Continuous Dynamic Constrained Optimisation - The Challenges

Trung Thanh Nguyen, *Member, IEEE* and Xin Yao, *Fellow, IEEE*

Abstract—Many real-world dynamic problems have constraints, and in certain cases not only the objective function changes over time, but also the constraints. However, there is no research in answering the question of whether current algorithms work well on continuous dynamic constrained optimisation problems (DCOPs), nor is there any benchmark problem that reflects the common characteristics of continuous DCOPs. This paper contributes to the task of closing this gap. We will present some investigations on the characteristics that might make DCOPs difficult to solve by some existing dynamic optimisation (DO) and constraint handling (CH) algorithms. We will then introduce a set of benchmark problems with these characteristics and test several representative DO and CH strategies on these problems. The results confirm that DCOPs do have special characteristics that can significantly affect algorithm performance. The results also reveal some interesting observations where the presence or combination of different types of dynamics and constraints can make the problems easier to solve for certain types of algorithms. Based on the analyses of the results, a list of potential requirements that an algorithm should meet to solve DCOPs effectively will be proposed.

Index Terms—Dynamic optimisation, dynamic environments, dynamic constraints, constraint handling, benchmark problems, evolutionary algorithms, performance measures.

I. Introduction

This research aims to answer some open questions about the characteristics, difficulty and solutions of a very common class of problem - dynamic constrained optimisation problems (DCOPs). DCOPs are constrained optimisation problems that have two properties: (a) the objective functions, the constraints, or both, may change over time, and (b) the changes are taken into account in the optimisation process[1]. It is believed that a majority of real-world dynamic problems are DCOPs. However, there are few studies on continuous dynamic constrained optimisation. Existing studies in continuous dynamic optimisation only focus on the unconstrained or domain constraint dynamic cases (which in this paper both are regarded as "unconstrained" problems). Likewise, existing research in

T.T. Nguyen is with The School of Engineering, Technology and Maritime Operations, Liverpool John Moores University, L3 3AF, United Kingdom; email: T.T.Nguyen@ljmu.ac.uk.

X. Yao is with The Centre of Excellence for Research in Computational Intelligence and Applications,(CERCIA), School of Computer Science,University of Birmingham, B15 2TT, United Kingdom; email: X.Yao@cs.bham.ac.uk.

[1]This definition is derived from the (more general) definition of dynamic optimisation problems in [1, section V].

constraint handling only focuses on the stationary constrained problems.

This lack of attention on DCOPs in the continuous domain raises some important research questions: What are the essential characteristics of these types of problems? How well would existing dynamic optimisation and constraint handling strategies perform in dynamic constrained environments if most of them are designed for and tested in either unconstrained dynamic problems or stationary constrained problems only? Why do they work well or not? How can one evaluate if an algorithm works well or not? And finally, what are the requirements for a "good" algorithm that effectively solves these types of problems?

As a large number of real-world applications are dynamic constrained, finding the answers to the questions above is essential. Such answers would help have better understanding about the practical issues of DCOPs and to solve this class of problem more effectively.

The paper is organized as follows. Section II identifies the special characteristics from real-world DCOPs and discuss how the characteristics make this type of problem different from unconstrained dynamic optimisation problems (DOPs). Section III reviews related literature about continuous benchmark problems, identifies the gaps between them and real-world problems and proposes a new set of DCO benchmark problems. Sections IV and V investigate the possibility of solving DCOPs using some representative DO/CH strategies. Experimental analyses about the strengths and weaknesses, and the effect of the mentioned characteristics on each strategy will be undertaken. Based on the experimental results, a list of requirements that algorithms should meet to solve DCOPs effectively is proposed. Finally, Section VI concludes the paper and points out future directions.

II. Characteristics of real-world dynamic constrained problems

The presence of constraints in DCOPs from real-world applications makes them very different from the unconstrained or domain constraint problems considered in academic research. In real-world DCOPs the objective function and constraint functions can be combined in three different types: (a) both the objective function and the constraints are dynamic [2], [3], [4]; (b) only the objective function is dynamic while the constraints are static [5], [6], [7]; and (c) the objective function is static and the constraints are dynamic [8], [9], [10]. In all three types, the presence of infeasible areas can affect how the global optimum moves, or appears after each change. This

leads to some special characteristics which are not found in the unconstrained cases and fixed constrained cases.

First, constraint dynamics can lead to changes in the shape/percentage/structure of the feasible/infeasible areas. Second, objective function dynamics might cause the global optima to switch from one disconnected feasible region to another on problems with disconnected feasible regions, which are very common in real-world constrained problems, especially the scheduling problems [11], [12], [13]. Third, in problems with fixed objective functions and dynamic constraints, the changing infeasible areas might expose new, better global optima without changing the existing optima. One example is the Dynamic 0-1 Knapsack Problem: significantly increasing the capacity of the knapsack can create a new global optimum without changing the existing optimum.

In addition to the three special characteristics above, DCOPs might also have the common characteristics of constrained problems such as global optima in the boundaries of feasible regions, global optima in search boundary, and multiple disconnected feasible regions. These characteristics are widely regarded as being common in real-world applications.

III. A REAL-VALUED BENCHMARK TO SIMULATE DCOPs CHARACTERISTICS

A. Related literature

In the continuous domain, there is no existing continuous benchmark that fully reflects the characteristics of DCOPs listed in Section II. Among existing continuous benchmarks, there are only two recent studies that are related to dynamic constraints. The first study was [14] in which two simple unimodal constrained problems were proposed. These problems take the time variable t as their only time-dependant parameter and hence the dynamic was created by the increase over time of t. These problems have some important disadvantages which prevent them from being used to capture/simulate the mentioned properties of DCOPs: they only capture a simple linear change. In addition, the two problems do not reflect common situations like dynamic objective + fixed constraints or fixed objective + dynamic constraints and other common properites of DCOPs.

The second study was [15]. In that research, a dynamic constrained benchmark problem was proposed by combining an existing "field of cones on a zero plane" dynamic fitness function with four dynamic norm-based constraints with the square/diamond/sphere-like shapes (see Figure 2 in [15]). Although the framework used to generate this benchmark problem is highly configurable, the current single benchmark problem generated by the framework in [15] was designed for a different purpose and hence does not simulate the properties mentioned in Section II. For example, the benchmark problem might not be able to simulate common properties of DCOPs such as optima in boundary; disconnected feasible regions; and moving constraints exposing optima in a controllable way. In addition, there is only one single type of benchmark problem and hence it might be difficult to use the problem to evaluate the performance of algorithms under different situations.

The lack of benchmark problems for DCOPs makes it difficult to (a) evaluate how well existing DO algorithms would work on DCOPs, and (b) design new algorithms specialising in DCOPs. Given that a majority of recent real-world DOPs are DCOPs [16], this can be considered an important gap in DO research.

This gap motivates the authors to develop general-purpose benchmark problems to capture the special characteristics of DCOPs. Some initial results involving five benchmark problems were reported in an earlier study [17]. This paper extends the framework to develop full sets of benchmark problems, which are able to capture all characteristics mentioned in the previous section. Two sets of benchmark problems, one with multimodal, scalable objective functions and one with unimodal objective functions, have been developed for this research. In this paper the benchmark set with unimodal objective functions (many problems in the set still have multiple optima due to the constraints) will be discussed in detail. Detailed descriptions of the multimodal, scalable set can be found in a technical report [18].

