



ELSEVIER

Artificial Intelligence ●●● (●●●●) ●●●-●●●

**Artificial
Intelligence**

www.elsevier.com/locate/artint

Multi-agent oriented constraint satisfaction

Jiming Liu *, Han Jing, Y.Y. Tang

Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong

Received 31 May 2001; received in revised form 2 August 2001

Abstract

This paper presents a multi-agent oriented method for solving CSPs (Constraint Satisfaction Problems). In this method, distributed agents represent variables and a two-dimensional grid-like environment in which the agents inhabit corresponds to the domains of the variables. Thus, the positions of the agents in such an environment constitute the solution to a CSP. In order to reach a solution state, the agents will rely on predefined local reactive behaviors; namely, *better-move*, *least-move*, and *random-move*. While presenting the formalisms and algorithm, we will analyze the correctness and complexity of the algorithm, and demonstrate the proposed method with two benchmark CSPs, i.e., *n*-queen problems and coloring problems. In order to further determine the effectiveness of different reactive behaviors, we will examine the performance of this method in deriving solutions based on behavior prioritization and different selection probabilities. © 2002 Published by Elsevier Science B.V.

Keywords: Constraint satisfaction; Multi-agent; Reactive moving behaviors; Behavior prioritization; Behavior selection; Experimental validation

1. Introduction

1.1. CSPs

Many problems in Artificial Intelligence (AI) as well as in other areas of computer science and engineering can be translated into a certain type of *constraint satisfaction problem* (CSP) [20,30]. Some examples of such problems include: spatial and temporal planning, qualitative and symbolic reasoning, diagnostics, decision support, computational

* Corresponding author.

E-mail address: jiming@comp.hkbu.edu.hk (J. Liu).

1 linguistics, scheduling, resource allocation and planning, graph problems, hardware design 1
2 and verification, configuration, real-time systems, and robot planning. 2
3

4 **Definition 1.1.** A constraint satisfaction problem (CSP) consists of: 4

- 5
6 (1) A finite set of variables, $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$. 6
7 (2) A domain set, containing a finite and discrete domain for each variable: 7

$$8 \quad \mathbf{D} = \{D_1, D_2, \dots, D_n\}, \forall i \in [1, n], X_i \in D_i. 8$$

- 9
10 (3) A constraint set, $\mathbf{C} = \{C(R_1), C(R_2), \dots, C(R_m)\}$, where each R_i is an ordered 10
11 subset of the variables, and each constraint $C(R_i)$ is a set of tuples indicating the 11
12 mutually consistent values of the variables in R_i . 12

13
14 **Definition 1.2.** The solution, S , for a CSP is an assignment to all variables such that the 14
15 assignment satisfies all given constraints. Specifically, 15

- 16
17 (1) S is an ordered set, $S = \langle v_1, v_2, \dots, v_n \rangle, S \in D_1 \times D_2 \times \dots \times D_n$. 17
18 (2) $(\forall j \in [1, m]) (\exists S' \subseteq S) \wedge (S' \in C(R_j))$ is true. S' is also an ordered set. 18

19
20 In this paper, we will focus our discussion on *binary CSPs* where each constraint is 20
21 either unary or binary [20]. A binary CSP has the same definition as Definition 1.1, except 21
22 $R_i = D_{i1} \times D_{i2}$. It is possible to convert a CSP with n -ary constraints to an equivalent 22
23 binary CSP [20,35]. 23

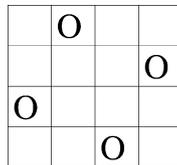
24 Let us now take a look at two typical CSP examples as follows: 24

25
26 **Example 1.1.** The n -queen problem is a classical CSP. It is generally regarded as 26
27 a benchmark for testing algorithms and has attracted a lot of attentions in the CSP 27
28 community [39]. This problem requires one to place n queens on an $n \times n$ chessboard, 28
29 so that no two queens are in the same row, the same column, or the same diagonal. There 29
30 exist solutions for the n -queen problems with n greater than or equal to 4 [2,39] (see Fig. 1). 30
31 The equivalent CSP can be stated as follows: 31

$$32 \quad \mathbf{X} = \{X_1, X_2, \dots, X_n\}. 32$$

$$33 \quad \mathbf{D} = \{D_1, D_2, \dots, D_n\}, \quad \forall i, D_i = [1, n]. 33$$

$$34 \quad \mathbf{C} = \{C(R_u) \mid \forall i, j \in [1, n], C(R_u) = \{\langle b, c \rangle \mid b \in D_i, c \in D_j, b \neq c, i - j \neq 34 \\ 35 \quad b - c, i - j \neq c - b\}\}. 35$$



O: queen

Fig. 1. A solution for a 4-queen problem.

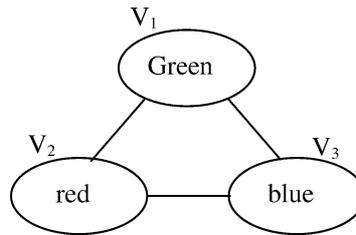


Fig. 2. An example solution for a coloring problem.

Example 1.2. The (vertex) coloring problem that is found in a variety of applications can readily be modeled as a CSP. In this problem, we need to color each vertex or node of a graph by a certain color (from a set of colors, suppose to be m colors), such that no two nodes incident to any edge have the same color. The equivalent CSP will represent each of the nodes in the graph into a variable. The domain of the variable corresponds to the given set of m colors. For each pair of nodes incident to an edge, there is a binary constraint between the corresponding variables that disallows identical assignments to these two variables. Here is an example: $\mathbf{X} = \{V_1, V_2, V_3\}$, $D_1 = \{green, red\}$, $D_2 = \{red, blue\}$, $D_3 = \{blue\}$, where constraints are: $V_1 \neq V_2$, $V_1 \neq V_3$, and $V_2 \neq V_3$, so $\mathbf{C} = \{\langle green, red \rangle, \langle green, blue \rangle, \langle red, blue \rangle\}$, $\{\langle green, blue \rangle, \langle red, blue \rangle\}$, $\{\langle red, blue \rangle\}$. Fig. 2 presents a possible solution to this problem.

1.2. Related work

General methods for solving CSPs include *generate-and-test* (GT) and *backtracking* (BT) methods [20]. GT generates each possible combination of the variables systematically and then checks to see whether it is a solution, i.e., whether it satisfies all the constraints. One limitation of this method is that it has to consider all instances of the Cartesian product of all the variable domains. In this respect, BT is more efficient than GT, as it assigns values to variables sequentially and then checks constraints for each variable assignment. If a partial assignment does not satisfy any of the constraints, it will backtrack to the most recently assigned variable and repeat the process again. Although this method eliminates a subspace from the Cartesian product of all the variable domains, its computational complexity for solving most nontrivial problems remains to be exponential.

Many studies have been conducted to investigate various ways of improving the above-mentioned BT method. In order to avoid *thrashing* [16,20] in BT, *consistency* techniques (Arc Consistency and k-Consistency) have been developed by Mackworth and other researchers [6,18,20,26,29], which are able to remove inconsistent values from the domains of the variables. In order to avoid both thrashing and *redundant-work* [20] in BT, a *dependency-directed backtracking* scheme and its improvements have been proposed [3, 20,34,40]. Other ways of increasing the efficiency of BT include the use of *search order* for variables, values, and consistency check [33]. Nevertheless, even with such improvements, BT is still unable to solve nontrivial large-scale CSPs in a reasonable runtime.

1 For the GT method, there have been some research efforts on making the solution 1
2 generator smarter. The representatives of such efforts are *stochastic* and *heuristic* 2
3 algorithms. Along this direction, one of the most popular ideas is to perform *local* 3
4 *search* [17,20]. For large-scale n -queen CSPs, it gives better results than a complete, or 4
5 even incomplete, systematic BT method. There are three key elements in local search [1], 5
6 they are: 6

- 7
- 8 (1) Configuration: one possible assignment of all variables, not required to be a solution. 8
- 9 (2) Evaluation value: the number of dissatisfied constraints. 9
- 10 (3) Neighbor: the configuration obtained by changing one variable's assignment of the 10
11 current configuration. 11

12
13 Local search generates an initial configuration and then incrementally uses "repair" 13
14 or "hill climbing" to modify the inconsistent configuration to move to a neighborhood 14
15 configuration that has the best or better evaluation value among its neighbors, until a 15
16 solution is found. In order to avoid falling into the *local optima*, it sometimes performs 16
17 *random-walk* and *tabu search* [13]. As related to the idea of local search, other heuristics 17
18 have also been developed, such as *hill-climbing* [1], *min-conflicts* [28], *MCRW* (Min- 18
19 Conflicts-Random-Walk) [42], and *GSAT* (GSAT is a randomized local search procedure 19
20 for solving propositional satisfiability problems) [37,38]. 20

21 22 23 1.2.1. *Min-conflicts heuristics* 23

24 Repair-based heuristics were originally used in AI problem-solving systems to debug 24
25 and modify initial solutions. Minton et al. [28] extended this approach to solving large- 25
26 scale constraint satisfaction problems, and proposed a value-ordering heuristic, called *min-* 26
27 *conflicts* heuristic. The *min-conflicts* heuristic attempts to select a new value that minimizes 27
28 the number of outstanding constraint violations after each step. 28

29 They argued that the effectiveness of the *min-conflicts* heuristic is largely due to the 29
30 repair of a complete but inconsistent assignment that is more informative in guiding search 30
31 than an incrementally constructed partial assignment as in the traditional backtracking 31
32 methods [28]. They also noted that the performance of this sequentially executed heuristic 32
33 is remarkably comparable to that of a parallelly implemented Guarded Discrete Stochastic 33
34 (GDS) network for solving constraint satisfaction problems (e.g., the Hubble Space 34
35 Telescope scheduling problem). The two implementations employed the same heuristic (in 35
36 fact, as stated in [28], the *min-conflicts* approach was intended to replicate the behavior of 36
37 the GDS network): the network reassigns a value for a variable by choosing the value that 37
38 violates the fewest constraints (i.e., flipping the neuron whose output is most inconsistent 38
39 with its current input). 39

40 Our multi-agent approach utilizes the idea of inconsistency reduction over a complete 40
41 initial assignment. However, our approach differs from the *min-conflicts* approach in a 41
42 number of ways. For instance, our approach explores heuristics in addition to violation 42
43 *minimization*, and relays on the combination of prioritized heuristics in order to improve 43
44 computational efficiency. A more detailed discussion on the distinctions between the two 44
45 approaches is provided in Section 5. 45

1 Other methods for solving CSPs have been based on Neural Networks [27] and Genetic 1
2 Algorithms [19]. 2

3 The above-mentioned methods and techniques have their advantages and drawbacks, 3
4 and no single algorithm has been found to be suited to solving all CSPs. For small-size 4
5 problems, we may use BT to readily find a solution, whereas for large-scale problems 5
6 we may use local search. The efficiency of local search for solving n -queen problems is 6
7 reported to be very efficient, among other algorithms [39]. However, we cannot prove that 7
8 it can find solutions for every case every time as it is stochastic in nature, while on the 8
9 other hand BT's performance is more stable and complete. Furthermore, local search is 9
10 not suitable for problems other than n -queen. It requires that the problems have a clear 10
11 neighborhood structure. 11

12 1.3. Multi-agent systems 12

13 Agent-based computation has been studied for some years in the field of artificial 13
14 intelligence and has been widely used in other branches of computer science. Multi-agent 14
15 systems are computational systems in which several agents interact or work together in 15
16 order to achieve goals. Agents in such systems may be homogeneous or heterogeneous, and 16
17 may have common goals or distinct goals [22]. Previous work on multi-agent systems has 17
18 generally focused on areas such as simulations of social and biological systems, problem 18
19 solving, communication, collective robotics, and electronic commerce on the Internet. 19
20 20
21 21
22 22

23 1.3.1. Distributed constraint satisfaction 23

24 A distributed constraint satisfaction problem (distributed CSP) is a constraint satisfac- 24
25 tion problem in which variables and constraints are semantically partitioned (or *distrib-* 25
26 *uted*) into *sub-problems*, each of which is to be solved by an agent. When multiple agents 26
27 are involved in solving a distributed CSP, the agents have to comply with certain constraints 27
28 among them. Thus, finding a solution to a distributed CSP requires that all agents find 28
29 the values for their variables that satisfy not only their own constraints but also interagent 29
30 constraints. Examples of distributed CSP research efforts include distributed scheduling, 30
31 planning, and reasoning [5,7,32,36]. 31

32 Yokoo et al. [43–46] have made significant contributions in the area of distributed 32
33 CSP. They developed an algorithm called *asynchronous backtracking* that guarantees the 33
34 completeness, and then later extended this algorithm into a more efficient *asynchronous* 34
35 *weak-commitment search* algorithm, by introducing dynamic ordering among agents. 35
36 Furthermore, they also proposed a *multi-agent real-time-A* algorithm with selection* to 36
37 solve an n -puzzle problem [46]. In those algorithms, the agents are individual solvers for 37
38 obtaining partial solutions. 38
39 39

40 1.3.2. Swarm-like systems 40

41 In addition to the above-mentioned distributed constraint satisfaction approaches, it is 41
42 worth mentioning a special instance of multi-agent systems for applications in computation 42
43 and simulation; namely, *swarm* [41]. 43

44 *Swarm* is a formulation for simulating distributed multi-agent systems, which involves 44
45 three key concepts: living environment, agents with reactive rules, and a schedule serving 45

1 as a timetable to update the changes and dispatch agents' actions. Based on this idea, Liu 1
2 et al. [24,25] developed an evolutionary autonomous agent system to adaptively extract 2
3 image features and segments. Recently, Liu and Han [23] proposed an energy-based 3
4 artificial-life model for solving n -queen problem. 4

5 6 1.4. The proposed approach 6 7

8 As inspired by the previous models of swarm, in this paper we will present a new 8
9 approach called ERA (i.e., Environment, Reactive rules, and Agents) to solving CSPs. The 9
10 CSPs considered here will not be limited to distributed CSPs. This approach is intended to 10
11 provide an alternative, multi-agent formulation that can be used to handle general CSPs and 11
12 to find approximate solutions without too much computational cost. The key idea behind 12
13 ERA lies in a distributed multi-agent system, having the same architecture as swarm, i.e., 13
14 *an environment, agents with moving behaviors, reactive rules, and a schedule*. This system 14
15 self-organizes itself, when each individual agent follows its behavioral rules, and gradually 15
16 evolves toward a global solution state. 16

17 From the point of view of solving CSPs, the proposed approach may be regarded as an 17
18 extended GT approach, somewhat like local search. However, the main difference between 18
19 ERA and local search is that the evaluation value of ERA is not the number of dissatisfied 19
20 constraints for the whole assignment as in local search, but the number of dissatisfied 20
21 constraints for the value of each variable—these numbers constitute an *environment* in the 21
22 ERA system. 22

23 If there exists a consistent solution, the ERA system will eventually find it. On the other 23
24 hand, if there is no complete solution, the ERA system can still generate an approximate 24
25 solution. As to be presented in this paper, our approach can solve both n -queen problems 25
26 and coloring problems. Furthermore, our experiments will show that ERA is efficient 26
27 in finding exact as well as approximate solutions to CSPs in few time steps. Generally 27
28 speaking, it is more efficient than the BT algorithms and more readily to solve different 28
29 CSPs than the local search algorithm. 29
30

31 1.5. Organization of the paper 31 32

33 The remainder of this paper is organized as follows: Section 2 describes the basic 33
34 ideas behind this distributed multi-agent oriented method. Section 3 discusses how to use 34
35 this method to find an approximate solution. Section 4 describes several experiments and 35
36 observations. Section 5 discusses the features of the proposed ERA approach and compare 36
37 them with the existing major approaches in the field. Finally, Section 6 concludes the paper 37
38 by highlighting the contribution of this work and some future extensions. 38
39

