sets $S_i = \{\sigma_{i1}, \dots, \sigma_{im_i}\}$, $i = 1, \dots, n$, $m_i = |S_i|$;
2. n real valued functions $U_i : \prod_{j=1}^n S_j \rightarrow \mathbb{R}$.

The elements $\sigma_{ij} \in S_i$ are called the *pure strategies* of player i , and the function U_i is called the *payoff* or *utility function* of the i -th player. Broadly speaking $U_i(\sigma_{1k_1}, \dots, \sigma_{nk_n})$ represents the payoff that the i -th player gains if the j -th player uses strategy $\sigma_{jk_j} \in S_j$, for all $j = 1, \dots, n$.

A *mixed strategy* q_i of the i -th player is a probability measure on S_i , that is $q_i : S_i \rightarrow \mathbb{R}$ such that $q_i(\sigma_{ij}) \geq 0$ and $\sum_{j=1}^{m_i} q_i(\sigma_{ij}) = 1$. In the following we shall identify a mixed strategy q_i with the vector $(q_{i1}, \dots, q_{im_i})$, where $q_{ij} = q_i(\sigma_{ij})$ for all $j = 1, \dots, m_i$. The j -th component q_{ij} of a mixed strategy $q_i = (q_{i1}, \dots, q_{im_i})$ rates the probability that the pure strategy σ_{ij} will be used by the i -th player. We shall denote with $\Delta(S_i)$ the set of all mixed strategies of the player i . Moreover, we can suppose $S_i \subset \Delta(S_i)$ through the canonical inclusion

$$\iota : S_i \longrightarrow \Delta(S_i),$$

$\iota(\sigma_{ij}) = (q_{i1}, \dots, q_{im_i})$ such that, $q_{ij} = 1$ and $q_{ik} = 0$ if $k \neq j$.

The utility functions U_i can be extended by linearity on $\prod_{i=1}^n \Delta(S_i) = \Delta(S)$. If $q = (q_1, \dots, q_n) \in \prod_{i=1}^n \Delta(S_i)$, and $q_i = (q_{i1}, \dots, q_{im_i})$, then

$$U_i(q) = \sum_{j_1=1}^{m_1} \dots \sum_{j_n=1}^{m_n} U_i(\sigma_{1j_1}, \dots, \sigma_{nj_n}) q_{1j_1} \dots q_{nj_n}.$$

We shall call $\Delta(S)$ the *strategy space* of the game. $\Delta(S_{-i}) = \Delta(S_1) \times \dots \times \Delta(S_{i-1}) \times \Delta(S_{i+1}) \times \dots \times \Delta(S_n)$ is called the *i -incomplete strategy space*. If $q \in \Delta(S)$, then we shall denote with q_{-i} the element of $\Delta(S_{-i})$ obtained from q removing the i -th component.

Let q_i be an element of $\Delta(S_i)$ and let q_{-i} be an element of $\Delta(S_{-i})$, we shall denote with $q_i q_{-i}$ the element in $\Delta(S)$ obtained by “filling” the i -th missing component of q_{-i} with q_i . We shall say that $q_i^* \in \Delta(S_i)$ is a *best reply* to $q_{-i} \in \Delta(S_{-i})$ if and only if

$$U_i(q_i^* q_{-i}) \geq U_i(q_i q_{-i}), \forall q_i \in \Delta(S_i).$$

A combination of mixed strategies $q \in \Delta(S)$ is a *Nash equilibrium* if no player can gain by unilaterally switching to any other mixed strategy. That is, $q = (q_1, \dots, q_n) \in \Delta(S)$ is a Nash equilibrium if and only if for all $i = 1, \dots, n$:

$$U_i(q) \geq U_i(q_i' q_{-i}) = U_i(q_1, \dots, q_i', \dots, q_n), \forall q_i' \in \Delta(S_i).$$

In other words a combination of mixed strategies q is a Nash equilibrium if and only if every mixed strategy component q_i is a best reply to q_{-i} . So we can also say that $q \in \Delta(S)$ is a Nash equilibrium iff it is a best reply to itself. A fundamental result of Nash[Nash, 1951] assures that every game has at least one equilibrium point (in mixed strategy).

3.2 Viewing CSP as Games

Now we can finally define the game associated to a given CSP. Although, there is no canonical way to do this, we shall concentrate on a particular and “natural” way to transform a CSP into a game.

Let $C = \langle X, D, R \rangle$ be a given CSP, we define the non-cooperative game $G(C) = (S_1, \dots, S_n; U_1, \dots, U_n)$, where:

1. $n = |X|$;
2. $S_i = D_i, \forall i = 1, \dots, n$;
3. $U_i(d_1, \dots, d_n) = \sum_{r \in R_{[i]}} k(r) \chi_r(d_{j_1}, \dots, d_{j_{k(r)}})$;

where $k(r)$ is the arity of the relation $r \subset D_{j_1} \times \dots \times D_{j_{k(r)}}$, and $R_{[i]}$ is the set of all relations which connect the i variable. In such a way, $G(C)$ has a player for every variable of C , and the pure strategies of this player are the elements of the label of the corresponding variable. The payoff function of a player/variable counts the number of satisfied constraints connecting that variable, taking every constraint along with its arity. In this way, the more constraints satisfied, the more utility a player gets.