B. Generating dynamic constrained benchmark problems

One useful way to create dynamic benchmark problems is to combine existing static benchmark problems with the dynamic rules found in dynamic constrained applications. This can be done by applying the dynamic rules to the parameters of the static problems, as described below.

Given a static function $f_P(x)$ with a set of parameters $P = \{p_1, ...p_k\}$, one can always generalise $f_P(x)$ to its dynamic version $f_{P_t}(x, t)$ by replacing each static parameter $p_i \in P$ with a time-dependent expression $p_i(t)$. The dynamic of the dynamic problem then depends on how $p_i(t)$ varies over time. One can use any type of dynamic rule to represent $p_i(t)$, and hence can create any type of dynamic problem. Details of the concept and a mathematical framework for the idea is described in [18]. Some additional information is provided in [19] (Section 3).

C. A dynamic constrained benchmark set

A set of 18 benchmark problems named G24 was introduced using the new procedure described in the previous subsection. The general form for each problem in the G24 set is as follows:

$$\text{minimise} \quad f(\mathbf{x})$$
$$\text{subject to} \quad g_i(\mathbf{x}) \leq 0, \, g_i(\mathbf{x}) \in G, \, i = 1, .., n$$

where the objective function $f(\mathbf{x})$ can be one of the function forms set out in equation (1), each constraint $g_i(\mathbf{x})$ can be one of the function forms given in equation (2), and G is the set of n constraint functions for that particular benchmark problem. The detailed descriptions of $f(\mathbf{x})$ and $g_i(\mathbf{x})$ for each problem are described in Table I and Table II.

Equation (1) describes the general function forms for the objective functions in the G24 set. Of these function forms, $f^{(2)}$ is used to design the objective function for G24_8a and G24_8b, and $f^{(1)}$ is used to design the objective functions for all other problems. $f^{(1)}$ is modified from a static function

TABLE I
THE OBJECTIVE FUNCTION FORM OF EACH BENCHMARK PROBLEM

Benchmark problem	objective function
G24_8a & G24_8b	$f(x) = f^{(2)}$
All other problems	$f(x) = f^{(1)}$

TABLE II
THE SET OF CONSTRAINT FUNCTION FORMS FOR EACH PROBLEM

Benchmark problem	Set G of constraints
G24_u; G24_uf; G24_2u; G24_8a	$G = \{\emptyset\}$
G24_6a	$G = \left\{ g^{(3)}, g^{(6)} \right\}$
G24_6b	$G = \left\{ g^{(3)} \right\}$
G24_6c	$G = \left\{ g^{(3)}, g^{(4)} \right\}$
G24_6d	$G = \left\{ g^{(5)}, g^{(6)} \right\}$
All other problems	$G = \left\{ g^{(1)}, g^{(2)} \right\}$

proposed in [20] and $f^{(2)}$ is a newly designed function.

$$f^{(1)} = -(X_1 + X_2) \qquad (1)$$

$$f^{(2)} = -3\exp\left(-\sqrt{\sqrt{(X_1)^2 + (X_2)^2}}\right)$$

where $X_i = X_i(x_i, t) = p_i(t)(x_i + q_i(t)); 0 \le x_1 \le 3; 0 \le x_2 \le 4$ with $p_i(t)$ and $q_i(t)$ $(i = 1, 2)$ as the dynamic parameters, which determine how the dynamic objective function of each benchmark problem changes over time.

Equation (2) describes the general function forms for the constraint functions in the G24 set. Of these function forms, $g^{(1)}$ and $g^{(2)}$ were modified from two static functions proposed in [20] and $g^{(3)}, g^{(4)}, g^{(5)}$ and $g^{(6)}$ are newly designed functions.

$$g^{(1)} = -2Y_1^4 + 8Y_1^3 - 8Y_1^2 + Y_2 - 2 \qquad (2)$$

$$g^{(2)} = -4Y_1^4 + 32Y_1^3 - 88Y_1^2 + 96Y_1 + Y_2 - 36$$

$$g^{(3)} = 2Y_1 + 3Y_2 - 9$$

$$g^{(4)} = \begin{cases} -1 & \text{if } (0 \le Y_1 \le 1) \text{ or } (2 \le Y_1 \le 3) \\ 1 & \text{otherwise} \end{cases}$$

$$g^{(5)} = \begin{cases} -1 & \text{if } (0 \le Y_1 \le 0.5) \text{ or } (2 \le Y_1 \le 2.5) \\ 1 & \text{otherwise} \end{cases}$$

$$g^{(6)} = \begin{cases} -1 & \text{if } [(0 \le Y_1 \le 1) \text{ and } (2 \le Y_2 \le 3)] \\ & \text{or } (2 \le Y_1 \le 3) \\ 1 & \text{otherwise} \end{cases}$$

where $Y_i = Y_i(x, t) = r_i(t)(x + s_i(t)); 0 \le x_1 \le 3; 0 \le x_2 \le 4$ with $r_i(t)$ and $s_i(t)$ $(i = 1, 2)$ as the dynamic parameters, which determine how the constraint functions of each benchmark problem change over time.

Each benchmark problem may have a different mathematical expression for $p_i(t)$, $q_i(t)$, $r_i(t)$ and $s_i(t)$. Note that although many benchmark problems share the same general function form in equation (1), their individual expressions for $p_i(t)$ and $q_i(t)$ make their actual dynamic objective functions very different. Similarly, the individual expressions for $r_i(t)$ and $s_i(t)$ make each actual dynamic constraint functions very different although they may share the same function form. The individual expressions of $p_i(t)$, $q_i(t)$, $r_i(t)$, and $s_i(t)$ for each benchmark function are described in Table III.

Two guidelines were used to design the test problems: (a) problems should simulate the common properties of DCOPs as

mentioned in Section II and (b) there should always be a pair of problems for each characteristic. The two problems in each pair should be almost identical except that one has a particular characteristic (e.g. fixed constraints) and the other does not. By comparing the performance of an algorithm on one problem with its performance on the other problem in the pair, it is possible to analyse whether the considered characteristic has any effect on the tested algorithm and to what extent that effect is significant.

Based on the two guidelines above, 18 different test problems were created (Table III). Each test problem is able to capture one or several of the mentioned characteristics of DCOPs, as shown in Table IV. In addition, the problems and their relationships are carefully designed so that they can be arranged in 21 pairs (Table V), of which each pair is a different test case to test a single characteristic of DCOPs (the two problems in each pair are almost identical except that one has a special characteristic and the other does not).

IV. CHALLENGES OF APPLYING CURRENT DYNAMIC OPTIMISATION STRATEGIES DIRECTLY TO SOLVING DCOPS

A. Analysing the performance of some common dynamic optimisation strategies in solving DCOPs

The strategies being considered are (1) introducing diversity, (2) maintaining diversity and (3) tracking the previous optima. These three are among the four most commonly used strategies (the other strategy is memory-based) to solve DOPs. The diversity-introducing strategy was proposed based on the assumption that by the time a change occurs in the environment, an evolutionary algorithm (EA) might have already converged to a specific area and hence would lose its ability to deal with changes in other areas of the search space. Consequently, it is necessary to increase the diversity level in the population, either by increasing the mutation rate or re-initialising/re-locating the individuals. This strategy was introduced years ago [21] but is still extensively used [22] [23].