40 41 42 2. The multi-agent model 42 43

44 In this section, we will describe the basic formulation and algorithm for our proposed 44
45 multi-agent model. Specifically, we will provide the definitions as well as formalisms for 45

1 agent environments, major policies for agent-environment interaction, reactive behaviors, 1
2 and a basic ERA algorithm. 2

3 2.1. ERA fundamentals 3

4 *Problem solving* is an area that many multi-agent-based applications are concerned with. 4
5 It includes the following subareas: *distributed solutions to problems*, *solving distributed* 5
6 *problems*, and *distributed techniques for problem solving* [10,22]. In this paper, we will 6
7 introduce an application of distributed techniques for solving CSPs. In our case, the domain 7
8 of a CSP is represented into a multi-agent environment. Thus, the problem of finding a 8
9 solution to the CSP is reduced to that of local behavior-governed moves within such an 9
10 environment. 10

11 Specifically, the notions of agent and multi-agent system can be defined as follows: 11
12

13 **Definition 2.1.** An agent, a , is a virtual entity that essentially has the following properties: 13
14

- 15 (1) Be able to live and act in the environment. 15
- 16 (2) Be able to sense its local environment. 16
- 17 (3) Be driven by certain objectives. 17
- 18 (4) Have some reactive behaviors. 18

19 **Definition 2.2.** A multi-agent system is a system that contains the following elements: 19
20

- 21 (1) An environment, E , a space in which the agents live. 21
- 22 (2) A set of reactive rules, R , governing the interaction between the agents and their 22
23 environment. They are the laws of the agent universe. 23
- 24 (3) A set of agents, $A = \{a_1, a_2, \dots, a_n\}$. 24

25 The goal of this work is to examine how exact or approximate solutions to CSPs can be 25
26 self-organized by a multi-agent system, consisting of $\{E, R, A\}$. 26
27

28 2.1.1. Overview of the ERA multi-agent formulation 28

29 The ERA method is meant to be a framework for interacting agents to achieve a global 29
30 solution state. In ERA, the environment records the number of constraint violations of the 30
31 current state for each value in the domains of all variables. Each agent represents a variable 31
32 and the position of the agent corresponds to the value of the respective variable. The agent 32
33 can move locally within a row and has its own reactive moving behaviors. Its objective is 33
34 to move to a position whose constraint violation number is *zero*, we call it *zero-position* 34
35 (for detail see Definition 2.3(2)). An exact solution state in ERA is reached when every 35
36 agent (variable) finds its *zero-position*. The reactive rules correspond to the schedules for 36
37 dispatching agents and updating the environment. 37

38 In this paper, we will first present the basic formulation and algorithm for the ERA 38
39 method, and then focus on the effectiveness of some extended ERA techniques that utilize 39
40 combined reactive behaviors as well as different selection probabilities. 40
41
42
43
44
45

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

X_1	1	2	3	☺	5	6
X_2	1	☺	3	4		
X_3	☺	2	3	4	5	

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

Fig. 3. An illustration of agent model for Example 2.1.

In the following paragraph, we will use an example to illustrate how a CSP can be translated into an ERA multi-agent system.

Example 2.1. A CSP is given as follows:

$$\mathbf{X} = \{X_1, X_2, X_3\}, \quad n = 3.$$

$$\mathbf{D} = \{D_1, D_2, D_3\}, \quad D_1 = \{1, 2, 3, 4, 5, 6\}, \quad D_2 = \{1, 2, 3, 4\},$$

$$D_3 = \{1, 2, 3, 4, 5\}.$$

$$\mathbf{C} = \{X_1 \neq X_2, X_1 > X_3\}.$$

Example 2.1 can be modeled as a multi-agent system as follows: The lattices represent an environment, where each row corresponds to a variable’s domain and the length of each row is equal to the domain size. In each row, there exists only one agent. In this case, the horizontal coordinate of the agent represents the corresponding variable’s value. As in Fig. 3, there are three agents all residing at *zero-positions*. The numbers that these agents occupy correspond to the values within the domains of the three variables. Fig. 3 shows a solution state of $S = \langle 4, 2, 1 \rangle$.

2.1.2. Environment

An environment, E , has n rows corresponding to the number of variables. For all $i \in [1, n]$, row_i has $|D_i|$ columns. It records two kinds of values: the domain value and the violation value.

Definition 2.3. The data structure of E can be defined as follows:

(1) Size

- n rows $\Leftrightarrow n$ variables. $E = \langle row_1, row_2, \dots, row_n \rangle$.
- $\forall i \in [1, n]$,
 $row_i \Leftrightarrow$ domain of $X_i \Leftrightarrow D_i$, so row_i has $|D_i|$ columns.
 $row_i = \langle lattice_{1i}, lattice_{2i}, \dots, lattice_{|D_i|i} \rangle$.
- E is an array of size $\sum |D_k|$. $e(i, j)$ refers to the position of $lattice_{ij}$.

(2) Values

- Domain value: $e(i, j).value$ records the i th value of domain D_j .
- Before we introduce $e(i, j).violation$, let us first define the notion of ‘attack’ between $position_1$ and $position_2$.
 We use (x_1, y_1) to represent $position_1$ and (x_2, y_2) to represent $position_2$. So,

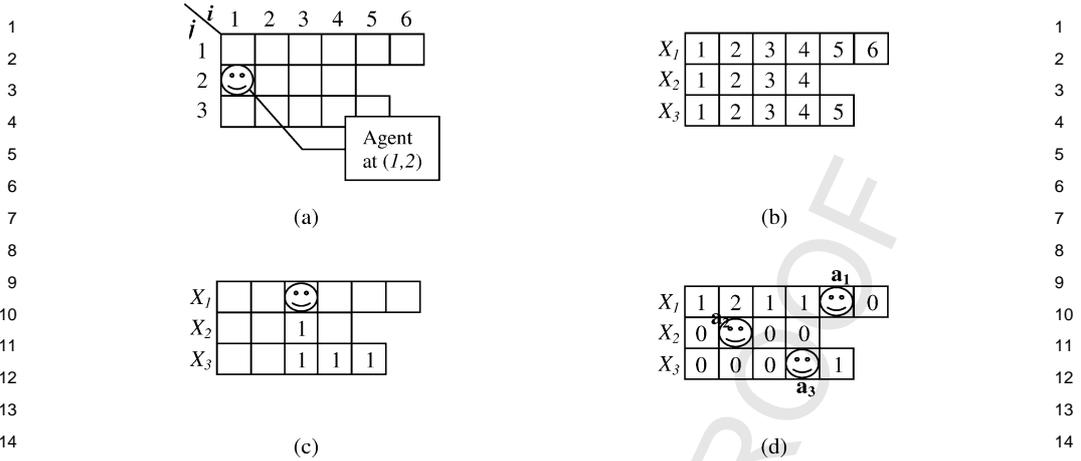


Fig. 4. An illustration of the agent environment. (a) the position of an agent, (b) the representation of domain values, and (c)–(d) violation numbers marked in the environment.

$$\begin{aligned}
 & \text{Attack}((x_1, y_1), (x_2, y_2)) \\
 &= \begin{cases} \text{true,} & \text{if there is constraint } C(R_t) \text{ between } X_{y_1} \text{ and } X_{y_2} \\ & \text{and } \langle e(x_1, y_1).value, e(x_2, y_2).value \rangle \notin C(R_t), \\ \text{false,} & \text{otherwise.} \end{cases} \quad (1)
 \end{aligned}$$

- Violation number: $e(i, j).violation$ records in the current state how many agents whose positions *attack* position (i, j) , i.e., $e(i, j).violation = m$ means there are m agents whose assignments dissatisfy the assignment of $X_j = e(i, j).value$. The $e(i, j).violation$ values are dynamically modified since the agents keep on moving and their corresponding state is changing. The violation numbers will be updated by applying an *updating-rule*, which will be described in Section 2.1.5.
- *zero-position*: position (i, j) , in which $e(i, j).violation = 0$. That means all other agents agree on $X_j = e(i, j).value$, according to constraints related to X_j .

Fig. 4(a) presents the position of an agent at $(1, 2)$. Fig. 4(b) shows the domain value of each lattice. row_1 contains values in domain $D_1 = \{1, 2, 3, 4, 5, 6\}$, row_2 represents $D_2 = \{1, 2, 3, 4\}$, and row_3 represents $D_3 = \{1, 2, 3, 4, 5\}$. Fig. 4(c) shows that if agent a_1 stays at $(3, 1)$, meaning $X_1 = 3$, according to the constraints of $X_1 \neq X_2$ and $X_1 > X_3$, it will violate $X_2 = 3$, $X_3 = 1$, $X_3 = 2$, and $X_3 = 3$. Therefore, it will contribute 1 to the violation number at position $(3, 2)$, $(3, 3)$, $(4, 3)$, and $(5, 3)$. Fig. 4(d) presents a snapshot for the state of the system with the violation numbers. Since all agents are at *zero-positions*, the state corresponds to an exact solution.

2.1.3. Agents

All agents inhabit in an environment, in which their positions indicate values of certain variables. During the operation of the system, the agents will keep on moving, based on certain reactive moving behaviors. At each time step, the positions of the agents provide a

1 consistent or inconsistent assignment for all variables. The agents are trying to find better 1
 2 positions that can lead them to a solution state. 2

3 Here is a summary of some major polices for agent-environment interaction in the ERA 3
 4 model: 4

5
 6 (1) $\forall i \in [1, n]$, a_i represents X_i . As shown in Fig. 4(d), three agents, a_1 , a_2 , and a_3 6
 7 represent X_1 , X_2 , and X_3 , respectively. 7

8 (2) Agent a_i moves locally in row_i . It can only move to its right or left, but not up 8
 9 and down. $a_i.x$ represents its x -coordinate. So the position of a_i can be denoted as 9
 10 $(a_i.x, i)$. 10
 11 11

12 **Example 2.2.** In Fig. 4(d), a_2 lives in row_2 , and it can move freely to position (1, 2), (2, 12
 13 2), (3, 2), or (4, 2) in one step, but not to other positions. 13
 14 14

15 In this paper, we use function ψ to define an agent's move. 15
 16 16

17
 18 **Definition 2.4.** $\psi : [1, n] \times [1, |D_i|] \rightarrow [1, |D_i|]$. $\psi(x, y)$ gives the x -coordination of the 18
 19 new position of agent a_i , after it moves from position (x, y) . So the new position can be 19
 20 represented as $(\psi(x, y), y)$. 20
 21 21

22
 23 (3) In any state of the system, the positions of all agents form an assignment for all 23
 24 variables. $\forall j \in [1, n]$, $X_j = e(a_j.x, j).value$. It may not be a consistent assignment, 24
 25 i.e., not an exact solution. 25

26 If an assignment satisfies all the constraints, i.e., $\forall j \in [1, n]$, $e(a_j.x, j).violation = 0$, 26
 27 it is an exact solution, $S = \langle e(a_1.x, 1).value, e(a_2.x, 2).value, \dots, e(a_n.x, n).value \rangle$. 27

28 (4) Agent a_i is able to 'perceive' the violation number for each lattice in row_i . Here, we 28
 29 define a function $\varphi(i)$ for returning a position (x -coordination) with the minimum 29
 30 violation number in row_i . 30
 31 31

32 **Definition 2.5.** A *minimum-position* is the position of (x, j) such that $j \in [1, n] \wedge (\forall i \in$ 32
 33 $[1, |D_j|]), e(x, j).violation \leq e(i, j).violation$. 33
 34 34

35 **Definition 2.6.** The function for finding the *first minimum-position* for agent a_i in row_i is 35
 36 defined as follows: 36

37 $\varphi : [1, n] \rightarrow [1, \max(|D_i|)], \varphi(i) = x \mid (x, i)$ is a *minimum-position* $\wedge (\forall j \in [1, x])(j, i)$ 37
 38 is not a *minimum-position*. 38
 39 39

40
 41 (5) In order to achieve a goal state, each agent uses its local reactive behaviors. Agents 41
 42 attempt to move toward *zero-positions* at each time step. But, in most cases, they 42
 43 cannot, or only some lucky agents can, find *zero-positions*, simply because some rows 43
 44 do not contain such positions at a certain time step. In such cases, the agents will have 44
 45 to perform other behaviors. 45

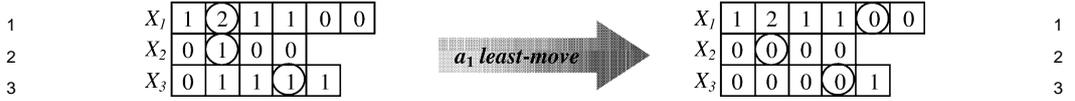


Fig. 5. The violation numbers are updated, when agent a_1 moves to a new position by executing a *least-move* behavior.

2.1.4. Local reactive behaviors

In order to reach a solution state, the agents will select and execute some predefined local reactive behaviors, namely, *better-move*, *least-move*, and *random-move*. Later in Section 4, we will investigate the effectiveness of these reactive behaviors by examining the performance of the ERA system with behavior prioritization and/or different selection probabilities.

2.1.4.1. *least-move*. An agent moves to a *minimum-position* with a probability of *least-p*. If there exists more than one *minimum-position*, we let the agent choose the first one on the left of the row. The *least-move* behavior can be expressed as follows:

$$\psi_{-l}(x, y) = \varphi(y). \quad (2)$$

Note that in this function, the result will not be affected by the current x , and the number of computational operations to find the position for each i is $|D_i|$.

Example 2.3. Fig. 5 shows that when agent a_1 performs a *least-move*, it will first compute $\psi_{-l}(2, 1) = \varphi(1) = 5$, and thereafter move to (5, 1).

2.1.4.2. *better-move*. An agent moves to a position that has a smaller violation number than its current position with a probability of *better-p*. It will randomly select a position and then compare its violation number to decide whether or not it should move to this position. We use function *Random(k)* to get a random number between 1 and k . This behavior can be defined using function ψ_{-b} :

$$\psi_{-b} = \begin{cases} x, & \text{when } e(\text{Random}(|D_y|), y).violation \geq e(x, y).violation, \\ \text{Random}(|D_y|), & \text{when } e(\text{Random}(|D_y|), y).violation < e(x, y).violation. \end{cases} \quad (3)$$

Although it may not be the best choice for the agent, the computational cost required for this behavior is much less than that of *least-move*. Only two operations are involved for deciding this move, i.e., producing a random number and performing a comparison. This behavior can readily find a position to move to especially when the agent is currently at a larger violation position.

As will be shown in Section 4, the *better-move* behavior plays an important role in bringing down the number of global constraint violations in a few time steps.

Example 2.4. Fig. 6 shows that when agent a_1 performs a *better-move*, it will compute $\psi_{-b}(2, 1)$. Suppose that $\text{Random}(6) = 3(|D_1| = 6)$. Thus, $\psi_{-b}(2, 1) = 3$, since $e(2, 1).violation > e(3, 1).violation$. The new assignment will become (3, 2, 4). Although

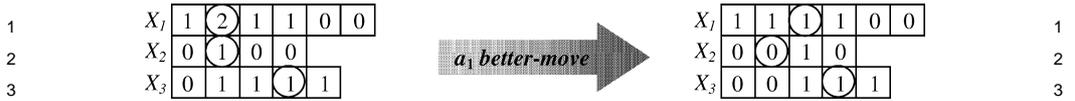


Fig. 6. The violation numbers are updated, when agent a_1 moves to a new position by executing a *better-move* behavior.



Fig. 7. The violation numbers are updated, when agent a_1 moves to a new position by executing a *random-move* behavior.

this assignment is not an exact solution, it is a better approximate solution than the assignment of $\langle 2, 2, 4 \rangle$ as in Fig. 5, because the new state has only one constraint, $X_1 > X_3$, dissatisfied.