In general, we could have a different weight $w(r)$ for every constraint, reflecting a hierarchy of constraints, and this weight could also depend on the player. Moreover we may want to produce a loss of utility when a constraint is not satisfied. In this case, given $a \geq 0$ and $b \leq 0$, the utility functions would be replaced with the following:

$$U_i(d_1, \dots, d_n) = \sum_{r \in R_{[i]}} w_i(r) \chi_r^{a,b}(d_{j_1}, \dots, d_{j_k}). \quad (1)$$

In the next section we shall discuss the relevance of equilibrium theory to CSP, and to the solution problem in particular. But as one can intuitively guess, the more the utility functions reflect the satisfaction of constraints, the more the equilibrium reflects the solution. Conversely, the more freely we define the weights, the farther the equilibrium is from a solution. If no otherwise specified, in the following we shall assume the game $G(C)$ utility functions are defined as in (1), but all the results proved in this paper holds if we replace the functions χ_r with $\chi_r^{a,b}$, assuming that $a \neq 0$ or $b \neq 0$.

In the same way as before, the utility functions can be linearly extended to the space $\Delta(S) = \prod_{i=1}^n \Delta(D_i)$. If $q = (q_1, \dots, q_n) \in \prod_{i=1}^n \Delta(D_i)$, then we have

$$\begin{aligned} U_i(q) &= \sum_{d_1 \in D_1} \dots \sum_{d_n \in D_n} U_i(d_1, \dots, d_n) q_1(d_1) \dots q_n(d_n) \\ &= \sum_{d_1 \in D_1} \dots \sum_{d_n \in D_n} \left(\sum_{r \in R_{[i]}} k(r) \chi_r(d_{i_1}, \dots, d_{i_{k(r)}}) \right) q_1(d_1) \dots q_n(d_n). \end{aligned}$$

The following lemma can be easily verified, and gives an idea of the resources needed to compute the utility functions.

Lemma 1 *Let $G(C)$ be the game associated to the CSP C . If the variable X_i is not connected to the variable X_j then $U_i(q_1, \dots, q_n)$ does not depend on q_j .*

4 Equilibria and Solutions

The equilibrium theory in non-cooperative games started with the work of J.Nash [Nash, 1951] who first introduced the concept of equilibrium and proved the existence of it, in every finite game.

Unfortunately a game may have many equilibria, and some of them are irrational, in the sense that no rational player will agree on them. Therefore the problem arises in selecting between different equilibria. In this section we shall follow the way to refine the concept of equilibrium, to remove some points, and we shall investigate the meaning of this refinement from the CSP viewpoint.

4.1 Nash Equilibria and Solutions

We start by illustrating the meaning of a Nash equilibrium for the game $G(C)$ associated to a CSP C . First of all, we can observe that if a CSP C has a solution $(d_1, \dots, d_n) \in \prod_{i=1}^n D_i$, then if $G(C) = (S_1, \dots, S_n; U_1, \dots, U_n)$ we have $U_i(d_1, \dots, d_n) = \max_{q \in \Delta(S)} [U_i(q)]$, for all $i = 1, \dots, n$. This immediately entails the following.

Proposition 1 *If C is a CSP, then every solution of C is a Nash equilibrium of $G(C)$.*

The converse of this proposition does not hold. In fact we may have a CSP C without solutions even if $G(C)$ always has equilibria. But the situation may be even worse. Let C be the 8-queen problem, where X_i represents the column position of the queen place on the i -th row. We have $S_i = D_i = \{1, 2, 3, 4, 5, 6, 7, 8\}$, and $U_i(d_1, \dots, d_n)$ is proportional to the number of queens which do not attack the i -th queen, when $X_i = d_i$ (we are assuming that three binary constraints exist between every pair of queens, and they forbid that two queens be placed on the same column or on the same diagonal). In Figure 1 we can see an example of equilibrium point that is not a solution ($X_1 = 4, X_2 = 1, X_3 = 5, X_4 = 2, X_5 = 8, X_6 = 6, X_7 = 3, X_8 = 7$).

In fact, let us indicate with q^* this point. We have $q_{14}^* = q_{21}^* = q_{35}^* = q_{42}^* = q_{58}^* = q_{66}^* = q_{73}^* = q_{87}^* = 1$ and all the other q_{ij} are 0. For example $U_1(q^*) = 40$ because the only constraint that is violated is that with the queen placed on the 5-th row (rows are counted from above and columns are counted from left). If we replace q_1^* , that is, if we change the position of the first queen we cannot improve U_1 , as the number of satisfied constraints does not increase. Therefore also for any convex combination of pure strategies we have $U_1(q^*) \geq U_1(q_1 q_{-1}^*)$, for all $q_1 \in \Delta(S_1)$. Reasoning analogously for the other queens we prove that q^* is a Nash equilibrium.

From this example it is clear that Nash equilibrium points represent an assignment to the variables that cannot be improved (in terms of number of satisfied constraints) modifying one variable at a time. The players/variables should bargain modifications, communicating their intentions. But this is only one possible way, in a following section we shall show how to achieve the more rational equilibrium simulating mutual repeated interactions that take into account the a priori beliefs about what a player will do and the actual outcome of the previous interaction.

4.2 Admissible Equilibrium and Partial Solutions

A first refinement of the concept of equilibrium is that of admissible equilibrium. The idea is to compare different points to choose a better equilibrium. If $q, q' \in \Delta(S)$ are two combinations of mixed strategies, we say that q is *better* than q' iff $U_i(q) \geq U_i(q')$, for all $i = 1, \dots, n$, and at least one of these inequalities is strict. We shall say that a Nash Equilibrium is *admissible* if there exists no better Nash equilibrium.