The diversity-introducing strategy requires that changes must be visible to the algorithm. To avoid this disadvantage, the diversity-maintaining strategy was introduced so that population diversity can be maintained without explicitly detecting changes [24]. This strategy is still the main strategy in many recent approaches [25] [26].

The third strategy, tracking-previous-optima, is used where the optima might only slightly change. The region surrounding the current optima is monitored to detect changes and "track" the movement of these optima. Similar to the two strategies above, the tracking strategy has also been used for years [21] and it has always been one of the main strategies for solving DOPs. Recently this strategy has been combined with the diversity maintaining/introducing strategy to achieve better performance. Typical examples are the multi-population/multi-swarm approaches, where multiple sub-populations are used to maintain diversity and each sub-population/sub-swarm focuses on tracking one single optimum [26] [27].

B. Chosen algorithms and experimental settings

1) Chosen algorithms: Two commonly used algorithms: *triggered hyper-mutation GA* (HyperM [21]) and *random-*

TABLE III
DYNAMIC PARAMETERS FOR ALL TEST PROBLEMS IN THE BENCHMARK
SET G24. EACH DYNAMIC PARAMETER IS A TIME-DEPENDANT
RULE/FUNCTION WHICH GOVERNS THE WAY THE PROBLEMS CHANGE

Prob	Parameter settings
G24_u	$p_1(t) = \sin\left(k\pi t + \frac{\pi}{2}\right); p_2(t) = 1; q_i(t) = 0$
G24_1	$p_2(t) = r_i(t) = 1; \; q_i(t) = s_i(t) = 0$
	$p_1(t) = \sin\left(k\pi t + \frac{\pi}{2}\right)$
G24_f	$p_i(t) = r_i(t) = 1; \; q_i(t) = s_i(t) = 0$
G24_uf	$p_i(t) = 1; \; q_i(t) = 1$
G24_2	if $(t \bmod 2 = 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\begin{cases} p_2(t-1) \text{ if } t>0 \\ p_2(0)=0 \text{ if } t=0 \end{cases} \end{cases}$
	if $(t \bmod 2 \neq 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\sin\left(\frac{k\pi(t-1)}{2}+\frac{\pi}{2}\right) \end{cases}$
	$q_i(t) = s_i(t) = 0; r_i(t) = 1$
G24_2u	if $(t \bmod 2 = 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\begin{cases} p_2(t-1) \text{ if } t>0 \\ p_2(0)=0 \text{ if } t=0 \end{cases} \end{cases}$
	if $(t \bmod 2 \neq 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\sin\left(\frac{k\pi(t-1)}{2}+\frac{\pi}{2}\right) \end{cases}$
	$q_i(t) = 0$
G24_3	$p_i(t) = r_i(t) = 1; q_i(t) = s_1(t) = 0$
	$s_2(t) = 2 + t.\frac{x_2\max - x_2\min}{S}$
G24_3b	$p_1(t) = \sin\left(k\pi t + \frac{\pi}{2}\right); \;\; p_2(t) = 1$
	$q_i(t) = s_1(t) = 0; \; r_i(t) = 1;$
	$s_2(t) = 2 + t.\frac{x_2\max - x_2\min}{S}$
G24_3f	$p_i(t) = r_i(t) = 1; q_i(t) = s_1(t) = 0; s_2(t) = 2$
G24_4	$p_2(t) = r_i(t) = 1; \; q_i(t) = s_1(t) = 0$
	$p_1(t) = \sin\left(k\pi t + \frac{\pi}{2}\right); s_2(t) = t.\frac{x_2\max - x_2\min}{S}$
G24_5	if $(t \bmod 2 = 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\begin{cases} p_2(t-1) \text{ if } t>0 \\ p_2(0) \text{ if } t=0 \end{cases} \end{cases}$
	if $(t \bmod 2 \neq 0)$ $\begin{cases} p_1(t)=\sin\left(\frac{k\pi t}{2}+\frac{\pi}{2}\right) \\ p_2(t)=\sin\left(\frac{k\pi(t-1)}{2}+\frac{\pi}{2}\right) \end{cases}$
	$q_i(t) = s_1(t) = 0; \; r_i(t) = 1;$
	$s_2(t) = t.\frac{x_2\max - x_2\min}{S}$
G24_6a/b/c/d	$p_1(t) = \sin\left(\pi t + \frac{\pi}{2}\right); p_2(t) = 1;$
	$q_i(t) = s_i(t) = 0; r_i(t) = 1$
G24_7	$p_i(t) = r_i(t) = 1; \; q_i(t) = s_1(t) = 0;$
	$s_2(t) = t.\frac{x_2\max - x_2\min}{S}$
G24_8a	$p_i(t) = -1; q_1(t) = -(c_1 + r_a.\cos(k\pi t))$
	$q_2(t) = -(c_2 + r_a.\sin(k\pi t));$
G24_8b	$p_i(t) = -1; q_1(t) = -(c_1 + r_a.\cos(k\pi t))$
	$q_2(t) = -(c_2 + r_a.\sin(k\pi t)); r_i(t) = 1; \; s_i(t) = 0$
k	k determines the severity of function changes.
	$k=1 \sim$ large; $k=0.5 \sim$ medium; $k=0.25 \sim$ small
S	S determines the severity of constraint changes
	$S=10 \sim$ large; $S=20 \sim$ medium; $S=50 \sim$ small
c_1, c_2, r_a	$c_1 = 1.470561702; c_2 = 3.442094786232;$
(G24_8a/b	$r_a = 0.858958496$.
only)	
i	i is the variable index, $i = 1, 2$

TABLE IV
PROPERTIES OF EACH TEST PROBLEM IN THE G24 BENCHMARK SET

Problem	ObjFunc	Constr	DFR	SwO	bNAO	OICB	OISB	Path
G24_u	Dynamic	NoC	1	No	No	No	Yes	N/A
G24_1	Dynamic	Fixed	2	Yes	No	Yes	No	N/A
G24_f	Fixed	Fixed	2	No	No	Yes	No	N/A
G24_uf	Fixed	NoC	1	No	No	No	Yes	N/A
G24_2*	Dynamic	Fixed	2	Yes	No	Yes&No	Yes&No	N/A
G24_2u	Dynamic	NoC	1	No	No	No	Yes	N/A
G24_3	Fixed	Dynamic	2-3	No	Yes	Yes	No	N/A
G24_3b	Dynamic	Dynamic	2-3	Yes	No	Yes	No	N/A
G24_3f	Fixed	Fixed	1	No	No	Yes	No	N/A
G24_4	Dynamic	Dynamic	2-3	Yes	No	Yes	No	N/A
G24_5*	Dynamic	Dynamic	2-3	Yes	No	Yes&No	Yes&No	N/A
G24_6a	Dynamic	Fixed	2	Yes	No	No	Yes	Hard
G24_6b	Dynamic	NoC	1	No	No	No	Yes	N/A
G24_6c	Dynamic	Fixed	2	Yes	No	No	Yes	Easy
G24_6d	Dynamic	Fixed	2	Yes	No	No	Yes	Hard
G24_7	Fixed	Dynamic	2	No	No	Yes	No	N/A
G24_8a	Dynamic	NoC	1	No	No	No	No	N/A
G24_8b	Dynamic	Fixed	2	Yes	No	Yes	No	N/A

DFR	number of Disconnected Feasible Regions
SwO	Switched global Optimum between disconnected regions
bNAO	better Newly Appear Optimum without changing existing ones
OICB	global Optimum is In the Constraint Boundary
OISB	global Optimum is In the Search Boundary
Path	Indicate if it is easy or difficult to use mutation to travel between feasible regions
Dynamic	The function is dynamic
Fixed	There is no change
NoC	There is no constraint
*	In some change periods, the landscape either is a plateau or contains infinite number of optima and all optima (including the existing optimum) lie in a line parallel to one of the axes

random solutions, the algorithm is able to maintain diversity throughout the search process to cope with dynamics. RIGA represents the "maintaining diversity" strategy in DO.