2.1.4.3. random-move. An agent moves randomly with a probability of *random-p*. *random-p* will be relatively smaller than the probabilities for selecting *better-move* and *least-move* behaviors. It is somewhat like a *random-walk* in local search. For the same reason as in local search, *random-move* is necessary because without randomized moves the system will get stuck in *local-optima*, that is, all the agents are at *minimum-positions*, but not all of them at *zero-positions*. In the state of *local-optima*, no agent will move to a new position if using the behaviors of *better-move* and *least-move* alone. Thus, the agents will lose their chance of finding a solution if without any techniques to avoid getting stuck in *local-optima*.

random-move can be defined using function ψ_{-r} :

$$\psi_{-r}(x, y) = \text{Random}(|D_y|). \tag{4}$$

Example 2.5. Fig. 7 shows that when agent a_1 performs a *random-move*, it will randomly produce a number. If $\text{Random}(6) = 1$, it will move to $(1, 1)$. If $\text{Random}(6) = 3$, it will move to $(3, 1)$.

2.1.5. System schedule

The multi-agent system proposed in this paper is concurrent and discrete in nature, with respect to its space, time, and state space. In the present *simulated* implementation, the system will use a *discrete clock* to synchronize its operations, as shown in Fig. 8. It works as follows:

- *time step* = 0: The system is initialized. We place n agents onto the environment, a_1 in row_1 , a_2 in row_2 , ..., a_n in row_n . The simplest way to place the agents is to randomly select positions. That is, for a_i , we set a position of $(\text{Random}(|D_i|), i)$.
- *time step* \leftarrow *time step* + 1: For each time step, which means one unit increment of the system clock, all agents will have a chance to decide their moves, that is, *whether to move or not and where to move, and then move synchronously*.

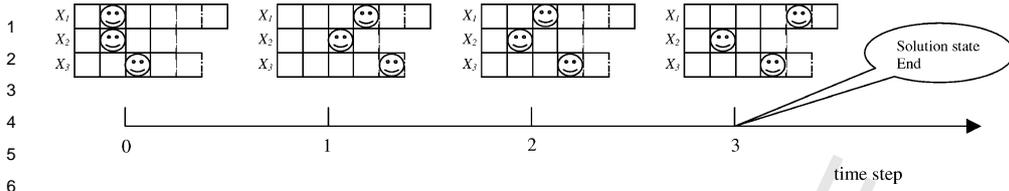


Fig. 8. Distributed agent-environment interaction at different time steps.

It should be pointed out that in the simulation, the multi-agent system dispatches the agents one by one. The order of dispatching is based on a random or a predefined sequence.

After the move of an agent from (x_1, y) to (x_2, y) , the violation number of the environment will be updated according to the following two *update-rules*:

(1) *update-rule 1*: Remove from (x_1, y)

$$(\forall x' \in [1, |D_y|]) (\forall y' \in [1, n]) (\text{Attack}((x_1, y), (x', y'))),$$

$$\text{execute } e(x', y').\text{violation} \leftarrow e(x', y').\text{violation} - 1.$$

(2) *update-rule 2*: Add to (x_2, y)

$$(\forall x' \in [1, |D_y|]) (\forall y' \in [1, n]) (\text{Attack}((x_2, y), (x', y'))),$$

$$\text{execute } e(x', y').\text{violation} \leftarrow e(x', y').\text{violation} + 1.$$

- *End*: After the moves of agents at each time step, the system will check whether all agents are at *zero-positions* and whether its clock exceeds a time threshold (i.e., time allowed). If one of these conditions is `true`, the system will stop its operations and return either an exact or an approximate solution.

Another way to terminate the operations is when q agents are staying at *zero-positions*.

2.2. The basic ERA algorithm

Now let us consider the basic ERA algorithm, from which a number of ERA properties and extended methods will be derived in the following sections.

Fig. 9 presents a function for initializing individual agents and then adding them to the environment randomly. This function also initializes the probabilities for *better-move*, *least-move*, and *random-move*. Fig. 10 provides the *RemoveFrom* function that updates environment numbers, according to *update-rule 1* when removing an agent from the environment. Fig. 11 shows the function of *AddTo* that updates violation numbers, according to *update-rule 2* when adding an agent to the environment. Function *SelectBehavior* in Fig. 12 selects a reactive behavior, according to the probabilities for various behaviors.

The complete listing of the basic ERA algorithm is given in Fig. 13.

1	Input: a_i	1
2	Output: the probabilities of a_i .	2
3	1. $a_i.random-p = p_i$;	3
4	2. $a_i.least-p = p_2$;	4
5	3. $a_i.better-p = p_3$;	5
6	4. $a_i.x = Random(I_d)$;	6
7	5. $AddTo(a_i.x, i)$;	7

Fig. 9. Function *Initialize*.

9	Input: position (x, y)	9
10	Output: updating violation numbers according to <i>update-rule 1</i> while removing a from (x, y) .	10
11	1. For all position $(i, j) \in environment$ do	11
12	2. If $Attack((i, j), (x, y))$ then $e(i, j).violation \leftarrow e(i, j).violation - 1$;	12
13	3. End for	13

Fig. 10. Function *RemoveFrom*.

16	Input: position (x, y)	16
17	Output: updating violation numbers according to <i>update-rule 2</i> while adding a to (x, y) .	17
18	1. $a_i.x = x$;	18
19	2. For all position $(i, j) \in environment$ do	19
20	3. If $Attack((i, j), (x, y))$ then $e(i, j).violation \leftarrow e(i, j).violation + 1$;	20
21	4. End for	21

Fig. 11. Function *AddTo*.

24	Input: i	24
25	Output: ψ	25
26	1. $p = Random(a_i.random_p + a_i.least_p + a_i.beter_p)$;	26
27	2. if $(p \leq a_i.random-p)$ then return ψ_i ;	27
28	3. if $(p \leq a_i.random-p + a_i.least-p)$ then return ψ_i ;	28
29	4. return ψ_b ;	29

Fig. 12. Function *SelectBehavior*.

2.3. Properties of the basic ERA algorithm

2.3.1. Termination

After each move, the ERA system will check whether all agents stay at *zero-positions*. Generally speaking, the termination condition for an exact solution can be stated as follows:

condition-1 : $(\forall a_i \in A) e(a_i.x, i).violation = 0$.

For an approximate solution, we can employ certain termination conditions as mentioned in Section 2.1.5, such as a threshold of time step. In this case, the algorithm will terminate if the clock exceeds *t-max*. Thus, the termination condition for an approximate solution can be stated as follows:

condition-2 : time step *t-max*.

```

1  Input:  $n$  variables, domains of variables, and constraints.
2  Output: an (approximate) solution.
3
4  Section-1. Initialization of the system:
5  1. time step = 0;
6  2. For all position  $(i, j) \in e$  do
7  3.    $e(i, j).value$  = the corresponding 1th value of domain  $D_j$ ;
8  4.    $e(i, j).violation=0$ ;
9  5. End for
10 6. For all  $a_i \in A$  do
11 7.   Initialize( $a_i$ );
12 8. End for
13 Section-2. Running of the system:
14 9. While (true) do
15 10.  For all  $a_i \in A$  do
16 11.    $\psi = \text{SelectBehavior}(\text{Random-Move}, \text{Least-Move}, \text{Better-Move})$ ;
17 12.   New position  $(x', i) = (\psi(a_i.x, i), i)$ ;
18 13.   If current-position  $(a_i.x, i) = (x', i)$  then stay
19 14.   Else
20 15.     RemoveFrom( $a_i.x, i$ );
21 16.     AddTo( $a_i.x, i$ );
22 17.   If current-state can satisfy us GoTo 21;
23 18.   End for
24 19.   time step ++;
25 20. End while
26 Section-3. Output solution:
27 21. For all  $a_i \in A$  do
28 22.    $X_i = e(a_i.x, i).value$ ;
29 23. End for

```

Fig. 13. The basic ERA algorithm for solving CSPs. It should be noted that the parallel operations of distributed agents are here simulated by means of sequentially dispatching agents and allowing them to sense the present state of their environment, containing violation numbers, and then decide where to move. In so doing, the movements of the agents will not interfere with each other. That is, their decisions will be independent. This is because in our implementation the sequentially updated violation numbers as a result of each agent's movement will be copied as the next state of the environment only after all the agents have been given the chance to move.

2.3.2. Correctness

We now give the correctness theorems for the basic ERA algorithm. Detailed proofs for the theorems can be found in Appendix A.

Note that when the system terminates at *condition-1*, all agents are at *zero-positions*.

Theorem 2.1. *If (x, y) is a zero-position, i.e., $e(x, y).violation = 0$, the following assertion is true: $(\forall a_i \in A, i \neq y) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow (e(a_i.x, i).value, e(x, y).value) \in C(R_t)$.*

Theorem 2.2. *The assignment of $S = \langle X_1, X_2, \dots, X_n \rangle$, $X_i = e(a_i.x, i).value$, is an exact solution when the system terminates at condition-1.*

2.3.3. Complexity

Now let us discuss the complexity of the basic ERA algorithm.

Theorem 2.3. *The space complexity of the basic ERA algorithm is $O(\sum |D_i|)$.*

1 **Proof.** The main contribution to the space complexity is from the storage for environment 1
 2 e . e has n rows, each row_i has $|D_i|$ lattices. So the total number of lattices in e is $\sum |D_i|$. 2
 3 For each lattice, it records 2 values, i.e., the domain value and the violation value. It 3
 4 requires $2 \sum |D_i|$ units to record environment e . Another contribution to the space is from 4
 5 the storage for n agents. For each agent, it records the current x -coordinate and three 5
 6 probabilities for *better-move*, *least-move*, and *random-move*, respectively. So it requires 6
 7 $4n$ space units in total for n agents. If all the agents have the same probabilities for *better-* 7
 8 *move*, *least-move*, and *random-move*, it needs only $n + 3$ units in total for n agents. In 8
 9 conclusion, the space complexity is $O(\sum |D_i|)$. \square 9

10
 11 **Theorem 2.4.** *The time complexity of the initialization is $O(\sum |D_i|)$.* 11

12
 13 **Proof.** For lines 2–5 in Fig. 13, there are in total $\sum |D_i|$ lattices to be initialized with the 13
 14 domain values and the violation numbers. So it needs $2 \sum |D_i|$ operators. And for lines 6– 14
 15 8, there are n agents to be initialized. For each agent, there are 3 operators for initializing 15
 16 the probabilities of three behaviors. There are in total $3n$ operations. So the complexity of 16
 17 initialization is $O(\sum |D_i|)$. \square 17

18
 19 **Theorem 2.5.** *The time complexity of each time step in the basic ERA algorithm is bounded 19
 20 by $O(n \sum |D_i|)$ in the worst case.* 20

21
 22 **Proof.** The main contribution to the time complexity is from the agent’s move, including 22
 23 the modification of violation numbers in the functions of *RemoveFrom* and *AddTo*, and 23
 24 the checking of the solution state. The number of operations in *RemoveFrom* and *AddTo* 24
 25 is $\sum |D_i|$ in the worst case that there exists a constraint between every two agents. And 25
 26 for the checking of a solution state, it needs n checks in the worst case. Now for each 26
 27 agent’s move, the total operation number is $\sum |D_i| + n$. So the complexity is bounded by 27
 28 $O(\sum |D_i| + n)$ in the worst case. In conclusion, for n agents, in the worst case, the time 28
 29 complexity is bounded by $O(n \times (\sum |D_i| + n)) = O(n \sum |D_i| + n^2) = O(n \sum |D_i|)$. \square 29
 30

31
 32
 33 **3. Approximate solution** 33

34
 35 One of the major motivations for developing the ERA multi-agent method is to be able 35
 36 to find an *anytime* solution, although it may be approximate, within the time allowed. In 36
 37 this section, we will discuss some properties of the basic ERA method, with respect to the 37
 38 goal of deriving approximate solutions: 38

- 39
 40 (1) Each state represents an approximate solution. 40
 41 In the BT method, variables are assigned with values sequentially. Unless the first k 41
 42 variables’ assignments satisfy constraints, the $(k + 1)$ th variable’s assignment will not 42
 43 be considered. Thus, we cannot get an assignment for all variables when the solution is 43
 44 not found. That is to say, in the process of BT, we cannot get an approximate solution. 44
 45 However, in the ERA method, every state, including the initial state, represents an 45

1 assignment to all variables, even though it may not be an exact solution. Therefore, 1
2 ERA is able to provide an approximate solution at *anytime*. 2

3 This property is useful for real-time systems that require a solution within a fixed time 3
4 interval while being not so demanding on the optimality of the solution. 4

5 (2) The system always evolves toward a better state in which more and more constraints 5
6 are satisfied. 6

7 Note that *random-p* is much smaller than *better-p* and *least-p* in most situations. 7
8 Thus, agents will have a greater chance to choose either *better-move* or *least-move*, in 8
9 order to effectively reduce the number of dissatisfied constraints. In the ERA method, 9
10 $e(a_i.x, i).violation$ records the violation number for the position at which agent a_i 10
11 resides. 11

12
13 **Definition 3.1.** $\lambda(s)$ represents in the current state s the sum of the violation numbers 13
14 for those positions at which agents reside, so $\lambda(s) = \sum_{i \in [1, n]} e(a_i.x, i).violation$. 14

15
16 If state s is not a solution state, that is, the positions of some agents are not *zero-* 16
17 *positions*, then $\lambda(s) > 0$. Otherwise, $\lambda(s) = 0$ for solution state s . 17

18
19 **Theorem 3.1.** For state s , $\lambda(s) = 0 \Leftrightarrow s$ is a solution state. 19

20
21 **Proof.** Because $\forall i \in [1, n], j \in [1, |D_i|], e(a_i.x, i).violation \geq 0$, we have 21
22 $\lambda(s) = 0 \Leftrightarrow \forall i \in [1, n], e(a_i.x, i).violation = 0 \Leftrightarrow s$ is a solution state (based on 22
23 Theorem 2.2). \square 23

24
25 Now we can see that the process of finding a solution is essentially a process of 24
26 minimizing the value of $\lambda(s)$. In initial state s_0 , $\lambda(s_0) > 0$ for most situations. 25
27 Gradually, as the system keeps on dispatching agents to move to smaller violation- 26
28 number positions, the value of $\lambda(s)$ will get minimized. When $\lambda(s)$ reaches *zero*, an 27
29 exact solution is found. 28

30
31 **Theorem 3.2.** When agent a_i moves from position (x_1, y) to position (x_2, y) , the 30
32 variation of λ can be computed as follows: $\Delta_\lambda = 2 \times (q_2 - q_1)$ where $q_1 = e(x_1,$ 31
33 $y).violation, q_2 = e(x_2, y).violation$. 32

34
35 **Proof.** We use Δ_{λ_1} to denote the variation after picking up a_i from (x_1, y) , Δ_{λ_2} to 34
36 denote the variation after placing a_i to (x_2, y) . When a_i is at position (x_1, y) , there 35
37 will be q_1 agents attacking a_i . When we pick up a_i from (x_1, y) , a_i 's contribution 36
38 to λ is *zero*, and all these q_1 agents' violation numbers are reduced by 1. Now 37
39 $\Delta_{\lambda_1} = -q_1 - q_1 \times 1 = -2 \times q_1$. Then we place a_i to (x_2, y) , a_i 's contribution 38
40 to λ is q_2 , and there will be q_2 agents attacking a_i . The violation numbers of all 39
41 these q_1 agents are increased by 1, i.e., $\Delta_{\lambda_2} = q_2 + q_2 \times 1 = 2 \times q_2$. So, we have 40
42 $\Delta_\lambda = \Delta_{\lambda_1} + \Delta_{\lambda_2} = -2 \times q_1 + 2 \times q_2 = 2 \times (q_2 - q_1)$. \square 41

43
44 **Theorem 3.3.** After agent a_i performs *better-move* or *least-move*, $\Delta_\lambda \leq 0$. 43

45
46 **Proof.** Because $\psi_{-l}(x, y) = \varphi(y)$, 44
45

$$e(x, y).violation \geq e(x, \varphi(y)).violation$$

(based on Definitions 2.5 and 2.6), (5)

$$\Delta_\lambda = 2 \times (e(x, \varphi(y)).violation - e(x, y).violation)$$

(based on Theorem 3.2). (6)

Thus, for *least-move*, $\Delta_\lambda \leq 0$.