Proposition 2 *If C is a CSP, every solution of C is an admissible Nash equilibrium of $G(C)$. Moreover, if the set of solutions of C is not empty then it coincides with the set of admissible Nash equilibrium points of $G(C)$.*

Proof. The proof follows trivially from the observation that if $q^* = (d_1, \dots, d_n) \in \prod_{i=1}^n D_i$ is a solution, then $U_i(q^*) = \max_{q \in \Delta(S)} [U_i(q)]$

If C has no solution, then the set of all instantiations $(d_1, \dots, d_n) \in \prod_{i=1}^n D_i$ which are admissible Nash equilibrium points in $G(C)$, might represent the set of “best” partial solutions of the CSP. From the previous proposition one can conclude admissible equilibrium points are the most interesting points to search. Unfortunately elementary methods to search admissible equilibrium points do not exist. The reason resides in the fact that whereas Nash equilibrium (and perfect equilibrium as we shall see in the next section) is a local concept, that is can be calculated looking only into a neighbourhood of the point, admissible equilibrium is a true global concept and it requires global maximization techniques.

4.3 Perfect Equilibrium

A second refinement of the Nash equilibrium concept has been proposed by R. Selten [Selten, 1975]. He argues that many equilibrium points require some or all of the players to use highly irrational strategies. As an example, let G be the 2 players game described by the matrix in Figure 2. In that game player 1 has the strategies $\{a_1, b_1\}$, and player 2 has the strategies $\{a_2, b_2\}$. In the upper left corner and in the lower right corner of a box are displayed the payoffs of the first and second player respectively. As one can easily show, (a_1, a_2) and (b_1, b_2) are both Nash equilibrium points, but the second one is irrational and will not be adopted if there is only a little probability that one player will adopt the first strategy. This is the essential idea behind Selten’s perfect equilibrium, that is no strategy should ever be given zero probability, since there is always a small chance that any strategy might be chosen, if only by mistake.

But let us turn to CSP, and note that the previous game may derive from a CSP with two variables $\{X_1, X_2\}$, labels $D_1 = \{a_1, b_1\}$, $D_2 = \{a_2, b_2\}$ and a unique binary constraint r which is satisfied by the pair (a_1, a_2) . The point (b_1, b_2) is an equilibrium but is not a solution.

We now go into a formal definition of perfectness and we show that the solutions of C are perfect equilibrium point of $G(C)$. To define a perfect equilibrium point we use a characterization due to Myerson [Myerson, 1978]. Let $\Delta^0(S)$ be the interior part of $\Delta(S)$, that is

$$\Delta^0(S) = \prod_{i=1}^n \Delta^0(S_i),$$

and

$$\Delta^0(S_i) = \{(q_{i1}, \dots, q_{im_i}) \in \Delta(S_i) : q_{ij} > 0, \forall j = 1, \dots, m_i\}.$$

Let ϵ be a positive constant, we define an ϵ -perfect equilibrium to be a point $q = (q_1, \dots, q_n) \in \Delta^0(S)$ such that:

$$\text{if } U_i(\sigma_{ij}q_{-i}) < U_i(\sigma_{ik}q_{-i}) \text{ then } q_{ij} \leq \epsilon, \forall i = 1, \dots, n, \forall j, k \leq m_i.$$

A perfect equilibrium is defined to be any limit point of ϵ -perfect equilibria. $q \in \Delta(S)$ is a *perfect equilibrium* point if and only if there exist two sequences $\{\epsilon_k\}_{k=1}^\infty$ and $\{q^k\}_{k=1}^\infty$ such that:

1. $\epsilon_k > 0, \forall k = 1, \dots, \infty$, and $\lim_k \epsilon_k = 0$;
2. q^k is a ϵ_k -perfect equilibrium $\forall k \in \mathbb{N}$;
3. $\lim_k q_{ij}^k = q_{ij}, \forall i = 1, \dots, n, \forall j = 1, \dots, m_i$.

Selten has proved that every perfect equilibrium is also a Nash equilibrium, but the converse does not hold. In the example shown above, only the point $q = (a_1, a_2)$ is perfect. To see that q is perfect, let $q_i^\epsilon = (1-\epsilon)a_i + \epsilon b_i$. We have $\lim_{\epsilon \rightarrow 0} (q_1^\epsilon, q_2^\epsilon) = q$, $U_i(a_i q_{-i}^\epsilon) = 2(1-\epsilon)$, and $U_i(b_i q_{-i}^\epsilon) = 0$. Therefore if $U_i(\sigma_{ij}q_{-i}) < U_i(\sigma_{ik}q_{-i})$ we get $\sigma_{ij} = b_i$ and from the definition of q_i^ϵ , $q_{i2}^\epsilon = \epsilon$. With the same arguments one sees that (b_1, b_2) is not perfect.

We now see the relation between solutions of a CSP and perfect equilibria.