One reason to choose these algorithms for the test is that their strategies are still commonly used in most current state-of-the-art DO algorithms. Another reason is the strategies in these algorithms are very simple and straightforward, making it easy to test and analyse their behaviour. In addition, because these two algorithms are very well studied, using them would help in comparing new experimental data with existing results. Finally, because both algorithms are developed from a basic GA (actually the only difference between HyperM/RIGA and a basic GA is the mutation strategy), it would be easier to compare/analyse their performance. The performance of HyperM and RIGA was also compared with a basic GA to see if they work well on the tested problems.[2]

2) Parameter settings: Table VI shows the detailed parameter settings for HyperM, RIGA and GA. All algorithms use real-valued representations. The algorithms were tested on 18 benchmark problems described in Section III. To create a fair testing environment, the algorithms were tested in a wide range of dynamic settings (different values of population size, severity of change and frequency of change) with five levels: *small, medium small, medium, medium large, large*.

The evolutionary parameters of all tested algorithms were set to similar values or the best known values if possible.

immigrant GA (RIGA [24]) were chosen to evaluate the performance of the three strategies mentioned above in DCOPs. HyperM is basically a simple GA with an adaptive mechanism to switch from a low mutation rate (standard-mutation-rate) to a high mutation rate (hyper-mutation-rate, to increase diversity) and vice versa depending on whether or not there is a degradation of the best solution in the population. It represents the "introducing diversity" and "tracking previous optima" strategies in DO.

RIGA is another derivative of a basic GA. After the normal mutation step, a fraction of the population is replaced with randomly generated individuals. This fraction is determined by a random-immigrant-rate (also named replacement rate). By continuously replacing a part of the population with

---

[2]Note that to save space, some tables/figures in this section include not only GA/RIGA/HyperM but also another algorithm: GA+Repair. This algorithm will be introduced in the later sections. This section only focuses on the data relating to GA, RIGA and HyperM.

TABLE V
THE 21 TEST CASES (PAIRS) TO BE USED IN THIS PAPER.

Static problems: Unconstrained vs Fixed constraints		
1 G24_uf (fF, noC)	vs	G24_f (fF, fC)
Fixed objectives vs Dynamic objectives		
2 G24_uf (fF, noC)	vs	G24_u (dF, noC)
3 G24_f (fF, fC, OICB)	vs	G24_1 (dF, fC, OICB)
4 G24_f (fF, fC, OICB)	vs	G24_2 (dF, fC, ONICB)
Dynamic objectives: Unconstrained vs Fixed constraints		
5 G24_u (dF, noC)	vs	G24_1 (dF, fC, OICB)
6 G24_2u (dF, noC)	vs	G24_2 (dF, fC, ONICB)
Fixed constraints vs Dynamic constraints		
7 G24_1 (dF, fC, OICB)	vs	G24_4 (dF, dC, OICB)
8 G24_2 (dF, fC, ONICB)	vs	G24_5 (dF, dC, ONICB)
9 G24_f (fF, fC)	vs	G24_7 (fF, dC, NNAO)
10 G24_3f (fF, fC)	vs	G24_3 (fF, dC, NAO)
No constraint vs Dynamic constraints		
11 G24_u (dF, noC)	vs	G24_4 (dF, dC, OICB)
12 G24_2u (dF, noC)	vs	G24_5 (dF, dC, ONICB)
13 G24_uf (fF, noC)	vs	G24_7 (fF, dC)
Moving constraints expose better optima vs not expose optima		
14 G24_3f (fF, fC)	vs	G24_3 (fF, dC, NAO)
15 G24_3 (fF, dC, NAO)	vs	G24_3b (dF, dC, NAO)
Connected feasible regions vs Disconnected feasible regions		
16 G24_6b (1R)	vs	G24_6a (2DR, hard)
17 G24_6b (1R)	vs	G24_6d (2DR, hard)
18 G24_6c (2DR, easy)	vs	G24_6d (2DR, hard)
Optima in constraint boundary vs Optima NOT in constr boundary		
19 G24_1 (dF, fC, OICB)	vs	G24_2 (dF, fC, ONICB)
20 G24_4 (dF, dC, OICB)	vs	G24_5 (dF, dC, ONICB)
21 G24_8b (dF, fC, OICB)	vs	G24_8a (dF, noC, ONISB)

dF	dynamic objective func	fF	fixed objective function
dC	dynamic constraints	fC	fixed constraints
OICB	optima in constraint bound	ONICB	opt. not in constraint bound
OISB	optima in search bound	ONISB	optima not in search bound
NAO	better newly appear optima	NNAO	No better newly appear opt
2DR	2 Disconn. feasible regions	1R	One single feasible region
Easy	easy for mutation to travel between disconn. regions	Hard	less easy to travel among regions
noC	unconstrained problem	SwO	Switched optimum between disconnected regions

The base mutation rate of the algorithms is 0.15, which is the average value of the best mutation rates commonly used for GA-based algorithms in various existing studies on continuous DO, which are 0.1 ([28] [29]) and 0.2 ([27] [30]). For HyperM and RIGA, the best *hyper-mutation-rate* and *random-immigrant-rate* parameter values observed in the original papers [21] [24] were used. The same implementations as described in [21] and [24] were used to reproduce these two algorithms. A crossover rate of 0.1 was chosen for all algorithms because, according to the analysis in [31], this value was one of the few settings where all tested algorithms perform well on this benchmark set.

A further study of the effect of different values of the base mutation rates, hyper-mutation-rates, random-immigrant rates and crossover rates on algorithm performance was also carried out. Detailed experimental results and discussion for this analysis can be found in [31] where it was found that the overall behaviours of the algorithms are not different from those using the default/best known settings, except for the followings: (i) When the base mutation rate is very low (≤ 0.01), the performance of GA and HyperM drop significantly; (ii) generally to work well in the tested DCOPs, algorithms need to use high base mutation rates. The range of best mutation rates

TABLE VI
TEST SETTINGS FOR ALL ALGORITHMS USED IN THE PAPER.