And since

$$\psi_{-b} = \begin{cases} x, & \text{when } e(\text{Random}(|D_y|), y).violation \\ & \geq e(x, y).violation \\ \text{Random}(|D_y|), & \text{when } e(\text{Random}(|D_y|), y).violation \\ & < e(x, y).violation \end{cases}$$

we have

$$e(x, y).violation \geq e(x, \psi_{-b}(x, y)).violation, \quad (7)$$

$$\Delta_\lambda = 2 \times (e(x, \psi_{-b}(x, y)).violation - e(x, y).violation)$$

(based on Theorem 3.2). (8)

Thus, for *better-move*, $\Delta_\lambda \leq 0$. \square

In the process of dispatching, agents will have a much higher probability to perform either a *better-move* or a *least-move*. So after each move, $\Delta_\lambda \leq 0$, which means that λ is decreasing and the system is improving the solution. However, $\lambda(s_1) < \lambda(s_2)$ does not always indicate that state s_1 is better than state s_2 . For instance, if s_1 is a local optimum state, which means all agents are at *minimum-positions*, but not all of them are at *zero-positions*, it is hard for s_1 to move to a new state except using a *random-move*. So, λ is just one of the important criteria.

- (3) After a few steps, the assignments of most variables will satisfy constraints.

Because we randomly place the agents at the initialization step, they will seldom be placed right at good positions. In other words, they will most likely be placed at positions that have large violation numbers. After one time step, many agents that apply the behavior of *better-move* or *least-move* will move to positions with smaller violation numbers, which means $\Delta_\lambda < 0$. The improvement achieved at the first time step will be very substantial because $|\Delta_\lambda|$ is large. While at the following time step, the chance for finding a smaller violation position is becoming less as many agents are already at the *minimum-positions* of their respective rows. After r time steps (to different problem r is different), the variation of λ will fluctuate up and down around a fixed value. In this state, the agents usually stay at their original positions and only the behavior of *random-move* will make them move to other positions. In such states, λ is very small and the corresponding solution can satisfy most of the constraints. So, if we do not require an exact solution, we may stop the system and get an approximate solution from the current state. This phenomenon can readily be observed in the experiments of solving n -queen problems and coloring problems as to be discussed in the next section.

4. Empirical studies on extended ERA methods with behavior prioritization and different selection probabilities

The preceding sections have presented the basic ERA formulation and algorithm, and discussed some of their key properties. In this section, we will further examine the extensions and performance of the ERA approach under various behavioral settings. Specifically, the goal of this section is threefold:

- (1) It presents several empirical results on solving different n -queen and coloring problems.
- (2) It discusses how to apply and implement this approach by choosing the probabilities of *least-move* and *random-move*.
- (3) It examines the effectiveness of prioritizing agent behaviors in order to efficiently derive an approximate solution.

In the experiments, we will initialize all agents with the same set of parameters, i.e., $(\text{better-}p, \text{least-}p, \text{random-}p)$. Specifically, $\forall i \in [1, n], a_i.\text{better-}p = \text{better-}p, a_i.\text{least-}p = \text{least-}p, a_i.\text{random-}p = \text{random-}p$.

4.1. n -queen problem

An n -queen problem has been stated in Example 1.1. It is required to place n queens on an $n \times n$ chessboard so that no two queens are in the same row, or the same column, or the same diagonal. This problem is a good benchmark because its problem size n can vary from 4 to a very large number. Also the solution to an n -queen problem can find many practical applications [39].

An n -queen problem can be translated into a binary CSP as described in Example 1.1 of Section 1.1:

- n queens: $X = \{X_1, X_2, \dots, X_n\}$.
- $n \times n$ chessboard: $\mathbf{D} = \{D_1, D_2, \dots, D_n\}, \forall i, D_i = \{1, 2, \dots, n\}$.
- *Placement requirement*:

$$\mathbf{C} = \{C(R_u) \mid \forall i, j \in [1, n], C(R_u) = \{\langle b, c \rangle \mid b \in D_i, c \in D_j, b \neq c, i - j \neq b - c, i - j \neq c - b\}\}.$$

In this CSP, each variable has the same domain $[1, n]$ and there is a constraint between every two variables. In the following paragraphs, we will show how to apply the basic ERA algorithm to solve this problem.

First, we use n agents to represent n queens (variables), a_i represents variable X_i , which is a queen in row_i . Second, we model the domains, i.e., the chessboard, as the environment of the agents (see Fig. 14(a)). Fig. 14(b) presents an example of the multi-agent system for a 4-queen problem.

At the initialization step, the domain values will be recorded as $e(i, j).\text{value}$ (see Fig. 14(a)) and the violation numbers for all positions will be set to *zero* (see Fig. 15(a)).

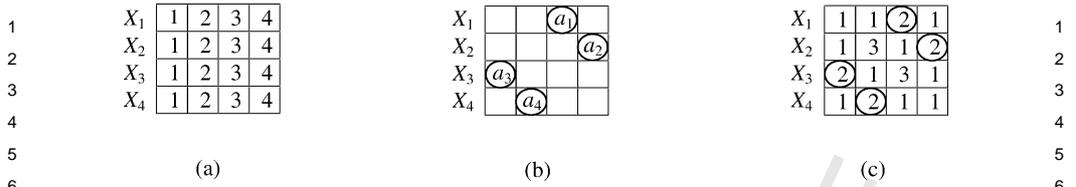


Fig. 14. (a) The representation of domain values for a 4-queen problem. (b) Four agents dispatched into the 4-queen environment. (c) Updated violation numbers corresponding to the positions of the four agents.

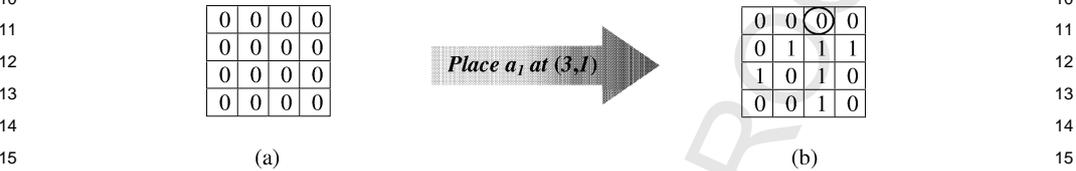


Fig. 15. (a) Violation numbers at the initialization step. (b) Violation numbers updated having placed a_1 at (3, 1).

After that, agents will be randomly placed onto different rows. For instance, if agent a_1 is placed at position (3, 1), the violation numbers in the environment will be updated accordingly, as shown in Fig. 15(b).

Fig. 16 presents a series of snapshots from an 8-queen problem experiment. Here each circle signifies an agent. The number on the lattice gives the corresponding violation number. The deeper the color of a lattice, the larger the violation number of that position will be. First, in the initialization of time step 0 in Fig. 16(a), eight agents are randomly placed onto the rows. In this particular case, none of the agents is at a *zero-position*. Five of them are at the positions of *violation* = 3. Two agents are at the positions of *violation* = 2. One agent is at the position of *violation* = 1. Obviously, the assignment according to this state is not a solution. For agent a_1 at position (4, 1), we can observe that agent a_2 at (3, 2) and agent a_3 at (2, 3) are both in the same diagonal as a_1 , and agent a_4 is in the same column as a_1 . So the position where a_1 stays has the violation number of 3. In this state of s_0 , $\lambda(s_0) = 3 + 3 + 3 + 3 + 3 + 2 + 2 + 1 = 20$.

Between time step 0 and time step 1, most agents can move to a better position that reduces the violation number. At time step 1, eight agents have moved to a better position. Now we have $\lambda(s_1) = 1 + 1 + 0 + 1 + 1 + 0 + 0 + 0 = 4$. In this assignment, four variables (i.e., X_3, X_6, X_7 , and X_8) satisfy all the constraints applicable to them. Two pairs, $\langle X_1, X_5 \rangle$ and $\langle X_2, X_4 \rangle$, cannot satisfy the constraints: X_1 and X_5 are in the same diagonal, and X_2 and X_4 are in the same column. Obviously, the assignment in state s_1 is much better than the assignment in state s_0 .

From time step 1 to time step 2, the following moves have occurred: $s_1 \Rightarrow a_4$ stays, a_5 least-move to (6, 5), a_3 stays, a_6 stays, a_2 stays, a_7 stays, a_1 stays, and a_8 stays $\Rightarrow s_2$. Now, $\lambda(s_2) = 0 + 1 + 0 + 1 + 0 + 0 + 0 + 0 = 2$. In this assignment, six variables (i.e., X_1, X_3, X_5, X_6, X_7 , and X_8) satisfy all the constraints related to them, while the pair of $\langle X_2, X_4 \rangle$ cannot satisfy each other.

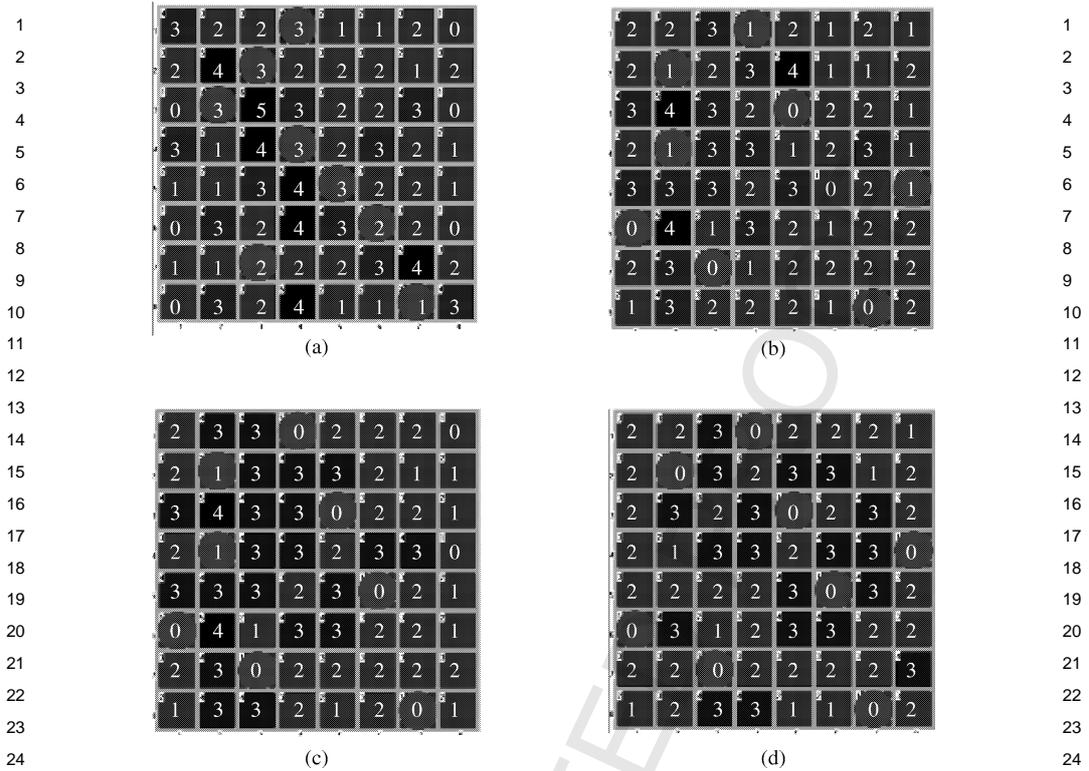


Fig. 16. (a) s_0 at time step 0 (initialization). (b) s_1 at time step 1. (c) s_2 at time step 2. (d) s_3 at time step 3, which is an exact solution state.

From time step 2 to time step 3, the moves can be summarized as follows: $s_2 \Rightarrow a_4$ *least-move* to (8, 4) but all other agents remain at the same positions $\Rightarrow s_3$. Now, $\lambda(s_3) = 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 = 0$. An exact solution state is reached.

4.1.1. The effects of the *least-p/random-p* ratio and behavior prioritization

Having illustrated the process of the basic ERA system in solving an n -queen problem, let us now consider the effects of behavior selection probabilities and behavior prioritization on the efficiency of finding an exact solution. In the experiments, we will let the system run until an exact solution is found. For the ease of comparison, we will record the average runtime required for generating the solution. Based on the observations made from the experiments, we will empirically show how the basic ERA method can be improved by adjusting *random-p* and by prioritizing behaviors.

4.1.1.1. The *least-p/random-p* ratio. From the above sections, we know that besides *better-move* and *least-move*, *random-move* is also necessary. If there is no *random-move*, i.e., $random-p = 0$, the system may get stuck in a local optimum and cannot find a solution. Now, the question that remains is how to set the probability ratio, *least-p/random-p*. It is

Table 1
Average runtime for different ratios of *least-p* to *random-p*

average runtime (s)		<i>least-p</i> : <i>random-p</i>			
		$0.5n$	n	$1.5n$	$2n$
n	1000	17.13	12.13	<u>8.63</u>	15.75
	2000	91.38	35	<u>31.13</u>	46.88
	3000	124.88	120.75	<u>79.75</u>	113.13
	4000	270.88	151.88	<u>150.25</u>	187.25
	5000	598.25	451.13	<u>370.75</u>	530.38
	6000	641.38	722.88	496.75	<u>426.38</u>
	7000	3478.75	1476.5	<u>1011.5</u>	1971.5

not so intuitive to see the most effective setting for solving a problem. In the following experiment, we will try to empirically determine a good *least-p/random-p* ratio.

Experiment 4.1. $n = \{1000, 2000, 3000, 4000, 5000, 6000, 7000\}$, *least-p/random-p* = $\{0.5n, n, 1.5n, 2n\}$, *type* = LR (10 runs). Note that the reason that we test n -queen problems up to 7000 queens is because beyond that number, the memory limitation of our computer becomes a problem. The behavioral type of LR indicates that in this experiment, the agents will only use the *least-move* and *random-move* behaviors.

Observation 4.1. The bold and underlined numbers in Table 1 correspond to the shortest ones among all the ratios for each n . As shown in the table, for all n except $n = 6000$, the ratio of *least-p* to *random-p* that results in the shortest runtime is $1.5n$. Therefore, we can obtain an empirical rule for setting the ratio of *least-p/random-p*: For n -queen problems, the good ratio is $1.5n$.

4.1.1.2. The high-priority better-move. Behaviors *better-move* and *least-move* are similar: move to a position based on the violation number. At each time step, it would be much easier for an agent using *better-move* to find a better position to move to than for the one using *least-move*. This is because *least-move* checks all the positions in its row, whereas *better-move* checks only one position. Therefore, the time complexity of *better-move* is much less than that of *least-move*.

In order to take the advantage of *better-move*, we decide to set this behavior to the highest priority, which means that an agent will first use *better-move* to compute its new position. If it fails to find a better position to move to, the agent will then turn to *least-move*. We call this behavior *better-least-move*, the probability of which will be the same as that of *least-p*.

Initially, most agents are at the positions with large violation numbers. The chance of successfully finding a position to move to with *better-move* is quite high. Thus at the first step, most agents will perform a *better-move* instead of a *least-move*.