Figure 1: An equilibrium point that is not a solution.

	a_2	b_2
a_1	2	0
	2	0
b_1	0	0
	0	0

Figure 2:

Proposition 3 *If $q = (x_1, \dots, x_n)$ is a solution of C then q is also a perfect equilibrium point of $G(C)$.*

Proof. Let $\{q^k\}_{k=m}^\infty$ be the succession in $\Delta^0(S)$, where $m = \max_i[m_i]$, defined by the following conditions:

$$q_{ij}^k = \begin{cases} 1 - (m_i - 1)/k & \text{if } \sigma_{ij} = x_i \\ 1/k & \text{otherwise} \end{cases}$$

It is clear that $\{q^k\}_{k=1}^\infty$ converges to the solution q . We show now that q^k is an ϵ_k -perfect equilibrium point, and $\lim_k \epsilon_k = 0$. Let us suppose

$$U_i(\sigma_{ij}q_{-i}^k) < U_i(\sigma_{il}q_{-i}^k). \quad (2)$$

If $\sigma_{ij} \neq x_i$ then $q_{ij}^k = 1/k$ therefore we can choose $\epsilon_k = 1/k$. We can conclude showing that we cannot have $\sigma_{ij} = x_i$. We shall reason as absurd, and we shall suppose $\sigma_{ij} = x_i$. We cannot have $\sigma_{ij} = \sigma_{il}$ because (2). Therefore let $j \neq l$, if $(x_1, \dots, x_{i-1}, \sigma_{il}, x_{i+1}, \dots, x_n)$ is a solution of C , we have $U_i(\sigma_{ij}q_{-i}^k) = U_i(\sigma_{il}q_{-i}^k)$ because the payoff functions depend only on the number of satisfied constraints and not on the particular value assigned to the variables. But if $(x_1, \dots, \sigma_{il}, \dots, x_n)$ is not a solution then $U_i(q) = M > M' = U_i(x_1, \dots, \sigma_{il}, \dots, x_n)$, and we have $U_i(\sigma_{il}q_{-i}^k) = (1 - (m_i - 1)/k)^n M' + R'_k$, with $\lim_k R'_k = 0$. Therefore there exists a k^* such that for all $k > k^*$ (2) does not hold because it must also hold that $U_i(\sigma_{ij}q_{-i}^k) = (1 - (m_i - 1)/k)^n M + R_k$ with $\lim_k R_k = 0$

5 Equilibrium Selection

5.1 The Tracing Procedure

Now everything is ready to define a procedure that models a process of mutual influence by which rational players will come to adopt and to expect each other to adopt, one particular equilibrium point. The following description is based on the fourth chapter of the book by Harsanyi and Selten [Harsanyi and Selten, 1988], and we invite the reader to consult the original work if interested in more details.

At the beginning of this process we can suppose that each player expresses his expectations about the pure strategies of every other player. Therefore, n probability vectors $p_i = (p_{i1}, \dots, p_{im_i}) \in \Delta(S_i)$ are given, where p_{ij} is the probability assigned by all the other players that i will adopt the pure strategy σ_{ij} . We call these p_i *prior probability distributions*. The prediction problem in game theory may be stated as follows. Suppose that the prior probability distributions are known. How can the strategies the players will actually use be predicted? From the CSP point of view this is equivalent in saying: if a CSP is in a state defined by a probability distribution for each variable, and supposing further that every variable/agent can perform only a certain amount of computation, what solution, if it does exist, will be ultimately reached? In this paper we have decided to devote ourselves to investigating only bayesian rationality, that is we assume each variable/agent makes choices using essentially classical decision theory. But we would like to note the very aim of this research is to experiment useful combinations of symbolic and numeric (bayesian) reasoning. We shall discuss this subject in more details in a future paper.

The naive bayesian approach says that every player i will use the strategy q_i that maximizes his payoff against the strategies p_{-i} supposed to be adopted by the other players. Therefore i will use the best reply strategy q_i^* :

$$U_i(q_i^* p_{-i}) = \max_{q_i \in \Delta(S_i)} [U_i(q_i p_{-i})].$$

But in general the point $q^* = (q_1^*, \dots, q_n^*)$ is not an equilibrium point, so it cannot be the rational outcome of the game. We should refine this approach, taking into consideration that the best reply strategies influence the players i not only to act against p_{-i} but to face a combination of p_{-i} with q_{-i}^* . We can imagine a recomputation of beliefs about the other players choices, provoked by the new information q^* . The tracing procedure is an example of continuous recomputation of the probability of every single choice, which has the fundamental property in “almost all” cases of converging to an equilibrium.

If G is a game, let $\{G(t, p)\}$, $0 \leq t \leq 1$, be the family of games with the same strategy space as G , but with the payoff functions:

$$U_i^t(q_i q_{-i}) = tU_i(q_i q_{-i}) + (1 - t)U_i(q_i p_{-i}).$$

Harsanyi and Selten have shown that for almost all p (with respect to the usual Lebesgue measure on $\Delta(S)$) $G(0, p)$ has exactly one equilibrium point q^0 . For each $t \in [0, 1]$ let $E^t(p)$ be the set of equilibrium points of $G(t, p)$. If $\gamma : [0, 1] \rightarrow \Delta(S)$ is a continuous path, such that $\gamma(t) \in E^t(p)$, $\gamma(0) = q^0$ and $\gamma(1) = q^*$, then we say that γ is a *feasible path* and we call $\gamma(1)$ the outcome of γ . The tracing procedure consists in selecting an outcome $q^* \in \Delta(S)$ by following a feasible path from its starting point to its end point.

The linear tracing procedure for (G, p) is called *feasible* if at least one feasible path exists, and is called *well defined* if exactly one feasible path exists. The following propositions are proved in [Harsanyi and Selten, 1988].