All algorithms (exceptions below)	Pop size (pop_size)	5, 15, 25 (medium), 50, 100
	Elitism	Elitism & non-elitism if applicable
	Selection method	Non-linear ranking as in [33]
	Mutation method	Uniform, $P = 0.15$.
	Crossover method	Arithmetic, $P = 0.1$.
HyperM	Triggered mutate	Uniform, $P = 0.5$ as in [21].
RIGA	Rand-immig. rate	$P = 0.3$ as in [24].
GA+Repair	Search pop size	pop_size $\times (4/5)$
	Reference pop size	pop_size $\times (1/5)$
	Replacement rate	0 (default is 0.25 as in [33]).
Benchmark problem settings	Number of runs	50
	Number of changes	$5/k$ (see below)
	Change frequency	250, 500, 1000 (med), 2000, 4000 evaluations
	ObjFunc severity k	0.25 (small), 0.5 (med), 1.0 (large)
	Constr. severity S	10 (small), 20 (medium), 50 (large)

is 0.3-0.8. (iii) Algorithms like RIGA and HyperM also need high random-immigrant/hyper-mutation rates to solve DCOPs. The best results are usually achieved with the rates of 0.6-0.8; (iv) The suitable range of crossover rate is 0.1-1.0.

3) Constraint handling: It is necessary to integrate existing DO algorithms with a CH mechanism to use these algorithms for solving DCOPs. That CH mechanism should not interfere with the original DO strategies so that it is possible to correctly evaluate whether the original DO strategies would still be effective in solving DCOPs. To satisfy this requirement, the penalty function approach in [32] was chosen because it is the simplest way to apply existing unconstrained DO algorithms directly to solving DCOPs without changing the algorithms. Also this penalty method can be effective in solving difficult numerical problems without requiring users to choose any penalty factor or other parameter [32].

4) Performance measures: For measuring the performance of the algorithms in this particular experiment, an existing measure: the *modified offline error*[27] was modified. The measure is calculated as the average over, at every evaluation, the error of the best solution found since the last change of the environment.

Because the measure above is designed for unconstrained environments, it is necessary to modify it to evaluate algorithm performance in constrained environments: At every generation, instead of considering the best errors/fitness values of *any* solutions regardless of feasibility as implemented in the original measure, only the best fitness values / best errors of *feasible* solutions at each generation are considered. If in any generation there is no feasible solution, the measure takes the *worst possible value* that a feasible solution can have for that particular generation. This measure is called the *modified offline error for DCOPs*, or *offline error* for short.

$$E_{MO} = \frac{1}{num\_of\_gen} \sum_{j=1}^{num\_of\_gen} e_{MO}(j) \qquad (3)$$

where $e_{MO}(j)$ is the best *feasible* error since the last change at the generation j.

Five new measures were also proposed to analyse why a particular algorithm might work well on a particular problem. The first two measures are the *recovery rate* (RR) and the *absolute recovery rate* (ARR) to analyse the convergence

behaviour of algorithms in dynamic environments. The RR measure is used to analyse *how quickly an algorithm recovers from an environmental change and starts converging to a new solution before the next change occurs*. The new solution is not necessarily the global optimum.

$$RR = \frac{1}{m} \sum_{i=1}^{m} \frac{\sum_{j=1}^{p(i)} \left[f_{best}(i,j) - f_{best}(i,1) \right]}{p(i) \left[f_{best}(i,p(i)) - f_{best}(i,1) \right]} \quad (4)$$

where $f_{best}(i,j)$ is the fitness value of the best feasible solution since the last change found by the tested algorithm until the jth generation of the change period i, m is the number of changes and $p(i), i = 1 : m$ is the number of generations at each change period i. The RR score would be 1 in the best case where the algorithm is able to recover and converge to a solution immediately after a change, and would be close to zero in case the algorithm is unable to recover from the change at all [3].

The RR measure only indicates if the considered algorithm converges to a solution and if it converges quickly. It does not indicate whether the converged solution is the global optimum. For example, RR can still be 1 if the algorithm does nothing but keep re-evaluating the same solution. Because of that, another measure is needed: the *absolute recovery rate* (ARR). This measure is very similar to the RR but is used to analyse *how quick it is for an algorithm to start converging to the global optimum before the next change occurs*:

$$ARR = \frac{1}{m} \sum_{i=1}^{m} \frac{\sum_{j=1}^{p(i)} \left[f_{best}(i,j) - f_{best}(i,1) \right]}{p(i) \left[f^*(i) - f_{best}(i,1) \right]} \quad (5)$$

where $f_{best}(i,j), i, j, m, p(i)$ are the same as in Eq. 4 and $f^*(i)$ is the global optimal value of the search space at the ith change. The ARR score would be 1 in the best case when the algorithm is able to recover and converge to the global optimum immediately after a change, and would be zero in case the algorithm is unable to recover from the change at all. Note that the score of ARR should always be less than or equal to that of RR. In the ideal case (converged to global optimum), ARR should be equal to RR.[4]

The RR and ARR measures can be used together to indicate if an algorithm is able to converge to the global optimum within the given time frame between changes and if so how quickly it takes to converge. The *RR-ARR diagram* in Figure 1 shows some analysis guidelines.

A third measure, *percentage of selected infeasible individuals*, is proposed to analyse algorithm ability to balance exploiting feasible regions and exploring infeasible regions in DCOPs. This measure finds the percent of infeasible individuals selected for the next generation. The average (over all tested generations) is then compared with the percentage of infeasible areas in the search space. If the considered algorithm is able to accept *infeasible* diversified individuals in the same way as it accepts *feasible* diversified individuals (and hence to maintain diversity effectively), the two percentage values should be equal.

---

[3]Note that RR will never be equal to zero because there is at least one generation where $f_{best}(i,j) = f_{best}(i,p(i))$

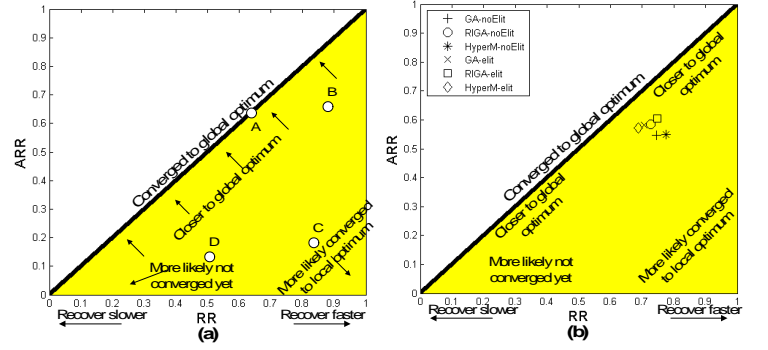[4]Note that to use the measure ARR it is necessary to know the global optimum value at each change period.

Fig. 1. **Diagram (a)** provides a guideline for analysing the convergence behaviour/recovery speed of an algorithm given its RR/ARR scores. These scores can be represented as the x and y coordinations of a point on the diagonal thick line or inside the shaded area. The position of the point represents the behaviour of the corresponding algorithm. The closer the point is to the right, the faster the algorithm was in recovering and re-converging, and vice versa. In addition, if the point lies on the thick diagonal line (where $RR = ARR$) like point A, the algorithm has been able to recover from the change and converged to the new global optimum. Otherwise, if the point lies inside the shaded area, the algorithm either has converged to a local solution (e.g. point C); or has not been converged yet (e.g. point D - recover slowly; and point B - recover quickly). **Diagram (b)** shows the mapping of the RR/ARR scores of GA, RIGA, and HyperM to the RR-ARR diagram.