Further to the above prioritized *better-move*, the next question is whether more *better-move* attempts will be helpful (since their complexity is low—a comparison operation). If so, how many *better-move* attempts will be most effective? In order to examine these issues, we will introduce another *better-move* right after the first *better-move* fails in finding a better position. For example, if a_i at position (x, y) performs a *better-move*, it will compute

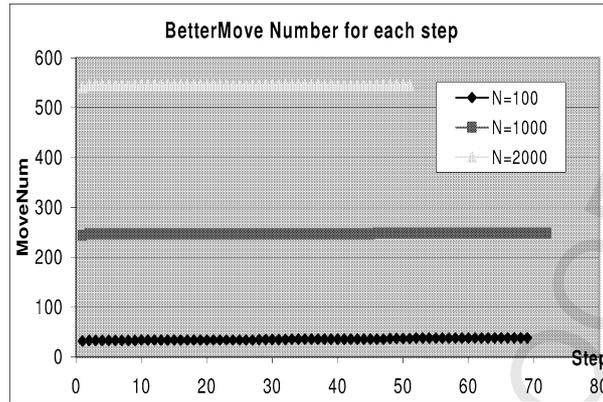


Fig. 17. Total better-move number vs clock step.

$\psi_{-b}(x, y)$. If $\psi_{-b}(x, y) = x$, then it will compute $\psi_{-b}(x, y)$ again. If at this time $\psi_{-b}(x, y) = x$, it will then turn to *least-move*. We call this behavior BBLR (i.e., one *better-move* and if it cannot find a better position, it will perform BLR) type of behaviors. Extending this concept, with the probability of *least-p*, the agents may perform at most r times *better-move* until it finds a better position to move to. We denote this behavior prioritization r BLR and its complexity is r .

4.1.1.3. The number of better-move at each step. Now, let us examine how many *better-move* attempts are needed in order to find a better position at each time step of the system. We will test and record the total number of *better-move* for all agents successfully finding a better position with *better-least-move* and *random-move* (this type of combined behaviors is called BLR). In this experiment, $n = 100, 1000, \text{ and } 2000$, respectively.

Experiment 4.2. $n = \{100, 1000, 2000\}$, *least-p/random-p* = n , *type* = BLR.

Observation 4.2. From Fig. 17, we note that except at time step 1, the lines for $n = 100, 1000, \text{ and } 2000$ go very smoothly without so much changes. That means *better-move* enables most agents to find better positions at the first step.

Specifically, 32 agents in the 100-queen problem, 247 agents in the 1000-queen problem, and 541 agents in the 2000-queen problem can move to a better position by using *better-move* at the first time step. And after the first step, very few agents can find a better position with *better-move*. Based on the observation, we may change the type of behavior prioritization into FBLR, which means agents will perform BLR *only at the first (F) time step* and then perform LR at the following steps. So, the type of combined behaviors changes during the process.

4.1.1.4. rBLR. Agents will apply this type for all steps. We will test the runtime for $n = 1000$ with $r = 1, 2, 3, \text{ and } 5$, respectively.

Experiment 4.3. $n = 1000$, *least-p/random-p* = n , *type* = {BLR, 2BLR, 3BLR, 5BLR} (10 runs).

Table 2
 Average runtime of $n = 1000$ with r BLR

	BLR	2BLR	3BLR	5BLR
average runtime (s)	6.13	4.5	5.25	10.25

Table 3
 Average runtime of $n = 1000$ with LR, r BLR, and Fr BLR

average runtime (s)		type		
		LR	r BLR	Fr BLR
r	1	12.25	6.13	5.25
	2		4.5	3.63
	3		5.25	4.75
	5		10.25	8.25

Observation 4.3. From Table 2, we note that the system has the best performance when $r = 2$. That means 2BLR will create more chance for agents to find a better position than BLR. The runtime complexity of 2BLR is less than 3BLR and 5BLR. So, 2BLR will be the best setting if we want to successfully find better positions, and at the same time, have less runtime complexity.

4.1.1.5. LR vs r BLR vs Fr BLR. The following experiment compares the average runtime among the prioritization types discussed above. LR means that an agent only has two moving behaviors: *least-move* (with the probability of *least-p*) and *random-move* (with the probability of *random-p*). Fr BLR means that the agent will perform r BLR behavior at the first step and then perform LR at all the following steps.

Experiment 4.4. $n = 1000$, $least-p/random-p = n$, $type = \{LR, BLR, 2BLR, 3BLR, 5BLR, FBLR, F2BLR, F3BLR, F5BLR\}$ (10 runs).

Observation 4.4. The results given in Table 3 show that Fr BLR has the best runtime performance among all the types. For $\forall r \in \{1, 2, 3, 5\}$, the runtime of Fr BLR is less than that of r BLR. So, we conclude that Fr BLR is better than r BLR, and r BLR is better than LR.

4.1.2. Approximate solution

In this section, we will study the performance of the system in finding an approximate solution by tracking at each time step the total number of agent moves, the number of agents at *zero-positions*, and the total number of *zero-positions*. The results of our experiments consistently indicate that $\forall n \in [100, 7000]$, the system will converge as follows:

- After 1 step, 80% agents are at *zero-positions*.
- After 2 steps, $n - c_1$ (where c_1 is a constant, $c_1 \approx 25$) agents are at *zero-positions*.
- After 3 steps, $n - c_2$ (where c_2 is a constant, $c_2 \approx 7$) agents are at *zero-positions*.

1 This property enables the system to efficiently find an approximate solution: 3 steps are 1
2 needed in order to find an approximate solution, in which about $n - 7$ queens will not 2
3 attack each other. 3
4 4

5 *4.1.2.1. Is the system stable?* In the following paragraphs, we will track the system step 5
6 by step and show how the system performs. The measures to be considered are as follows: 6
7 7

- 8 ● *Move-Num* (number of moves): First, we observe the total number of moves that 8
9 all agents have performed, from the beginning of the system till the current time 9
10 step (i.e., accumulative number). Before an exact solution is found, the agents in the 10
11 system should move to other positions to improve the current assignment. If no agent 11
12 moves, the system will not improve itself. Thus, the number of moves at one time step 12
13 measures the *improvement speed* of the system. 13
14 ● *Zero-Agent-Num* (number of agents at *zero-positions*): This measure expresses how 14
15 good the assignment in the current state is. If $\text{Zero-Agent-Num} = n$, the current state 15
16 is an exact solution state. A larger *Zero-Agent-Num* means more variables satisfy the 16
17 constraints. 17
18 ● *Zero-Position-Num* (*zero-position* number in the environment): If there are a lot of 18
19 *zero-positions* in the environment, the agents will have a good chance to find and move 19
20 to *zero-positions*. So, the *Zero-position-Num* implies the degree of difficulty to find a 20
21 *zero-position* to move to, and the difficulty for the system to improve itself. 21
22 22

23 Now let us examine the cases of $n = 100, 1000, \text{ and } 2000$. 23
24 24

25 **Experiment 4.5.** $n = \{100, 1000, 2000\}$, *least-p/random-p* = n , *type* = 2BLR. 25
26 26

27 **Observation 4.5.** 27
28 28

- 29 (1) *Move-Num*: As shown in Figs. 18 and 19, especially in large problem sizes such as 29
30 $n = 1000$, the agents move more drastically at the first time step, and thereafter the 30
31 total number of moves increases slowly. This means that the agents can easily find 31
32 a better position to move to at the first step, and then the chance of finding a better 32
33 position decreases as time goes by. 33
34 (2) *Zero-Position-Num*: As also shown in Figs. 18 and 19, except in small problem sizes 34
35 such as $n = 100$, the number of *zero-positions* drops sharply at the first 3 steps and 35
36 then fluctuates at the following steps. These may explain why *Move-Num* increases at 36
37 the first 3 steps and then remains unchanged. More *zero-positions* means more chance 37
38 for agents to find a better position. 38
39 (3) *Zero-Agent-Num*: The plots as shown in Figs. 18–19 all increase at the first 3 steps, 39
40 nearly reaching n , and then slightly fluctuate. This means that the system improves 40
41 itself quickly during the first 3 steps. 41
42 42

43 *4.1.2.2. Approximate solution: the first 3 steps.* Since we have observed that the system 43
44 will converge and its improvement will slow down after 3 steps, we can just let the system 44
45 45

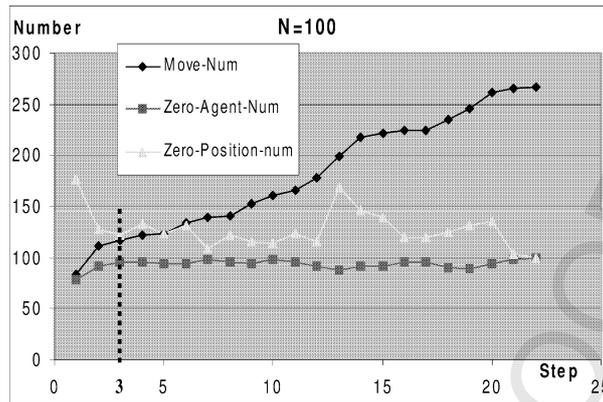


Fig. 18. Move-Num, Zero-Agent-Num, and Zero-Position-Num ($n = 100$).

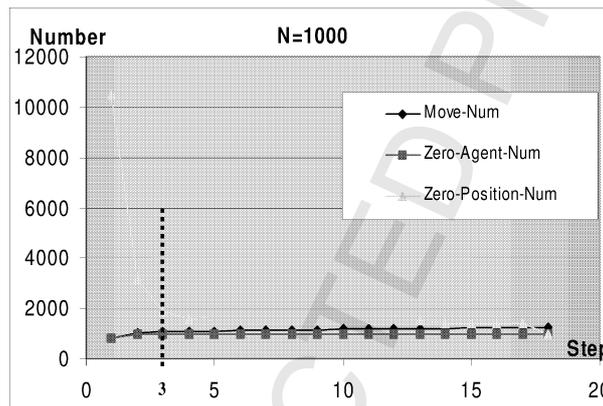


Fig. 19. Move-Num, Zero-Agent-Num, and Zero-Position-Num ($n = 1000$).

run for 3 steps and then get an approximate solution. Experiment 4.6 examines how well the system does at the first 3 steps.

Experiment 4.6. $n = [100, 7000]$, $\Delta n = 100$, $least-p/random-p = n$, $type = F2BLR$ (10 runs).

Observation 4.6.

- (1) Our experiments have shown that after the initialization, nearly 10% agents stay at *zero-positions*.
- (2) After the 1st step, nearly $80\% \times n$ agents stay at *zero-positions* when n is larger than 1000, as shown in Fig. 20. That means when $n > 1000$, there are 80% variables' assignments can satisfy the constraints. This result is desirable because it is obtained by just one step.

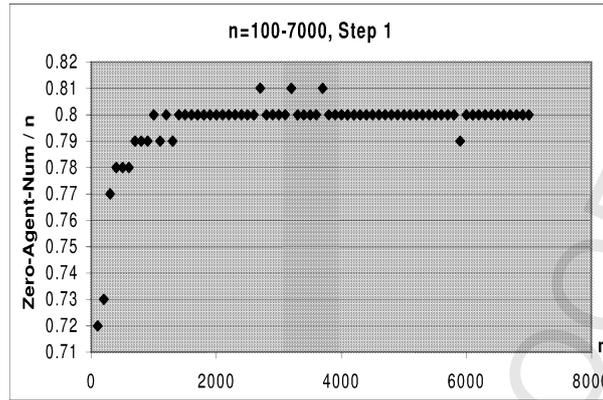


Fig. 20. Average Zero-Agent-Num/n of the 1st step for $n = 100$ to 7000 , $\Delta n = 100$.

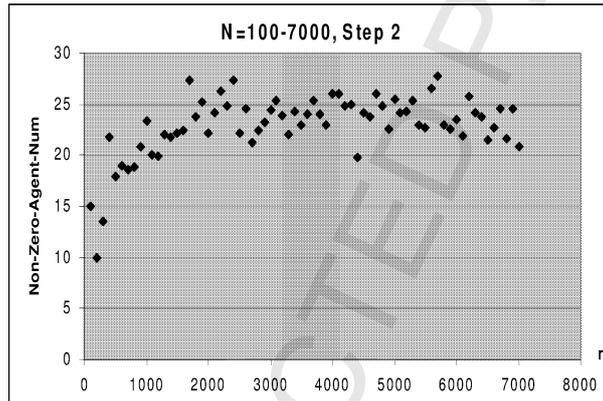


Fig. 21. Average Non-Zero-Agent-Num of the 2nd step for $n = 100$ to 7000 , $\Delta n = 100$.

- (3) After the 2nd step, nearly $n - 25$ agents stay at *zero-positions* when n is larger than 1000, as shown in Fig. 21. That means when $n > 1000$, there are about $n - 25$ variables' assignments can satisfy the constraints. On the other hand, there are about 25 variables' assignments cannot satisfy the constraints. This result is obtained by just two steps no matter how large n is.
- (4) After the 3rd step, nearly $n - 7$ agents stay at *zero-positions* when n is larger than 1000, as shown in Fig. 22. That means when $n > 1000$, there are about $n - 7$ variables' assignments can satisfy the constraints. This is a good approximate solution obtained in just three steps.

Based on Theorem 2.5, we know that the complexity for the first time step is $O(n \sum |D_i|)$. Thus, we can tell that the complexity for the approximate solution will be bounded by $O(n \sum |D_i|)$. For n -queen problems, $|D_i| = n$, the complexity will be bounded

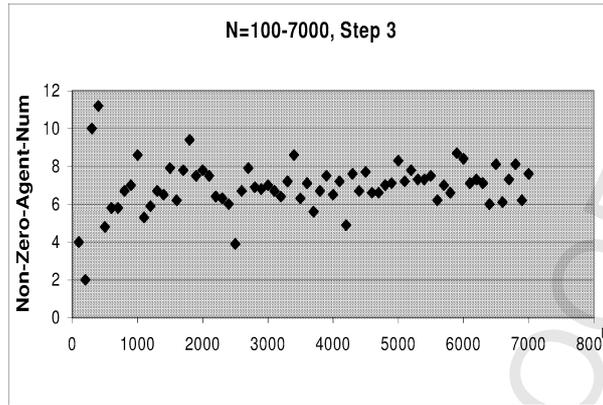


Fig. 22. Average *Non-Zero-Agent-Num* of the 3rd step for $n = 100$ to 7000, $\Delta n = 100$.

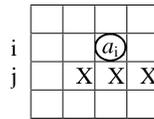


Fig. 23. Attacking 3 positions at most.

by $O(n^3)$. In fact, for such problems, the single time step complexity of the ERA method is $O(n^2)$.

Theorem 4.1. *The single time step complexity of the ERA method in solving n -queen problems is $O(n^2)$.*

Proof. The total number of operations in *RemoveFrom* and *AddTo* is $6n$ in the worst case. This is because between agent a_i and all other agents in *row* _{j} , there exist at most 3 positions in *row* _{j} (see Fig. 23), in which the agents will attack a_i . So the computation of *RemoveFrom* and *AddTo* is $3n$, in total $6n$. As for checking a solution state, it needs n tests in the worst case. Now for each agent's move, the total number of operations is $7n$. So the complexity for one time step is $7n^2$, bounded by $O(n^2)$, and the complexity for 3-step approximate solution finding is also $O(n^2)$. □

4.1.2.3. Runtime for finding an approximate solution. Now let us take a look at the runtime in finding an approximate solution in 3 steps.

Experiment 4.7. $n = \{100, 500, 1000, 2000, 3000, 4000, 5000, 6000, 7000\}$, *least-p/random-p* = n , *type* = F2BLR (10 runs for 3 steps each).

Observation 4.7. Fig. 24 shows that the empirical results are consistent with Theorem 4.1. To find an approximate solution in 3 time steps, in which nearly $n - 7$ variables satisfy the constraints, the runtime complexity is $O(n^2)$.