Proposition 4 *For any possible pair (G, p) , the linear tracing procedure is always feasible but is not always well defined.*

Proposition 5 *For any specific vector p of prior probability distributions, almost all games G will give rise to a well defined linear tracing procedure.*

To overcome the limitations indicated in the above propositions, Harsanyi and Selten have proposed a modification of $G(t, p)$, in which they add to every U_i^t a logarithmic term. This term eliminates bifurcations in $\bigcup_{0 \leq t \leq 1} E^t(p)$. For our purpose, it is enough to say that one can use knowledge to choose between alternatives, and in every case, with a slight perturbation of the game we can obtain (with probability 1) a new game G_ϵ for which the linear procedure is well defined. Moreover when the game comes from a CSP, payoffs are integer valued functions, and from continuity considerations one can prove this perturbation does not modify the equilibrium points. Therefore solutions of the CSP C will still be equilibrium points of $G(C)_\epsilon$.

The major difficulty of the tracing procedure is that in general we are not able to compute $E^t(p)$ and the computation cannot be done in parallel by the n players. In fact, whereas $G^0(p)$ decomposes into n mutually independent maximization problems, the equilibrium points of $G^t(p)$ in general cannot be found so easily.

5.2 Rational Processes

The tracing procedure can be viewed as a first example of a large spectrum of evolutionary procedures that model the progressive updating of the probability that a strategy will be used. In every single step the player is assumed to compute, using only a certain amount of computational resource, what the other players will presumably do. Then he tries to maximize his utility choosing the best strategy at his disposal. This process of iterated updating of beliefs (what the others will do and what I should do) in some cases converges to a stationary point which is the simulated solution of the game. In the following we shall see how some simple hypothesis on the computation done by the players in each step can yield strong conclusions on the properties of the point reached by the process. In a certain sense the more rational the players the more profitable the equilibrium.

We have been strongly inspired by L. Samuelson's paper [Samuelson, 1988], devoted to 2-player games, but we have modified many concepts to better face our problem, which originates from constraint networks. Moreover, we have proved different results.

Let us start with some definitions. A differentiable path $q : [0, +\infty) \rightarrow \Delta(S)$ will be called a *process*. A process models the time dependent nature of the belief a strategy is to be used.

We shall say that a process q is *monotonic* if and only if, for all $i = 1, \dots, n$, when $U_i(\sigma_{ij}q_{-i}(t)) \leq U_i(\sigma_{ik}q_{-i}(t))$, for some $k = 1, \dots, m_i$, then $dq_{ij}/dt \leq 0$.

We shall say that a process q is *rational* if and only if, for all $i = 1, \dots, n$, when $U_i(\sigma_{ij}q_{-i}(t)) < \max_k [U_i(\sigma_{ik}q_{-i}(t))]$ then $dq_{ij}/dt < 0$.

We shall say that a process q is *asymptotically rational* if and only if it converges and if there exist $\tau > 0$ s.t. for all $t \geq \tau$:

1. $U_i(\sigma_{ij}q_{-i}(t)) \leq U_i(\sigma_{ik}q_{-i}(t))$;
2. $\lim_{t \rightarrow \infty} q_{ij}(t) > 0$;

then $\lim_{t \rightarrow \infty} U_i(\sigma_{ij}q_{-i}(t)) = \lim_{t \rightarrow \infty} U_i(\sigma_{ik}q_{-i}(t))$.

Asymptotical rationality does not imply rationality and there exist rational processes which converge and are rational but that are not asymptotically rational. For example, let G be a 2 players game where player 1 has two pure strategies σ_{11}, σ_{12} which always give payoffs α_1, α_2 respectively, with $\alpha_1 < \alpha_2$. If q is a process with $dq_{11}/dt = -1/t^2$, for $t \geq 2$, and $q_{11}(2) = 3/4$, then we have $q_{11}(t) = 1/t + 1/4$, and $q_{12}(t) = 1 - q_{11}(t) = 3/4 - 1/t$. Such a q is rational but not asymptotically rational.

The following lemma shows that in a monotonic process the probability/belief to adopt a best reply strategy does not decrease. That is, in a monotonic process a best reply strategy will never be punished.

Lemma 2 *If q is a monotonic process and σ_{ij} is a best reply strategy for $t = t_0$, then $dq_{ij}/dt(t_0) \geq 0$.*

Proof. The proof follows immediately from the fact $\sum_j q_{ij} = 1$. In fact, let σ_{ij} be a best reply strategy for player i at time t_0 , then we have $U_i(\sigma_{ij}q_{-i}(t_0)) \geq U_i(\sigma_{ik}q_{-i}(t_0))$, for all $k \neq j$. From the monotonic assumption this yields $dq_{ik}/dt(t_0) \leq 0$, for all $k \neq j$, and from $\sum_k dq_{ik}/dt = 0$ we get $dq_{ij}/dt \geq 0$

In the following we shall consider discrete processes, that is maps from \mathbb{N} to $\Delta(S)$. All the previous definitions have a discrete translation. For example if $\{q^n\}_{n=0}^\infty$ is a discrete process we shall call it monotonic iff $U_i(\sigma_{ij}q_{-i}^n) \leq U_i(\sigma_{ik}q_{-i}^n)$, for some $k = 1, \dots, m_i$, implies $q_{ij}^{n+1} \leq q_{ij}^n$.

Asymptotically rational processes exploit just the right amount of rationality to yield Nash equilibrium points. This is proved in the following proposition.