To analyse the behaviour of algorithms using triggered-mutation mechanisms such as HyperM, a fourth measure: *triggered-time count*, which counts the number of times the hyper-mutation-rate is triggered by the algorithm, and a fifth measure: *detected-change count*, which counts the number of triggers actually associated with a change, are also proposed. For HyperM, triggers associated with a change are those that are invoked by the algorithm within ν generations after a change, where ν is the maximum number of generations (five in this implementation) needed for HyperM to detect a drop in performance. These two measures indicate how many times an algorithm triggers its hyper-mutation; whether each trigger time corresponds to a new change; and if there is any change that goes undetected during the search process.

Note that the five measures above are all needed for our analysis because they are used to investigate different aspects of the algorithms. Furthermore, all of the measures used here are specifically designed for dynamic problems. This creates a problem for the experiments in this paper because in the G24 benchmark set there are not only dynamic problems but also stationary problems. To overcome this issue, in this study stationary problems are considered a special type of dynamic problem which still have "changes" with the same change frequency as other dynamic problems. However, in stationary problems the "changes" do not alter the search space.

C. Experimental results and analyses

The full *offline-error* results of the tested algorithms on all 18 benchmark problems for all test scenarios are presented in the tables in [34]. These data were further analysed from different perspectives to achieve a better understanding of how existing DO strategies work in DCOPs and how each characteristic of DCOPs would affect the performance of existing DO algorithms. First of all, the average performance

of the tested algorithms on each major group of problems under different parameter settings and dynamic ranges were summarised to have an overall picture of algorithm behaviour on different types of problems (see Figure 2). Then the effect of each problem characteristic on each algorithm was analysed in 21 test cases (each case is a pair of almost identical problems, one with a particular characteristic and one without) as shown in Table V of Section III (see test results in Figures 3 and 4). For each particular algorithm, some further analyses were also carried out using the five newly proposed measures mentioned above. Details of these analyses will be described in the next subsections. Only the summarised results are presented in Figure 2 with different settings (small / medium / large). For other detailed figures and tables, the results will only be presented in the default settings (all parameters and dynamic range are set to medium). For detailed results in other settings, readers are referred to [34].

A statistical *t*-test with a significance level of 0.05 was done to evaluate the level of siginfance of the possible impacts that each characteristic of DCOPs can have on the performance of the tested algorithms[5]. The summarised results of this statistical test can be found in Figures 3 and 4.

The experimental results show some interesting, and in some cases, surprising findings.

1) The impact of different dynamic ranges on algorithm performance: The summarised results in groups of problems (Figure 2) show that (i) generally the behaviour of algorithms and their relative strengths/weaknesses in comparison with other algorithms still remain roughly the same when the dynamic settings change; and (ii) as expected in most cases algorithms' performance decrease when the conditions become more difficult (magnitude of change becomes larger; change frequency becomes higher; population size becomes much smaller). Among the variations in dynamic settings, it seems that the variations in frequency of change affect algorithms' performance the most, followed by variations in magnitude of changes. Variations in population size have the least impact on algorithm performance.

2) The effect of elitism on algorithm performance: The summarised results in groups of problems (Figure 2) and the pair-wise comparisons in Figure 3 and Figure 4 reveal an interesting effect of elitism on both unconstrained and constrained dynamic cases: the elitism versions of GA/RIGA/HyperM perform better than their non-elitism counterparts in most tested problems. The reason for this effect (with evidence shown in the next paragraph) is that elitism helps algorithms with diversity-maintaining strategies to converge faster. This effect is independent of the combined CH techniques.

Two measures proposed in Section IV-B4: *recovery rate* (RR) and *absolute recovery rate* (ARR) were used to study the inefficiency of GA/RIGA/HyperM in the non-elitism case. The scores of the algorithms on these measures are given in Figure 1b. The figure shows that none of the algorithms are close to the optimum line, meaning there are problems/ change periods where the algorithms were unable to converge to the global

---

[5]*t*-test is considered robust under the conditions of this experiment[35, ch. 37].
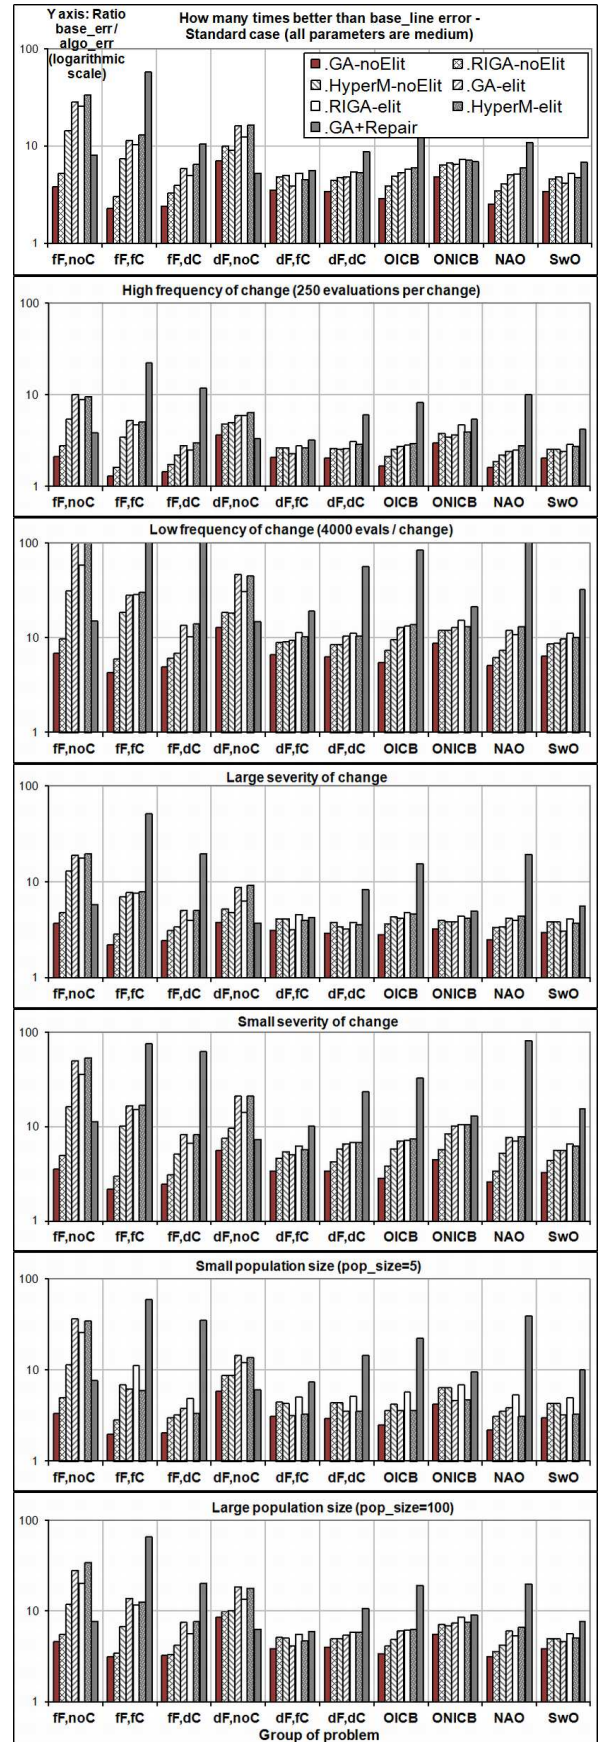