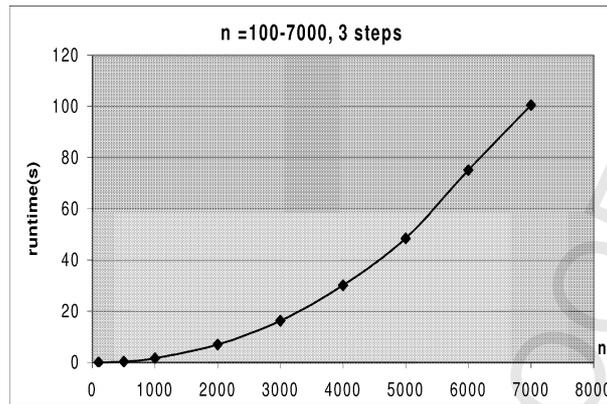


Fig. 24. Average runtime (second) for 3 steps ($n = 100 \sim 7000$).

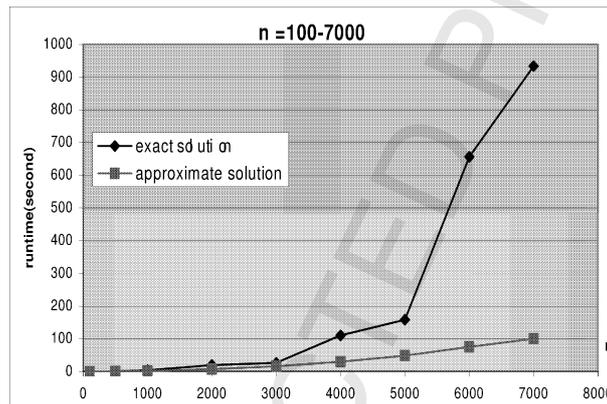


Fig. 25. Average runtime (second) for finding an exact solution and an approximate solution ($n = 100 \sim 7000$).

We have also compared the runtime for deriving an exact solution with that for finding an approximate solution. The results are presented in Fig. 25.

4.2. Coloring problem

A coloring problem has been defined in Example 1.2. Many problems of practical interest can be modeled as coloring problems, such as time tabling and scheduling [21], frequency assignment [14], register allocation [4], printed circuit board testing[15], and pattern labeling [31].

Given an undirected graph $G = (V, E)$, an m -coloring problem can be translated into a binary CSP as follows:

- n nodes, $\mathbf{V}: \mathbf{X} = \{X_1, X_2, \dots, X_n\}$, X_i represents $v_i \in \mathbf{V}$.
- m colors: $\mathbf{D} = \{D_1, D_2, \dots, D_m\}$, $\forall i, D_i = \{1, 2, \dots, m\}$.

1 • *Coloring requirement:*

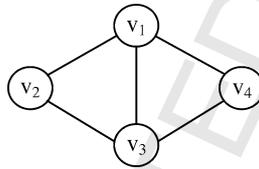
$$C = \{C(R_u) \mid \forall i, j \in [1, n] \langle v_i, v_j \rangle \in E, C(R_u) = \{\langle b, c \rangle \mid b \in D_i, c \in D_j, b \neq c\}\}.$$

2
3
4
5
6
7 In this case, all variables have the same domain, $[1, m]$, and there is a constraint between
8 two nodes incident to an edge. An example coloring problem is given in Fig. 26. This
9 instance can be colored by using three colors. Now, let us see how the ERA method works
10 in solving this problem.

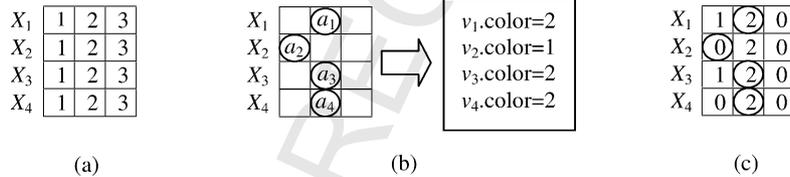
11 First, we use four agents to represent four nodes (variables). a_i represents node v_i
12 (variable X_i). Second, we model the domains as the environment of the agents (see Fig. 27).

13 Initially, the domain values (labels of color) will be recorded as $e(i, j).value$ (see
14 Fig. 27(a)) and the violation numbers for all positions will be set to *zero*. After that, we
15 will randomly place the agents onto different rows. For instance, if we place agent a_1
16 at position (2, 1), the violation numbers will be updated accordingly as shown in Fig. 28.

17 In what follows, we will examine the performance of the ERA method in solving a set
18 of large-scale coloring problems from DIMACS (the Center for Discrete Mathematics and
19



20
21
22
23
24
25
26 Fig. 26. A coloring problem.



27
28
29
30
31
32
33
34
35 Fig. 27. (a) The representation of domain values into the environment of a multi-agent system for the coloring
36 problem of Fig. 26. (b) Four agents dispatched into the environment, the positions of which correspond to a
37 specific color assignment. (c) Violation numbers updated.



38
39
40
41
42
43
44
45 Fig. 28. (a) Violation numbers at the initialization step. (b) Violation numbers updated having placed a_1 at (2,1).

1 Theoretical Computer Science) [47].¹ In particular, we will use the test problems from 1
2 *Donald Knuth's Stanford GraphBase* [48]. 2

3 Each problem includes the information of (nodes, edges), number of optimal colors, 3
4 and source. Thus, each problem has n variables, where n is the number of nodes, k values 4
5 for each domain, where k is the number of optimal colors, and m constraints, where m 5
6 is the number of edges. In the following listing, the source of SGB (from Michael Trick 6
7 (trick@cmu.edu)) refers to Donald Knuth's Stanford GraphBase: 7

- 8
9 (1) *miles250.col* (128,387), 8, SGB 9
10 (2) *miles500.col* (128,1170), 20, SGB 10
11 (3) *miles750.col* (128,2113), 31, SGB 11
12 (4) *miles1000.col* (128,3216), 42, SGB 12
13 (5) *miles1500.col* (128,5198), 73, SGB 13
14 (6) *anna.col* (138,493), 11, SGB 14
15 (7) *david.col* (87,406), 11, SGB 15
16 (8) *huck.col* (74,301), 11, SGB 16
17 (9) *jean.col* (80,254), 10, SGB 17
18 (10) *games120.col* (120,638), 9, SGB 18
19 (11) *inithx.i.1.col* (864,18707), 54, REG 19
20

21 In the above problems, *miles graphs* are similar to geometric graphs in that the nodes are 21
22 placed in space with two nodes connected if they are close enough. The nodes represent a 22
23 set of United States cities and the distance between them corresponds to the road mileage 23
24 recorded in 1947. Book graphs are created where each node represents a character and 24
25 two nodes are connected by an edge if the corresponding characters encounter each other 25
26 in the book. The book graphs were created by *Anna*, *David*, *Huck* and *Jean*. *games120* 26
27 from Knuth represents the 1990 college football season. In *games120*, the nodes represent 27
28 college teams and two teams are connected by an edge if they played against each other 28
29 during the season. *inithx.i.1* is a problem based on register allocation (named REG) as 29
30 contributed by Gary Lewandowski (gary@cs.wisc.edu). 30
31

32 4.2.1. Approximate solution 32

33 In the previous experiments on n -queen problems, we have shown that the system can 33
34 get a good approximate solution after 3 steps. Similarly, for coloring problems, we will also 34
35 test the performance of 3 steps. We will examine the runtime for 3 steps and the number of 35
36 agents at *zero-positions* (*Zero-Agent-Num*) at each time step. 36
37

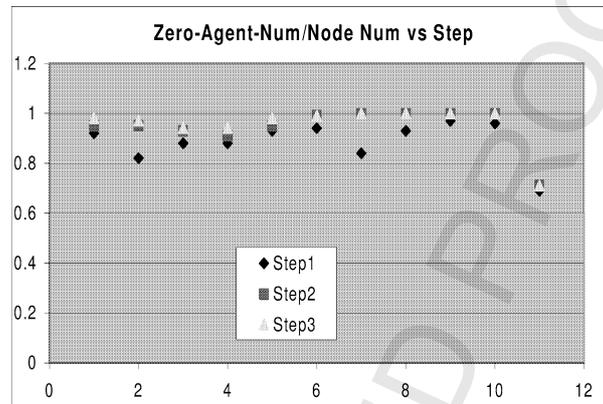
38 **Experiment 4.8.** *least-p/random-p = n*, *type = BBBLR* (10 runs for 3 steps each). 38
39

40 **Observation 4.8.** The system can quickly find an approximate solution. As shown in 40
41 Table 4, the average runtime measurement for all problems are close to *zero*. 41
42

43
44 ¹ More information can be found from the OR-Library [49] for test data sets for a variety of Operations 44
45 Research (OR) problems. 45

1 Table 4
 2 *Zero-Agent-Num* for coloring problems after 3 steps

measurements	<i>miles</i> 250	<i>miles</i> 500	<i>miles</i> 750	<i>miles</i> 1000	<i>miles</i> 1500	<i>anna</i>	<i>david</i>	<i>huck</i>	<i>jean</i>	<i>games</i> 120	<i>inithx.i.1</i>
runtime (s)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
<i>zero-agent-num</i>	124	124	122	122	120	134	86.8	74	80	120	604.2
nodes	128	128	128	128	128	138	87	74	80	120	864
<i>zero-agent-num</i> (%)	0.97	0.97	0.95	0.95	0.94	0.97	1	1	1	1	0.7



23 Fig. 29. *Zero-Agent-Num* (%) at each time step for 11 coloring problems. The number on the horizontal axis
 24 corresponds to an individual problem: (1) *miles250*, (2) *miles500*, (3) *miles750*, (4) *miles1000*, (5) *miles1500*, (6)
 25 *anna*, (7) *david*, (8) *huck*, (9) *jean*, (10) *games120*, and (11) *inithx.i.1*.

27 **Observation 4.9.** After the first step, the system can find an approximate solution for all
 28 problems except *inithx.i.1*, with more than 80% variables satisfying the constraints. For the
 29 *miles* problems, almost 95% variables (nodes) can satisfy the constraints after 3 steps. For
 30 the problems of *david*, *huck*, *jean*, and *games120*, the system can find an exact solution
 31 within 3 steps. In fact, in this experiment, they need only 2 time steps to find an exact
 32 solution (see Fig. 29). An exception is the problem of *inithx.i.1*, where about 70% variables
 33 can satisfy all the constraints.

35 4.2.1.1. *Zero-Agent-Num* at each time step. In order to study the performance of this
 36 system, we will record *Zero-Agent-Num* at each time step. The problems of *david*, *huck*,
 37 *jean*, and *games120* will be excluded from the following experiment, since the system can
 38 find an exact solution for these problems within 2 steps.

40 **Experiment 4.9.** *least-p/random-p = n*, *type = FBBLR* (1 run).

42 **Observation 4.10.** The results from this experiment are shown in Fig. 30, where the
 43 *Zero-Agent-Num* curves fluctuate around 120 after 3 steps. The value at the last time
 44 step corresponds to an exact solution found. This observation has been quite consistent,
 45

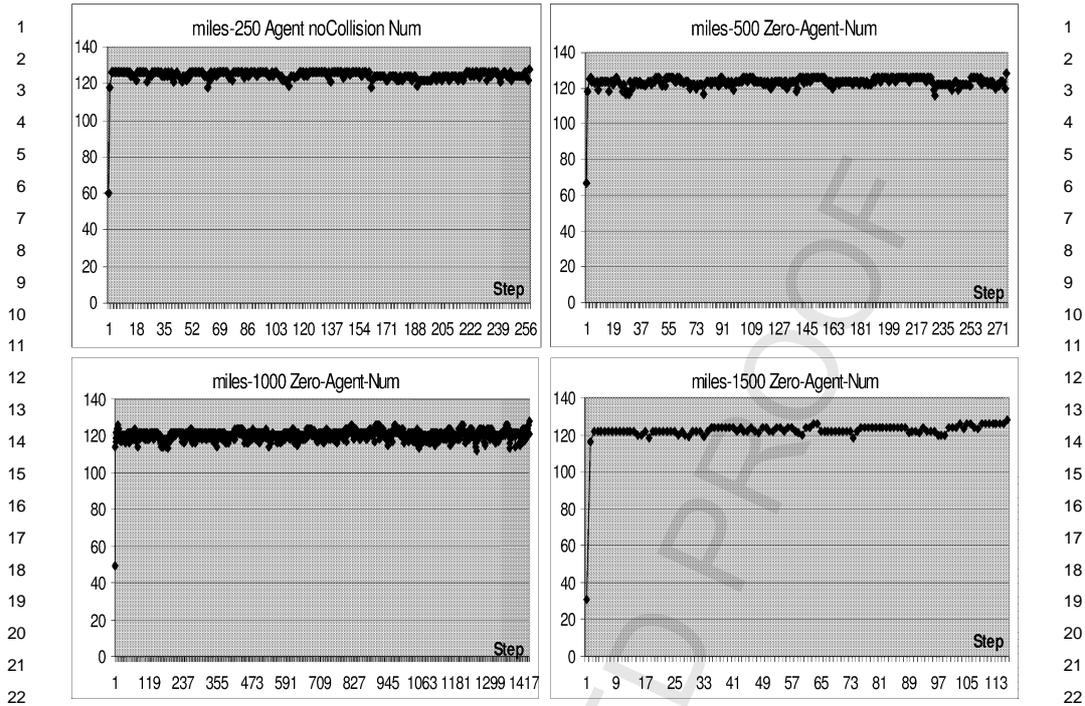


Fig. 30. Zero-Agent-Num at different time steps, obtained from the experimental studies on the miles problems.

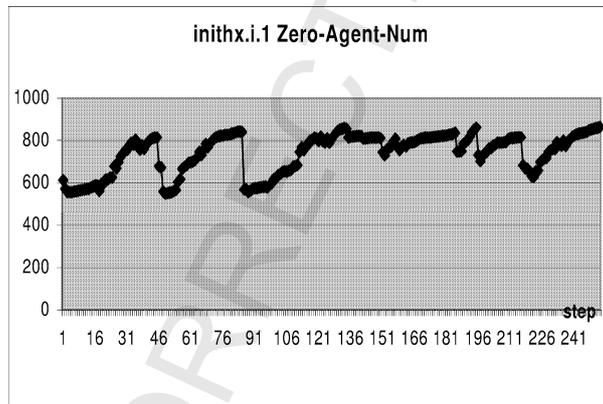


Fig. 31. Zero-Agent-Num at each time step in the case of inithx.i.1.

meaning that we can get an approximate solution with about 120 (95%) variables satisfying the constraints at anytime after 3 steps.