Proposition 6 *Let q be a convergent process. $q^* = \lim_{t \rightarrow \infty} q(t)$ is a Nash equilibrium point if and only if q is asymptotically rational.*

Proof. To prove this proposition we recall that a point q^* is a Nash equilibrium point iff for every pure strategy σ_{ij} s.t. $q_{ij}^* > 0$ we have $U_i(\sigma_{ij}q_{-i}^*) \geq U_i(\sigma_{ik}q_{-i}^*)$, for all $k = 1, \dots, m_i$. Let q^* be a Nash equilibrium point, $U_i(\sigma_{ij}q_{-i}(t)) \leq U_i(\sigma_{ik}q_{-i}(t))$ and $\lim_{t \rightarrow \infty} q(t)_{ij} = q_{ij}^* > 0$. From this last we get $U_i(\sigma_{ij}q_{-i}^*) \geq U_i(\sigma_{ik}q_{-i}^*)$ and from the two inequalities we get

$$\lim_{t \rightarrow \infty} U_i(\sigma_{ij}q_{-i}(t)) = U_i(\sigma_{ij}q_{-i}^*(t)) = U_i(\sigma_{ik}q_{-i}^*(t)) = \lim_{t \rightarrow \infty} U_i(\sigma_{ik}q_{-i}(t)).$$

Vice versa, let q be asymptotically rational and $\lim_{t \rightarrow \infty} q_{ij}(t) = q_{ij}^* > 0$. Suppose further that $U_i(\sigma_{ij}q_{-i}^*) = \max_k [U_i(\sigma_{ik}q_{-i}^*)]$ does not hold. Then there must exist σ_{ik} such that $U_i(\sigma_{ij}q_{-i}^*) <$

$U_i(\sigma_{ik}q_{-i}^*)$. This fact implies that there exists $\tau > 0$ s.t. for all $t \geq \tau$, $U_i(\sigma_{ij}q_{-i}(t)) < U_i(\sigma_{ik}q_{-i}(t))$. But from the definition of asymptotic rationality we must have

$$\lim_{t \rightarrow \infty} U_i(\sigma_{ij}q_{-i}(t)) = \lim_{t \rightarrow \infty} U_i(\sigma_{ik}q_{-i}(t)),$$

and this is a contradiction

Now we are going to define a particular process that can be viewed as the “natural” version of the Harsanyi-Selten tracing procedure with no upper temporal bound and with a time delay between the evaluation of the best strategy and the updating of beliefs. First of all if G is a game and $q \in \Delta(S)$, we shall indicate with $BR(q)$ one of the best reply strategies to q , that is $BR(q) \in \Delta(S)$ and if $U_i(\sigma_{ij}q_{-i}) < U_i(q)$ then $BR(q)_{ij} = 0$. We note $BR(q)$ is always in pure strategies, namely $BR(q)_{ij} \in \{0, 1\}$, as the payoff functions U_i are linear and the constraints $\sum_j BR(q)_{ij} = 1$ define a convex set.

Now let p be a prior probability distribution and $0 < \epsilon \leq 1$ a positive constant. We define inductively a discrete process $P_{HS}(G, p, \epsilon) = \{p^n\}_{n=0}^\infty$, letting:

$$\begin{aligned} p^0 &= p; \\ p^n &= \epsilon BR(p^{n-1}) + (1 - \epsilon)p^{n-1}. \end{aligned}$$

We observe $P_{HS}(G, p, \epsilon)$ is a synchronous process, where at every step each agent recomputes his current probability distribution over the set of his strategies using the probability distributions of the other agents at the step before. In the following lemmas we prove two useful properties of such a process.

Lemma 3 *Let G be a game, p a prior distribution probability, and $\{p^n\} = P_{HS}(G, p, \epsilon)$. If $p_{ij}^{n+1} > p_{ij}^n$ then σ_{ij} is a best reply strategy to p^n .*

Proof. We can note the following condition holds:

$$p_{ij}^{n+1} = \begin{cases} \epsilon + (1 - \epsilon)p_{ij}^n & \text{if } BR(p^n)_{ij} = 1 \\ (1 - \epsilon)p_{ij}^n & \text{otherwise} \end{cases}$$

Moreover we have $\epsilon + (1 - \epsilon)p_{ij}^n > p_{ij}^n$ and $(1 - \epsilon)p_{ij}^n < p_{ij}^n$. Therefore if $p_{ij}^{n+1} > p_{ij}^n$ we have $p_{ij}^{n+1} = \epsilon + (1 - \epsilon)p_{ij}^n$, which yields the conclusion

Lemma 3 says that, within process $P_{HS}(G, p, \epsilon)$, players evaluate the perceived payoff for different strategies and increment the probability to adopt one of those with highest payoff. It is worth noting that if $\epsilon < 1$ players do not shift to one of the most profitable strategies at once. But they increase their intention to use that one.

Lemma 4 *Let G be a game, p a prior distribution probability, and $\{p^n\} = P_{HS}(G, p, \epsilon)$. If $\{p^n\}$ converges, then we have $\lim_n p_{ij}^n \in \{0, 1\}$ and there exists $\bar{n} > 0$, such that for all $n \geq \bar{n}$ $\{p_{ij}^n\}$ are strictly monotonic sequences.*

Proof. Let us assume that $\lim_n p_{ij}^n$ there exists. We have $\lim_n |p_{ij}^{n+1} - p_{ij}^n| = 0$. Moreover

$$p_{ij}^{n+1} - p_{ij}^n = \begin{cases} -\epsilon p_{ij}^n \\ \epsilon(1 - p_{ij}^n) \end{cases}$$

and this cannot converge to 0 unless p_{ij}^n converges to 1 or to 0. We see also that if $\lim_n p_{ij}^n = 1$ there must exist an $\bar{n} > 0$ such that for all $n > \bar{n}$, $p_{ij}^{n+1} - p_{ij}^n = \epsilon(1 - p_{ij}^n)$ and $p_{ij}^{n+1} > p_{ij}^n$. One can

analogously prove that if $\lim_n p_{ij}^n = 0$ there must exist an $\bar{n} > 0$ such that for all $n > \bar{n}$ $p_{ij}^{n+1} < p_{ij}^n$