Fig. 2. Algorithm performance in groups of problem. Performance (vertical axis in logarithmic scale) is evaluated by calculating the ratio between the *base line* (worst error among all scenarios) and the error of each algorithm in each problem to see how many times their performance is better (smaller) than the *base line*. Explanations for abbreviations can be found in Table V.
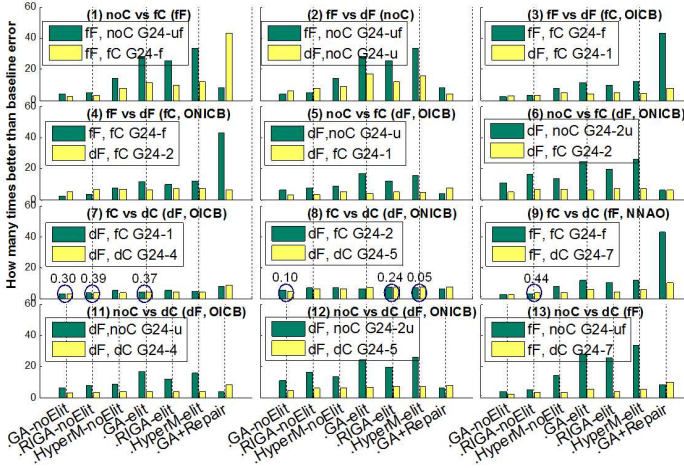
Fig. 3. The effect of twelve different problem characteristics on algorithm performance (medium case). Performance (vertical axis) is evaluated based on the ratio between the base line error (described in Figure 2) and algorithm errors. Each subplot represents algorithm performance (pair of adjacent bars) in a pair of almost identical problems (one has a special characteristic and the other does not). The larger the difference between the bar heights, the greater the impact of the corresponding DCOP characteristic on performance. Subplots' title represent the test case numbers (in brackets) followed by an abbreviated description. Explanations for the abbreviations are in the last rows of Table V. Pairs where the impact of a characteristic on an algorithm is *not significant* (according to a *t*-test with significance level of 0.05) are circled and in such cases the *t*-test scores are also given to highlight the level of insignificance.


Fig. 4. The effect of the other eight different problem properties on algorithm performance (medium case). Instructions to read this figure can be found in Figure 3. All eight characteristics have statistically significant impacts on the algorithms, and hence there is no bar with circles.

optimum. In addition, for RIGA, its elitism version is closer to the top-right corner while its non-elitism version is closer to the bottom-left corner, meaning that non-elitism makes RIGA converge slower/less accurately. Finally, for GA/HyperM, their elitism versions are closer to the global optimum while their non-elitism versions are closer to the bottom-right corner, meaning that the non-elitism versions of GA/HyperM are more suceptible to premature convergence. The results hence show that the high diversity maintained by the random-immigrant rate in RIGA and the high mutation rate in GA/HyperM comes with a trade-off: the convergence speed is affected. In such a

TABLE VII
AVERAGE *percentage of selected infeasible individuals* OVER 18 PROBLEMS.
THE LAST ROW SHOWS THE AVERAGE *percentage of infeasible areas*.

Algorithms	Percent of infeasible solutions
.GA-elit	23.0%
.RIGA-elit	37.6%
.HyperM-elit	26.4%
.GA-noElit	46.3%
.RIGA-noElit	49.1%
.HyperM-noElit	45.3%
Percentage of infeasible areas	**60.8%**

situation, elitism can be used to speed up the convergence process. Elite members can guide the population to exploit the good regions faster while still maintaining diversity.

3) Effect of infeasible areas on maintaining/introducing diversity: Another interesting observation is that the presence of constraints makes the performance of diversity-maintaining/introducing strategies less effective when used in combination with the tested penalty functions. This behaviour can be seen in Figure 2 where the performance of all algorithms in the unconstrained dynamic case (dF+noC) is significantly better than their performance in all dynamic constrained cases (dF+fC, fF+dC, dF+dC). This behaviour can also be seen in the more accurate pair-wise comparisons in Figure 3 and Figure 4: for each pair of problems in which one has constraints and the other does not, GA, RIGA and HyperM always perform worse on the problem with constraints (see pairs 1, 5, 6, 11, 12, 13 in Figure 3 and pair 21 in Figure 4).

The reason for this inefficiency is the use of tested penalty functions prevents diversity-maintaining/introducing mechanisms from working effectively. In solving unconstrained dynamic problems, all diversified individuals generated by the diversity maintaining/introducing strategies are useful because they contribute to either (1) detecting newly appearing optima or (2) finding the new place of the moving optima. In DCOPs, however, there are two difficulties that prevent diversified individuals that are infeasible from being useful in existing DO strategies. One difficulty is many diversified but infeasible individuals might not be selected for the next generation population because they are penalised with lower fitness values by the penalty functions. Consequently, these diversified individuals cannot be used for maintaining diversity unless they are re-introduced again in the next generation. To demonstrate this drawback, the previously proposed measure *percentage of selected infeasible individuals* was used. As can be seen in Table VII, in the elitism case the percentage of infeasible solutions in the population (23 - 37.6%) is much smaller than the percentage of infeasible areas over the total search space (60.8%). This means only a few of the diversified, infeasible solutions are retained and hence the algorithms are not able to maintain diversity in the infeasible regions.[6]

The second difficulty is that, even if a diversified but infeasible individual is selected for the next generation, it

---

[6]Non-elitism algorithms are able to retain more infeasible individuals, of which some might be diversified solutions. However, as shown in Subsection IV-C2, in the non-elitism case this higher percentage of infeasible individuals comes with a trade-off of slower/less accurate convergence, which leads to the generally poorer performance.

might no longer have its true fitness value. Consequently, environmental changes might not be accurately detected or tracked.

4) Effect of switching global optima (between disconnected feasible regions) on strategies that use penalty functions: The results show existing DO methods become less effective when they are used in combination with the tested penalty functions to solve a special class of DCOPs: problems with disconnected feasible regions where the global optimum switches from one region to another whenever a change occurs. In addition, the more separated the disconnected regions are, the more difficult it is for algorithms using penalty functions to solve.

The reason for this difficulty is it is necessary to have a path through the infeasible areas that separate the disconnected regions to track the moving optimum. This path might not be available if penalty functions are used because penalties make it unlikely infeasible individuals are accepted. Obviously the larger the infeasible areas between disconnected regions, the harder it is to establish the path using penalty methods.

Three test cases (pairs of almost identical problems) 16, 17, 18 in Table V were used to verify the statement above. In all three test cases the objective functions are the same and the global optimum switches between two locations whenever a change occurs. However, each case represents a different dynamic situation. Case 16 tests the situation where in one problem of the pair (G24_6b) there is a feasible path connecting the two locations and in the other problem (G24_6a) the path is infeasible, i.e., there is an infeasible area separating two feasible regions. Case 17 is the same as case 16 except that the infeasible area separating two feasible regions has a different shape. Case 18 tests a different situation where in one problem (G24_6c) the infeasible area separating the two feasible regions is small whereas in the other problem (G24_6d) this infeasible area is large.

The experimental results in these three test cases (pairs 16, 17, 18 in Figure 4) confirm the hypotheses stated in the beginning of this subsection. In cases 16 and 17, the performance of the tested algorithms did decrease when the path between the two regions is infeasible. In case 18, the larger the infeasible area separating the two regions, the worse the performance of the tested algorithms.