Observation 4.11. In Fig. 31, Zero-Agent-Num increases during the first 45 steps and then falls down. This process is repeated for several times before it finally reaches an exact

1 solution state. The speed of convergence is not as high as that in other problems. This 1
2 indicates that not all the problems can converge to 80% at the first step and to 95% at the 2
3 third step, while using the ERA method. The performance is, to some extent, affected by 3
4 the structure of the problem. 4

5 5. Discussion 5

6
7
8
9 In this section, we will compare the basic as well as extended ERA approach to 9
10 the existing heuristic and distributed approaches, and discuss their distinct features and 10
11 advantages. 11

12 5.1. Comparison with min-conflicts heuristics 12

13
14
15 The proposed ERA approach differs from the min-conflicts approach in the following 15
16 aspects: 16

17
18 (1) In the min-conflicts hill-climbing system reported in [28], the system chooses a 18
19 variable at each step that is currently in conflict and reassign its value by searching the 19
20 space of possible assignments and selecting the one with the minimum total conflicts. 20
21 The hill-climbing system can get trapped in a local minimum (note that the same 21
22 phenomenon can also be observed from the GDS network for constraint satisfaction). 22
23 On the other hand, in our approach, an agent is given a chance to select a *random-move* 23
24 behavior according to its probability, and hence it is capable of escaping from a local 24
25 trap. In our present work, we also note that the extent to which the agents can most 25
26 effectively avoid the local minima and improve their search efficiency is determined 26
27 by the probabilities (i.e., behavior selection probabilities) of the *least-move* (as well as 27
28 *better-least-move*) and *random-move* behaviors. 28

29 (2) Another system introduced in [28] is called *informed backtracking*. It argues a 29
30 standard backtracking method with the *min-conflicts* ordering of the variables and 30
31 values. This system attempts to find a sequence of repairs, such that no variable is 31
32 repaired more than once. If there is no way to repair a variable without violating 32
33 a previously repaired variable, the algorithm backtracks. It incrementally extends a 33
34 consistent partial assignment in the same way as a constructive backtracking program, 34
35 however, it uses information from the initial assignment to guide its search. 35

36 The key distinction between this approach and ours is that our approach does not 36
37 require backtracking. As stated by Minton et al. [28], their system trades search 37
38 efficiency for completeness; for large-scale problems, terminating in a no-solution 38
39 report will take a very long time. 39

40 (3) In both *min-conflicts* hill-climbing and informed backtracking systems proposed 40
41 in [28], the key is to compute and *order* the choice of variables and values to consider. 41
42 It requires to test all related constraints for each variable and to test all its possible 42
43 values. 43

44 This step is similar to the *RemoveFrom* and *AddTo* operations in our approach, except 44
45 that we only test a selected position (one value for each variable) and do not sort the 45

1 variables. The use of the ordering heuristic can lead to excessive assignment evaluation 1
2 preprocessing and therefore will increase the computational cost at each step. 2

3 (4) In our present approach, we examine the use of a *fewer-conflicts* repair, by introducing 3
4 the *better-move* behavior, that requires only one violation number evaluation for each 4
5 variable. The empirical evidence has shown that the use of the high-priority *better-* 5
6 *move* when combined with other behaviors can achieve more efficient results. We 6
7 believe that the reason that using the currently-available min-conflicts value at each 7
8 step can compromise the systems performance is because the min-conflicts values 8
9 quickly reduce the number of inconsistencies for some variables but at the same time 9
10 also increase the difficulties (e.g., local minima) for other variables. 10

11 12 5.2. Comparison with Yokoo et al.'s distributed constraint satisfaction 12

13 Our multi-agent approach has several fundamental distinctions from Yokoo et al.'s 13
14 distributed constraint satisfaction approach, as listed below: 14

15 (1) Yokoo et al.'s approach does not require a global broadcasting mechanism or data 15
16 structure. It allows agents to communicate their constraints to others by sending and 16
17 receiving messages such as *ok?*, and *nogood*. In other words, their methods handle 17
18 the violation checking among agents (variables) through agent-to-agent message 18
19 exchanges, such that each agent knows all instantiated variables relevant to its own 19
20 variables. 20
21

22 In our approach, the notion of agent-to-agent communication is implicit—we assume 22
23 that for violation updating, each agent (representing the value of a variable) is 23
24 'informed' about the values from relevant agents (representing the values of relevant 24
25 variables) either by means of accessing an $n \times n$ look-up memory table or via 25
26 pairwise value exchange—both implementations enable an agent to obtain the same 26
27 information, but the latter can introduce significant communication overhead costs 27
28 (i.e., longer *cycles* required [44]²) to the agents. 28
29

30 As the communication in Yokoo et al.'s approach is quite distributed, we believe that 30
31 their approach will work well under a large number of constrained conditions. 31

32 (2) In the *asynchronous weak-commitment search* algorithm developed by Yokoo et al. [43, 32
33 44], a consistent partial solution is incrementally extended until a complete solution 33
34 is found. When there exists no value for a variable that satisfies all the constraints 34
35 between the variables included in the partial solution, this algorithm abandons the 35
36 whole partial solution and then constructs a new one. Although *asynchronous weak-* 36
37 *commitment search* is more efficient than *asynchronous backtracking*, abandoning par- 37
38 tial solutions after one failure can still be costly. In the case of the ERA approach, the 38
39 high-level control mechanism for maintaining or abandoning consistent partial solu- 39
40 tions does not exist. 40

41 Yokoo et al. [43] have also developed a non-backtracking algorithm called *distributed* 41
42 *breakout*, which provides a distributed implementation for the conventional breakout. 42

43
44 ² As stated in [44], "one drawback of this model is that it does not take into account the costs of 44
45 communication". 45

Table 5
Comparison (in averaged number of cycles) between ERA and Yokoo et al.'s distribution constraint satisfaction in solving benchmark n -queen problems

n	Asynchronous backtracking	Asynchronous backtracking with min-conflicts heuristic	Asynchronous weak-commitment	ERA
100	510	504	51	22
1,000	–	324	30	18
2,000	–	–	–	30

(3) In *asynchronous weak-commitment search*, each agent utilizes the *min-conflicts* heuristic as mentioned in Section 5.1 to select a value from those consistent with the *agent_view* (those values that satisfy the constraints with variables of high-priority agents, i.e., value-message senders).

On the other hand, the ERA approach utilizes a combination of value-selection heuristics that involves a *better-move* behavior for efficiently finding *fewer-conflicts* repairs.

(4) As related to the above two remarks, the *asynchronous weak-commitment search* and *asynchronous backtracking* algorithms are designed to achieve completeness and thus the steps of backtracking and incremental solution constructing/abandoning are necessary, whereas the ERA approach is aimed at more efficiently finding an *approximate* solution, which is useful when the amount of time available for an exact solution is limited.

(5) Last but not the least, we have also systematically compared the performance of the ERA system with that of Yokoo et al.'s algorithms, namely, asynchronous backtracking, asynchronous backtracking with min-conflicts heuristic, and asynchronous weak-commitment, in solving benchmark n -queen problems where $n = 100, 1000, 2000$, respectively [44]. The averaged numbers of cycles used in each case are summarized in Table 5. We can establish that as demonstrated in solving the benchmark n -queen problems, ERA is an effective approach and the number of cycles used in the ERA system is competitive with those by Yokoo et al.'s approach, given that our formulation utilizes different behavior prioritization and violation checking schemes. Note that the '–' symbol in the table indicates that the data item is presently unavailable.

In summary, as complementary to each other, *both Yokoo et al.'s asynchronous approach and the ERA approach can be very efficient and robust when applied in the right context*. For instance, in some practical applications such as distributed telecommunication networks, Yokoo et al.'s formulation involving agent information exchange offers a natural way of modeling and solving the distributed CSP, whereas in CSPs that do not lend themselves so well to partitioning variables and constraints into sub-problems, the ERA formulation becomes straightforward to implement and execute. A distinct feature of Yokoo et al.'s asynchronous approach is, like other standard backtracking techniques, its completeness, whereas the feature of the ERA approach lies in its efficiency and robustness in obtaining an approximate solution within a few time steps (although it empirically always produces an exact solution when enough time steps are allowed). The ERA approach

1 is not guaranteed to be complete since it involves random moves. Another feature of the 1
2 ERA approach is that its strategies are quite easy to implement. 2

3 4 5.3. Remarks on partial constraint satisfaction 4

5
6 Partial constraint satisfaction is a very desirable way of solving CSPs that are either 6
7 overconstrained or too difficult to solve [42]. It is also extremely useful in situations where 7
8 we want to find the best solution obtainable within fixed resource bounds or in real-time. 8
9 Freuder and Wallace [11] are the pioneers in systematically studying the effectiveness of a 9
10 set of partial constraint satisfaction techniques using random problems of varying structural 10
11 parameters. The investigated techniques included basic branch and bound, backjumping, 11
12 backmarking, pruning with arc consistency counts, and forward checking. Based on the 12
13 measures of constraint checks and total time to obtain an optimal partial solution, forward 13
14 checking was found to be the most effective. Also of general interest is that their work has 14
15 offered a model of partial constraint satisfaction problems (PCSPs) involving a standard 15
16 CSP, a partially ordered space of alternative problems, and a notion of distances between 16
17 these problems and the original CSP. 17

18 Our present work attempts to develop, and empirically examine, an efficient technique 18
19 that is capable of generating partial constraint satisfaction solutions. This work shares the 19
20 same motivation as that of Freuder and Wallace's work [11,42], and also emphasizes that 20
21 the costs of calculating (communicating) and accessing constraint violation information 21
22 should be carefully considered in developing a practically efficient technique. That is 22
23 also part of the reason why much attention in our work has been paid to (1) the use of 23
24 environmentally updated and recorded violation numbers (without testing from the scratch 24
25 for each variable) and (2) the effectiveness of the *better-move* behavior in finding an 25
26 approximate solution. 26

27 28 5.4. Remarks on agent information and communication for conflict-check 28

29
30 Yokoo et al.'s approach and the ERA approach have a common thread; both formula- 30
31 tions employ multiple agents that reside in an *environment* of variables and constraints (al- 31
32 though in the ERA approach, the environment also contains violation information, which 32
33 is analogous to the 'artificial pheromone' in an *ant system* [8,9]) and make their own deci- 33
34 sions in terms of how the values of certain local variables should be searched and selected 34
35 in the process of obtaining a globally consistent solution. 35

36 Nevertheless, it should be pointed out that the present implementations of the two 36
37 approaches differ from each other in the way in which the agents record and access their 37
38 conflict-check information. The former utilizes a sophisticated communication protocol 38
39 to enable the agents representing different groups of variables and constraints to exchange 39
40 their values. By doing so, the agents are capable of evaluating constraint conflict status with 40
41 respect to other relevant agents (variables). On the other hand, our implementation utilizes a 41
42 feature of agent current-value broadcast to enable other agents to compare with their values 42
43 and to update the violation numbers in their local environment. Although the formulations 43
44 may seem different, the objectives as well as effects of them are fundamentally similar. 44
45 The reasons that we decided to use value broadcast and sharing are threefold: First, the 45

1 implementation can make use of a common storage space of complexity $O(n)$ where n 1
 2 corresponds to the number of variables and by doing so avoid introducing the same space 2
 3 requirement to every agent; secondly, it can reduce the overhead costs incurred during the 3
 4 pairwise information exchange, which can be quite significant; and thirdly, since our ERA 4
 5 method extensively uses *fewer-conflicts* moves, such behaviors can be triggered based on 5
 6 only one possible violation evaluation instead of n assignment evaluations, and hence the 6
 7 access to such a broadcast information source is not demanding. 7

8 9 5.5. Remarks on sequential-iteration implementation 9

10 11 Theoretically, the ERA method using swarm-like agents can be implemented in a 11
 12 parallel fashion, however, in light of our resource limitation, we used a sequential 12
 13 computation implementation to simulate the multi-agent concurrent or synchronous actions 13
 14 and to test the effectiveness of our approach. 14

15 16 Our sequential simulation utilizes a global simulated clock, called *time step*. The state 16
 17 of the environment as well as the agents (i.e., the positions of the agents) will be changed 17
 18 only at each discrete time step. In order to simulate the concurrent or synchronous actions 18
 19 of the agents at time step k , we let the individual agents perform their cycles of behavior 19
 20 selection, value selection, and violation updating. In so doing, the agents are dispatched in 20
 21 a sequential fashion. Once this is completed, the state of the system will then be refreshed 21
 22 with the new positions of the agents corresponding to the newly-selected values, and 22
 23 thereafter, the simulated clock will be incremented to $k + 1$. 23

24 25 Here, it is worth mentioning that apart from the fact that our implementation 24
 26 simulates the operations of a parallel system, the empirical results of our sequential 25
 27 ERA implementation are still comparable to those reported in [28,44], if we evaluate the 26
 28 performance using the measures of number of constraint checks, as introduced by Freuder 27
 29 and Wallace [11], and space complexity. 28

29 30 31 6. Summary 31

32 33 In this paper, we have described a multi-agent oriented approach to solving constraint 33
 34 satisfaction problems, such as n -queen problems and coloring problems. The key ideas 34
 35 behind this approach rest on three notions: Environment, Reactive rules, and Agents 35
 36 (ERA). Each agent can only sense its local environment (i.e., violation numbers) and apply 36
 37 some behavioral rules for governing its value-selection moves. The environment records 37
 38 and updates the local values that are computed and affected according to the moves of 38
 39 individual agents (analogous to the idea of laying 'artificial pheromone' in an *ant system* 39
 40 [8,9]). 40

41 42 In solving a CSP with either a basic or an extended ERA method, each agent represents 41
 43 a variable and its position corresponds to a value assignment for the variable. The 42
 44 environment for the whole multi-agent system contains all the possible domain values 43
 45 for the problem, and at the same time, it also records the *violation numbers* for all the 44
 46 positions. An agent can move within its row, which represents its domain. So far, we have 45

1 introduced three reactive behaviors: *better-move*, *least-move*, and *random-move*. The move 1
2 of an agent will affect the violation numbers of other rows in the environment. 2

3 While formally describing the ERA approach as well as its properties, we have also 3
4 presented several empirical studies that examine how the basic algorithm can be extended 4
5 to effectively find a solution to an n -queen or a coloring problem. Some practical rules for 5
6 behavioral settings have been established following our observations from the experiments. 6

7 We have compared the ERA approach with some of the existing heuristic or distributed 7
8 approaches in order to highlight their main features and limitations. Some features of the 8
9 ERA approach include: 9

- 10 (1) The move of an agent will affect the whole environment. Thus, the interaction among 10
11 agents is indirectly carried out through the medium of their environment. In this sense, 11
12 we may regard that the agents can self-organize themselves in finding a solution. 12
- 13 (2) The performance of this approach in solving CSPs for approximate solutions is 13
14 efficient. It can find a reasonably good solution in just few steps. This property is 14
15 quite useful in situations where a solution is required with a hard deadline. 15
- 16 (3) The ERA approach is also open and flexible: We may set or combine various reactive 16
17 behaviors for each agent, or modify their parameter settings. 17

18 Although we have tested several n -queen problems and coloring problems in our present 18
19 work, we hope that in our future work we will be able to further improve the ERA approach 19
20 by explicitly introducing communication and/or cooperation mechanisms and to discover 20
21 new properties of this approach in solving other types of CSPs. 21
22 22
23 23
24 24

25 Acknowledgements 25

26 We would like to thank the reviewers and associate editor for their valuable comments 26
27 and suggestions, which have led us to examine several related issues and at the same time 27
28 make the paper easier to follow. Our work has been partially supported by a HKBU FRG 28
29 research grant. 29
30 30
31 31
32 32

33 Appendix A 33

34 In what follows, we will present the proofs for the ERA correctness theorems. In the 34
35 proofs, we will adopt symbol ' \rightarrow ' to represent 'such that' and symbol ' \Rightarrow ' to denote 35
36 'imply'. 36
37 37
38 38
39 39

40 A.1. Proof for Theorem 2.1 40

41 **Proof.** There two sections of the algorithm that will have an effect on the violation 41
42 numbers: Initialization (Section 1 in Fig. 13) and the running (Section 2 in Fig. 13). 42
43 Therefore, our proof can be divided into two sections. 43
44 44
45 45