The two previous lemmas enable us to prove the following result

Proposition 7 *Let G be a game, p a prior distribution probability and $\{p^n\} = P_{HS}(G, p, \epsilon)$. If $\{p^n\}$ converges then it is asymptotically rational.*

Proof. Let us suppose that $U_i(\sigma_{ij} p_{-i}^n) \leq U_i(\sigma_{ik} p_{-i}^n)$ and $\lim_n p_{ij}^n > 0$. Then, from Lemma 4 $\lim_n p_{ij}^n = 1$ and there exists $\bar{n} > 0$ such that, for all $n > \bar{n}$ $p_{ij}^{n+1} > p_{ij}^n$. Then from lemma 3 we have $U_i(\sigma_{ij} p_{-i}^n) \geq U_i(\sigma_{ik} p_{-i}^n)$, and this yields $\lim_n U_i(\sigma_{ij} p_{-i}^n) = \lim_n U_i(\sigma_{ik} p_{-i}^n)$

Corollary 1 *Let G be a game, p a prior distribution probability and $\{p^n\} = P_{HS}(G, p, \epsilon)$. If $\{p^n\}$ converges then the limit point is a Nash equilibrium.*

Proposition 8 *Let G be a game, p a prior distribution probability and $\{p^n\} = P_{HS}(G, p, \epsilon)$. If $\{p^n\}$ converges and $\epsilon < 1$ then the limit point is a perfect equilibrium.*

Proof. We must show that a sequence $\{\epsilon^n\}$ exists, such that $\lim_n \epsilon^n = 0$ and p^n is an ϵ^n -perfect equilibrium point. From lemma 4 we know there exists $\bar{n} > 0$, such that for all $n > \bar{n}$, the sequences $\{p_{ij}^n\}$ are strictly monotonic. Let $\epsilon^n = \max_{i,j} \{p_{ij}^n : \lim_n p_{ij}^n = 0\}$. We have $\lim_n \epsilon^n = 0$ and if $U_i(\sigma_{ij} p_{-i}^n) < U_i(\sigma_{ik} p_{-i}^n)$ then $p_{ij}^{n+1} < p_{ij}^n$ and $\lim_n p_{ij}^n = 0$. Therefore we have $p_{ij}^n \leq \epsilon^n$

In conclusion, $P_{HS}(G, p, \epsilon)$ is an evolutionary process that can be computed in parallel by the variable/agent and when it converges the limit point is a perfect equilibrium. Unfortunately, as we have seen in a previous section, in general the set of perfect equilibrium points only contains the solution set of the CSP the game has originated. However, it is worth noting this limitation enable us to define, when the process converges, the “outcome” of a CSP even if that particular CSP does not have a solution. Moreover, in some cases (for example in resource management) one is often interested in a quasi solution, that is a solution that does not satisfy all the constraints, but that can be accepted by the players as the best thing they can do.

5.3 Experimental Results

In this section we shall illustrate a very simple example in order to give a practical feeling of what happens during an evolutionary process such as $P_{HS}(G, p, \epsilon)$. All the computations here reported have been performed using the object-oriented language COOL [Avesani *et al.*, 1990] which was expressly developed to face CSP.

We shall deal again with the previously mentioned 8 Queens problem. According to the definition of $G(C)$ (Section 3.2) this CSP can be translated into a non-cooperative game with eight players with identical strategy sets $S_i = \{s_{i0}, \dots, s_{i7}\} = \{0, \dots, 7\}$, $i = 0, \dots, 7$ and utility functions

$$U_i(s_{0j_0}, \dots, s_{7j_7}) = \sum_{j=0}^7 \chi_{r_{ij}}^{0,-1}(d_i, d_j),$$

where r_{ij} is the constraint between the i -th and j -th queen. We note that a function $\chi^{0,-1}$ has been used, that is a loss of utility is gained when a constraint is not satisfied.

We have generated 80 instances of $P_{HS}(G, p, \epsilon)$ evolutionary processes with $\epsilon = 0.1$. Every case is characterized by an initial prior distribution probability p . Half of those start with a uniform

prior distribution, that is we have $p_{ij} = 1/8$. In the other cases p_{ij} are taken as normalized random numbers in the interval $[0, 100]$.

When the prior distribution is uniform, 40% of the processes converge to a solution of the CSP, and the others come up with precisely one constraint not satisfied. In the second set of examples the percentage of processes ending with a solution increases a bit and we have 55% of the processes ending with a solution.

Some observations are in order. First of all we note that all of the processes converge, which is not a fact we have been able to prove in general. Second, it seems that a random initial distribution has a better median behaviour than uniform prior distribution. And this could be interpreted saying that an initial noise may be of some help in finding a solution. Finally we observe that the choice of $\chi^{0,-1}$ ($a = 0$, $b = -1$) is motivated by a partial evaluation of different choices of the parameters a and b . We think that a more extensive study could reveal interesting differentiations in the behaviour of processes which adopt different utility functions.