1	(1) Stage 1: Initialization	1
2	(a) Lines 1–5 in Fig. 13.	2
3	Before we place the agents into the environment, because there is no agent in	3
4	it, there will be no assignment. In such a case, there is no constraint for each	4
5	domain value. The precondition $(\forall a_i \in A, i \neq y) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge$	5
6	$(R_t = D_i \times D_y)$ is <i>false</i> , so $(\forall a_i \in A, i \neq y, a_i \text{ one}) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge$	6
7	$(R_t = D_i \times D_y) \rightarrow \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)$ is <i>true</i> .	7
8	(b) Lines 6–8.	8
9	To each agent, the <i>Initialize</i> function is defined in Fig. 9, and the operations on	9
10	violation are in the <i>AddTo</i> function. After adding a_i to position (g, i) , if position	10
11	(x, y) is still a <i>zero-position</i> , that means	11
12	<i>Attack</i> $((g, i), (x, y))$ is <i>false</i>	12
13	$\Rightarrow \nexists t C(R_t)$ between X_i and $X_y \vee$	13
14	$\exists t C(R_t)$ between X_i and $X_y \wedge$	14
15	$\langle e(g, i).value, e(x, y).value \rangle \in C(R_t)$.	15
16	$\Rightarrow (\exists t C(R_t)$ between X_i and $X_y \rightarrow$	16
17	$\langle e(g, i).value, e(x, y).value \rangle \in C(R_t))$.	17
18	$\Rightarrow (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow$	18
19	$\langle e(g, i).value, e(x, y).value \rangle \in C(R_t)$.	19
20		20
21	After all agents are placed onto the environment, if $\exists(x, y) \in \text{environment}$, position	21
22	(x, y) is still a <i>zero-position</i> , that means:	22
23	$(\forall a_i \in A, i \neq y) ((\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow$	23
24	$\langle e(g, i).value, e(x, y).value \rangle \in C(R_t))$	24
25	$\Rightarrow (\forall a_i \in A, i \neq y) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow$	25
26	$\langle e(g, i).value, e(x, y).value \rangle \in C(R_t)$.	26
27		27
28	So Theorem 2.1 is true in this section.	28
29	(2) Stage 2: Running	29
30	We need to prove: If Theorem 2.1 is true before line 15, Theorem 2.1 is still true	30
31	after line 15–16. We will prove it in two steps: (1) Theorem 2.1 is true before line	31
32	15 \rightarrow Theorem 2.1 is true after line 15; (2) Theorem 2.1 is true after line 15 \rightarrow	32
33	Theorem 2.1 is true after line 16.	33
34	[Theorem 2.1 is true.]	34
35	(a) Line 15 <i>RemoveFrom</i> $(a_i.x, i)$.	35
36	After remove a_i from $(a_i.x, i)$, if $\exists(x, y) \in \text{environment}$, (x, y) is still a <i>zero-</i>	36
37	<i>position</i> , that means:	37
38	<i>Attack</i> $((a_i.x, i), (x, y))$ is <i>false</i>	38
39	$\Rightarrow \nexists t C(R_t)$ between X_i and $X_y \vee$	39
40	$\exists t C(R_t)$ between X_i and $X_y \wedge$	40
41	$\langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)$.	41
42	$\Rightarrow (\exists t C(R_t)$ between X_i and $X_y \rightarrow$	42
43	$\langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t))$.	43
44	$\Rightarrow (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow$	44
45	$\langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)$.	45
	[-I]	

1 Because during this process, other agents do not move, and Theorem 2.1 is true 1
 2 before this process. So all these agents 2

$$3 \quad (\forall a_j \in A, j \neq y \wedge j \neq i) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \\
 4 \quad \rightarrow \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t) \text{ is true.} \\
 5$$

6 Combining with [-I], we have: 6

$$7 \quad (\forall a_i \in A, i \neq y) ((\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow \\
 8 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)). \\
 9 \quad \Rightarrow (\forall a_i \in A, i \neq y) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow \\
 10 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t). \\
 11$$

[Theorem 2.1 is true.] 12

13 (b) Line 16 *AddTo*($a_i.x, i$). 13

14 After adding a_i to $(a_i.x, i)$, if $\exists(x, y) \in$ environment, (x, y) is still a *zero-position*, 14
 15 that means: 15

$$16 \quad \text{Attack}((a_i.x, i), (x, y)) \text{ is false} \\
 17 \quad \Rightarrow \nexists t C(R_t) \text{ between } X_i \text{ and } X_y \vee \\
 18 \quad \exists t C(R_t) \text{ between } X_i \text{ and } X_y \wedge \\
 19 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t). \\
 20 \quad \Rightarrow (\exists t C(R_t) \text{ between } X_i \text{ and } X_y \rightarrow \\
 21 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)). \\
 22 \quad \Rightarrow (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow \\
 23 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t). \quad [-I] \\
 24$$

25 Because during the process of *AddTo*($a_i.x, i$), other agents do not move, and 25
 26 Theorem 2.1 is true before this process. So all other agents 26

$$27 \quad (\forall a_j \in A, j \neq y \wedge j \neq i) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \\
 28 \quad \rightarrow \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t) \text{ is true.} \\
 29$$

30 Combining with [-I], we get: 30

$$31 \quad (\forall a_i \in A, i \neq y) ((\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow \\
 32 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t)). \\
 33 \quad \Rightarrow (\forall a_i \in A, i \neq y) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_y) \rightarrow \\
 34 \quad \langle e(a_i.x, i).value, e(x, y).value \rangle \in C(R_t). \\
 35$$

[Theorem 2.1 is true.] 36

37 So during the running process, Theorem 2.1 is true. 37

38 In conclusion, Theorem 2.1 is true. \square 38
 39 39

40 A.2. Proof for Theorem 2.2 40

41 **Proof.** We will prove it based on Definition 1.2 and Theorem 2.1. 41
 42 42

43 (1) $X_i = e(a_i.x, i).value \in D_i$, so S is an n -tuple that $S \in D_1 \times D_2 \times \dots \times D_n$. 43
 44 44
 45 45

1 (2) For each $u \in [1, m]$, $C(R_u) \in \mathbf{C}$, suppose $\exists g, h \in [1, n]$ that $R_u = D_g \times D_h$, 1
 2
 3 $e(a_g.x, g).violation = 0$ 3
 4 $\Rightarrow ((\forall a_i \in A, i \neq g) (\exists t \in [1, m]) (C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_g) \rightarrow$ 4
 5 $\langle e(a_i.x, i).value, e(a_g.x, g).value \rangle \in C(R_t))$, 5
 6 (based on Theorem 2.1) 6
 7 then 7
 8 $(C(R_u) \in \mathbf{C}) \wedge ((\forall a_i \in A, i \neq g) (\exists t \in [1, m])$ 8
 9 $(C(R_t) \in \mathbf{C}) \wedge (R_t = D_i \times D_g) \rightarrow$ 9
 10 $\langle e(a_i.x, i).value, e(a_g.x, g).value \rangle \in C(R_t))$ 10
 11 $\Rightarrow \langle e(a_h.x, h).value, e(a_g.x, g).value \rangle \in C(R_t)$. 11
 12

13 We have: 13
 14 $(C(R_u) \in \mathbf{C}) \wedge e(a_g.x, g).violation = 0 \rightarrow$ 14
 15 $\langle e(a_h.x, h).value, e(a_g.x, g).value \rangle \in C(R_t)$. 15
 16

17 So for 17
 18 $\forall C (R_u) \in C, R_u = D_g \times D_h,$ 18
 19 $\langle \langle e(a_h.x, h).value, e(a_g.x, g).value \rangle \subseteq S \rangle \wedge$ 19
 20 $\langle \langle e(a_h.x, h).value, e(a_g.x, g).value \rangle \in C(R_u) \rangle$ is true. 20
 21

22 So the assignment of $S = \langle X_1, X_2, \dots, X_n \rangle, X_i = e(a_i.x, i).value$, is an exact solution 22
 23 when the system terminates at *condition-1*: $(\forall a_i \in A) e(a_i.x, i).violation = 0$. \square 23
 24

References

25
 26
 27
 28 [1] R. Barták, Heuristics and stochastic algorithms, <http://ktlinux.ms.mff.cuni.cz/~bartak/constraints/stochastic.html>. 28
 29
 30 [2] J.R. Bitner, E. Reingold, Backtrack programming techniques, *Comm. ACM* 18 (11) (1991) 651–656. 30
 31 [3] M. Bruynooghe, Solving combinatorial search problems by intelligent backtracking, *Inform. Process. Lett.* 12 (91) (1981) 36–39. 31
 32 [4] F.C. Chow, J.L. Hennessy, The priority-based coloring method to register allocation, *ACM Trans. Programming Languages and Systems* 12 (4) (1990) 501–536. 32
 33 [5] B. Clement, E. Durfee, Scheduling high-level tasks among cooperative agents, in: *Proceedings of the Third International Conference on Multi-Agent Systems*, 1998. 33
 34 [6] M.C. Cooper, An optimal k -consistency algorithm, *Artificial Intelligence* 41 (1989) 89–95. 34
 35 [7] K. Decker, J. Li, Coordinated hospital patient scheduling, in: *Proceedings of the Third International Conference on Multi-Agent Systems*, 1998. 35
 36 [8] M. Dorigo, G. Di Caro, L.M. Gambardella, Ant algorithms for discrete optimization, *Artificial Life* 5 (2) (1999) 137–172. 36
 37 [9] M. Dorigo, V. Maniezzo, A. Colomi, The ant system: An autocatalytic optimizing process, Technical Report No. 91-016 Revised, Politecnico di Milano, Italy, 1991. 37
 38 [10] J. Ferber, *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*, Addison-Wesley, New York, 1999, pp. 31–35. 38
 39 [11] E.C. Freuder, R.J. Wallace, Partial constraint satisfaction, *Artificial Intelligence* 58 (1–3) (1992) 21–70. 39
 40 [12] E.C. Freuder, Backtrack-free and backtrack-bounded search, in: L. Kanak, V. Kumar (Eds.), *Search in Artificial Intelligence*, Springer, New York, 1988, pp. 343–369. 40
 41
 42
 43
 44
 45

- 1 [13] P. Galinier, J.-K. Hao, Tabu search for maximal constraint satisfaction problems, in: Proceedings of CP-97, 1
Linz, Austria, 1997, pp. 196–208. 2
- 3 [14] A. Gamst, Some lower bounds for a class of frequency assignment problems, IEEE Trans. Vehicular 3
Technology 35 (1) (1986) 8–14. 4
- 5 [15] M.R. Garey, D.S. Johnson, H.C. So, An application of graph coloring to printed circuit testing, IEEE Trans. 5
Circuits Systems CAS-23 (1976) 591–599. 6
- 7 [16] J. Gaschnig, Performance measurement and analysis of certain search algorithm, Ph.D. Thesis, Carnegie- 7
Mellon University, 1979. 8
- 9 [17] J. Gu, Efficient local search for very large-scale satisfiability problem, SIGART Bull. 3 (1992) 8–12. 8
- 10 [18] C.C. Han, C.H. Lee, Comments on Mohr and Henderson's path consistency algorithm, Artificial 9
Intelligence 36 (1988) 125–130. 10
- 11 [19] A. Homaifar, J. Turner, S. Ali, The n -queen problem and genetic algorithms, in: Proceedings of IEEE 10
SOUTHEASTCON-92, 1992, pp. 262–267. 11
- 12 [20] V. Kumar, Algorithm for constraint satisfaction problem: A survey, AI Magazine 13 (1) (1992) 32–44. 12
- 13 [21] F.T. Leighton, A graph coloring algorithm for large scheduling problems, J. Research of the National Bureau 13
of Standards 84 (1979) 489–506. 14
- 15 [22] J. Liu, Autonomous Agents and Multi-Agent Systems: Explorations in Learning, Self-Organization, and 15
Adaptive Computation, World Scientific, 2001. 16
- 17 [23] J. Liu, J. Han, ALIFE: A multi-agent computing paradigm for constraint satisfaction problems, Internat. J. 16
Pattern Recognition and Artificial Intelligence 15 (3) (2001) 475–491. 17
- 18 [24] J. Liu, Y.Y. Tang, Adaptive segmentation with distributed behavior-based agents, IEEE Trans. Pattern 18
Analysis and Machine Intelligence 21 (6) (1999) 544–551. 19
- 20 [25] J. Liu, Y.Y. Tang, Y. Cao, An evolutionary autonomous agents approach to image feature extraction, IEEE 20
Trans. Evolutionary Comput. 1 (2) (1997) 141–158. 21
- 22 [26] A.K. Mackworth, Consistency in networks of relations, Artificial Intelligence 8 (1) (1977) 99–118. 22
- 23 [27] J. Mandziuk, Solving the n -queen problem with a binary Hopfield-type network: Synchronous and 23
asynchronous models, Biological Cybernetics 72 (1995) 439–446. 24
- 24 [28] S. Minton, M.D. Johnston, A.B. Philips, P. Laird, Minimizing conflicts: A heuristic repair method for 24
constraint-satisfaction and scheduling problems, Artificial Intelligence 58 (1992) 161–205. 25
- 26 [29] R. Mohr, T.C. Henderson, Arc and path consistency revisited, Artificial Intelligence 28 (1986) 225–233. 26
- 27 [30] B. Nadel, Some applications of the constraint satisfaction problem, Technical Report CSC-90-008, 27
Computer Science Department, Wayne State University, 1990. 28
- 28 [31] H. Ogawa, Labeled point pattern matching by Delaunay triangulation and maximal cliques, Pattern 28
Recognition 19 (1) (1986) 35–40. 29
- 30 [32] S. Park, E. Durfee, W. Birmingham, Emergent properties of a market based digital library with strategic 30
agents, in: Proceedings of the Third International Conference on Multi-Agent Systems, 1998. 31
- 32 [33] V.G. Paul, Introduction to constraint satisfaction, Lecture 3: Search Order in CSPs and Solution Synthesis, 32
1997. 33
- 34 [34] W. Rosiers, M. Bruynooghe, On the equivalence of constraint satisfaction problem, in: Proceedings of 33
AIMSA-86, North-Holland, Amsterdam, 1989. 34
- 35 [35] F. Rossi, C. Petrie, On the equivalence of constraint satisfaction problems, Technical Report ACT-AI-222- 35
89, MCC, Austin, TX, 1989. 36
- 37 [36] F. Seghrouchni, S. Haddad, A recursive model for distributed planning, in: Proceedings of the Second 37
International Conference on Multi-Agent Systems, 1996. 38
- 38 [37] B. Selman, H. Kautz, B. Cohen, Local search strategies for satisfiability testing, DIMACS Series in Discrete 38
Mathematics and Theoretical Computer Science 26 (1996) 521–532. 39
- 40 [38] B. Selman, H. Kautz, An empirical study of greedy local search for satisfiability testing, in: Proceedings of 40
AAAI-93, Washington, DC, MIT Press, Cambridge, MA, 1993, pp. 46–51. 41
- 42 [39] R. Susic, J. Gu, Efficient local search with conflict minimization: A case study of the n -queen problem, 42
IEEE Trans. Knowledge and Data Engineering 6 (5) (1994) 661–668. 43
- 43 [40] R. Stallman, G.J. Sussman, Forward reasoning and dependency directed backtracking, Artificial Intelli- 43
gence 9 (2) (1977) 135–196. 44
- 44 [41] Swarm Development Group, Swarm simulation system, <http://www.swarm.org/index.html>. 44
- 45 45

1 [42] R. Wallace, Analysis of heuristic methods for partial constraint satisfaction problem, in: Principles and 1
2 Practice of Constraint Programming (CP-1996), 1996, pp. 482–496. 2
3 [43] M. Yokoo, K. Hirayama, Algorithms for distributed constraint satisfaction: A review, Autonomous Agents 3
4 and Multi-Agent Systems 3 (2000) 185–207. 4
5 [44] M. Yokoo, E.H. Durfee, T. Ishida, K. Kuwabara, The distributed constraint satisfaction problem: 5
6 Formalization and algorithms, IEEE Trans. Knowledge and Data Engineering 10 (5) (1998) 673–685. 6
7 [45] M. Yokoo, K. Hirayama, Distributed constraint satisfaction algorithm for complex local problems, in: 7
8 Proceedings of the Third International Conference on Multi-Agent Systems, 1998, pp. 372–379. 7
9 [46] M. Yokoo, Y. Kitamura, Multi-Agent real-time-a* with selection: Introducing competition in cooperative 8
9 search, in: Proceedings of the Second International Conference on Multi-Agent Systems, 1996. 9
10 [47] <http://dimacs.rutgers.edu/Challenges/Seventh/#PC>. 10
11 [48] <http://mat.gsia.cmu.edu/COLOR/instances.html>. 11
12 [49] <http://mscmga.ms.ic.ac.uk/jeb/orlib/colourinfo.html>. 11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

UNCORRECTED PROOF