6 Conclusions

In this paper we have shown how to translate a CSP into a non-cooperative game. We have studied the relationships between the equilibrium theory in games and the solution concept in CSP. Evolutionary processes, first introduced in application of game theory to economics [Brown, 1951] and biology [Smith, 1982], can be usefully adopted to model the development of a mutual agreement on a particular instantiation of the search variables. Each variable is to be considered an elementary processing unit, that can increase or decrease the probability to assume a certain value, according to the computation of the expected utility of every possible choice. An evolutionary process simulates a step by step interaction and updating cycle. If the $P_{HS}(G, p, \epsilon)$ process converges then for each variable there exists a value that has probability 1 to be chosen. The set of those values represents the outcome of the CSP. The computation can be done in parallel by the variables and only bayesian reasoning is needed. The limit point of such a process can be viewed as the outcome of a course of mutual interaction between bayesian reasoners. Even if this limit point is not a solution of the CSP it can be considered as the best agreement the variables/agents have managed to get.

The major limitation of $P_{HS}(G, p, \epsilon)$ resides on the computation of the best reply routine, where a large number of configurations have to be generated. More precisely, in computing $BR(q)_i$, the best reply strategy of X_i , all the configurations of the variables connected to X_i have to be produced (see lemma 1). This procedure is practically feasible only when a variable is connected to a small fraction of the total number of variables.

For the more general case, we are conducting some experiments where the computation of the utility functions is approximated with a stochastic sampling. In this case we generate only a few configurations, using for each variable the probability distribution computed in the previous step. On a future paper we will mostly concentrate on the integration of bayesian reasoning with other forms of symbolic reasoning.

References

- [Avesani *et al.*, 1990] Paolo Avesani, Anna Perini, and Francesco Ricci. COOL: An object system with constraints. In Jean Bézivin, Bertrand Meyer, and Jean-Marc Nerson, editors, *Technology of Object-Oriented Languages and Systems*, pages 221–228, 1990.
- [Borning *et al.*, 1987] Alan Borning, Robert Duisberg, Bjorn Freeman-Benson, Axel Kramer, and Michael Woolf. Constraint hierarchies. In *Proceedings of OOPSLA '87*, pages 48–60, 1987.

- [Brown, 1951] G. W. Brown. Iterative solutions of games by fictitious play. In T. C. Koopmans, editor, *Activity analysis of production and allocation*. John Wiley, 1951.
- [de Kleer, 1989] Johan de Kleer. A comparison of the ATMS and CSP techniques. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 290–296, 1989.
- [Dechter *et al.*, 1989] Rina Dechter, Italy Meiri, and Judea Pearl. Temporal constraint networks. In *Proceedings of the first international conference on principles of knowledge representation and reasoning*, pages 83–93, 1989.
- [Fox, 1987] Mark S. Fox. *Constraint-Directed Search: A Case Study of Job-Shop Scheduling*. Morgan Kaufmann, 1987.
- [Gasser and Huhns, 1989] Les Gasser and Michael N. Huhns, editors. *Distributed artificial intelligence, volume II*, London, 1989. Pitman.
- [Genesereth *et al.*, 1986] Michael R. Genesereth, Matthew L. Ginsberg, , and Jeffrey S. Rosen-schein. Cooperation without communication. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 51–57, 1986.
- [Harsany and Selten, 1988] John C. Harsany and Reinhard Selten. *General theory of equilibrium selection in games*. M.I.T. press, 1988.
- [Huberman and Hogg, 1988] Bernardo A. Huberman and Tad Hogg. The behaviour of computational ecologies. In B. A. Huberman, editor, *The ecology of computation*, pages 77–115. North-Holland, 1988.
- [Mackworth, 1977] Alan K. Mackworth. Consistency in network of relations. *Artificial Intelligence*, 8:99–118, 1977.
- [Mackworth, 1987] Alan K. Mackworth. Constraint satisfaction. In Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 205–211. Wiley, 1987.
- [Montanari and Rossi, 1989] Ugo Montanari and Francesca Rossi. Exact solution in linear time of network of constraints using perfect relaxation. In Ronald Brachman, Hector Levesque, and Ray Reiter, editors, *Proceedings of the first international conference on principles of knowledge representation and reasoning*, pages 394–399, 1989.
- [Myerson, 1978] R. B. Myerson. Refinements of the Nash equilibrium concept. *Int. Journal of Game Theory*, 7:73–80, 1978.
- [Nadel, 1989] Bernard A. Nadel. Constraint satisfaction algorithms. *Computational Intelligence*, 5:188–224, 1989.
- [Nash, 1951] John Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- [Ricci, 1990] Francesco Ricci. Constraint networks. *Note di Software*, (47):3–14, 1990.
- [Samuelson, 1988] Larry Samuelson. Evolutionary foundations of solution concepts for finite, two-player, normal-form games. In Moshe Y. Vardi, editor, *Proceedings of the second conference on theoretical aspects of reasoning about knowledge*, pages 211–225. Morgan Kaufman, 1988.
- [Sathi and Fox, 1989] Arvind Sathi and Mark S. Fox. Constraint-directed negotiation of resource reallocations. In Les Gasser and Michael N. Huhns, editors, *Distributed artificial intelligence, volume II*, pages 163–193. Pitman, 1989.

- [Selten, 1975] Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.
- [Smith, 1982] John Maynard Smith. *Evolution and the theory of games*. Cambridge University Press, 1982.
- [Waltz, 1975] David L. Waltz. Understanding line drawings of scenes with shadows. In *The psychology of computer vision*, pages 19–91. McGraw-Hill, New York, 1975.