5) Effect of moving infeasible areas on strategies that track the previous optima: Algorithms relying on tracking previous global optimum such as HyperM might become less effective when the moving constraints expose new, better optima without changing the existing optima. The reason is HyperM cannot to detect changes in such DCOPs and hence might not be able to trigger its hyper-mutation rate. With the currently chosen base mutation of 0.15, HyperM is still able to produce good results because the mutation is high enough for the algorithm to maintain diversity. However, in a previous study [17], when a much smaller base mutation rate was used, HyperM becomes significantly worse compared to other algorithms in solving problems like G24_3.

To illustrate this drawback, the newly proposed measures *triggered-time count* and *detected-change count* were used to analyse how the triggered-hypermutation mechanism works on problem G24_3. As can be seen in Table VIII, Hy-

| Algorithms | G24_3 (NAO+fF) | | | | G24_3b (NAO+dF) | | | |
| | Trigger Count | | Detected Change Count | | Trigger count | | Detected Change Count | |
	Value	stdDev	Value	stdDev	Value	stdDev	Value	stdDev
HyperM-noElit	188.70	8.40	1.74	0.78	199.83	5.88	11.00	0.00
HyperM-elit	0.00	0.00	0.00	0.00	30.43	0.57	11.00	0.00

NAO - Newly Appearing Optimum

fF / dF - fixed / dynamic objective Function

perM either was not able to trigger its hyper-mutation rate to deal with changes (elitism case, *triggered-time count*=0 & *detected-change count*=0) or was not able to trigger its hyper-mutation rate correctly when a change occurs (non-elitism case, *triggered-time count*∼188.7 & *detected-change count*∼1.74). It is worth noting in the non-elitism case, most of the trigger times are caused by the selection process because in non-elitism selection the best solution in the population is not always selected for the next generation.

Table VIII also shows that in problem G24_3b, which is almost identical to G24_3 except it has its existing optima changed, HyperM was able to detect changes and hence trigger its hyper-mutation timely whenever a change occurs. It shows HyperM only becomes less effective where environmental changes do not change the value of existing optima.

D. Possible suggestions to improve current dynamic optimisation strategies in solving DCOPs

The experimental results suggest some directions for addressing the drawbacks listed in the previous subsections:

(i) Based on the observation that elitism is useful for diversity-maintaining strategies in solving DCOPs, it might be useful to develop algorithms that support both elitism and diversity maintaining mechanisms.

(ii) Given that methods like HyperM are not able to detect changes because they mainly use change detectors (the best solution in case of HyperM) in the feasible regions, it might be useful to use change detectors in both regions and infeasible regions.

(iii) Because experimental results show that tracking the existing optima might not be effective in certain cases of DCOPs, it might be useful to track the moving feasible regions instead. Because after a change in DCOPs the global optimum always either moves along with the feasible areas or appears in a new feasible area, an algorithm able to track feasible areas would have higher chance of tracking the actual global optimum.

V. CHALLENGES OF SOME CONSTRAINT HANDLING STRATEGIES IN SOLVING DCOPs

A. Difficulties in handling dynamics

The most obvious reason for the difficulties in applying existing CH strategies to solving DCOPs is these strategies are not designed to handle environmental dynamics. One might then question whether these difficulties can be overcome by combining existing CH strategies with existing DO strategies.

Unfortunately, as will be shown below, not all difficulties can be resolved by combining existing CH strategies with existing DO strategies. In addition, this combination might also bring some new challenges due to the conflict of the optimisation goals of the two types of strategies. These are the challenges in maintaining diversity, introducing diversity, and detecting changes based on performance drop.

1) Impacts on maintaining/introducing diversity: As already discussed, one of the important strategies in DO is to maintain/introduce diversity in the whole search space to detect changes and to find newly-appearing/moving optima. However, diversity might no longer be maintained this way when combined with some CH techniques.

In many CH techniques, the original space is specifically transformed so algorithms only focus on certain areas instead of the whole original space. In such cases, even if a diversity-introducing strategy such as HyperM is used to generate individuals in the whole search space, diversified individuals generated in the unfocused areas might be neglected by the algorithms and hence do not contribute to maintaining diversity. Typical examples of CH strategies that adopt this search space transformation approach are penalty methods where the constrained search space is transformed to an unconstrained search space with penalised fitness values. Another example is some approaches use special representations/operators. In these approaches, the algorithms might be restricted to searching only in the feasible regions, in a transformed feasible search space, or in the boundaries of feasible regions. Detailed reviews/references for representative penalty approaches and special representations/operators approaches can be found in [36], [37].

In some other CH techniques, individuals are selected not exclusively based on their actual fitness values but also on some special specifications. For example, in Stochastic Ranking [38] infeasible individuals might have a better chance of being accepted based on the given stochastic parameter. A contrary example can be found in Simple Multimembered ES [39] where infeasible solutions are less likely to be accepted even if they have higher fitness values than the feasible ones. Another example is in a CH multi-objective approach [40] where individuals are ranked not entirely based on their original fitness but also on the number of violated constraints. In CH techniques like these, diversified individuals generated by DO strategies might not be selected in the same way as they were originally designed for, i.e., the number of infeasible diversified individuals might become too large or too small. The way a diversity maintaining strategy works might not be the same as in the unconstrained case.

Experimental evidence for the inefficiency mentioned above has already been shown in Section IV-C, where the diversity-maintaining/introducing strategies become less effective when combined with the tested penalty methods.

In [31], it was shown that current state-of-the-art in CH such as SRES [38], [41] and SMES [39] become much less effective in DCOPs and could not maintain enough diversity to deal with the dynamics in DCOPs.

2) Impacts on change detection: Another possible difficulty of combining CH strategies with DO strategies is the use of some existing CH techniques might make change detection based on performance drop, a common DO technique, less effective. As already mentioned in Section IV-B, algorithms like HyperM assume that during the search process, if there is a degradation in the fitness values of the best solution found in each generation, there might be a change in the search space. However, when DO algorithms are combined with some CH techniques to solve DCOPs, such degradation in best fitness values might no longer be caused by an actual change in the search space. Instead, the degradation might be caused either by an increase in penalty values or by the elimination of the current good solutions from the population.

One example can be found in some CH techniques such as dynamic penalty or adaptive penalty [42], [43], [44], where the degradation of (modified) fitness values is not caused by environmental changes but by the increase over time of the penalty values. The consequence of this dynamic/adaptive scheme is that if the detector solutions used by the change detection method are infeasible or become infeasible, over time their fitness value will decrease.[7]

In some other CH techniques which use ranking-based methods [38], [39], [40], during the selection process the current better solutions might be dropped in favour other solutions, which might have worse fitness values but are more useful for the CH process. In these situations there might also be a drop in the values of the best solutions at each generation.

The drop in fitness values of the detector solutions in both cases above might be incorrectly considered by DO strategies like HyperM to be a change in the environment and this might consequently trigger the DO strategies to react inappropriately.

B. Difficulties in handling constraints (empirical evidence shown in Section V-D)

The difficulties of applying some existing CH strategies to solving DCOPs are also caused by that their CH ability become less effective. This is due to two reasons.

1) The issue of outdated information: In DOPs, after a change, all existing information that an algorithm has acquired or has been given about the problem might become outdated and consequently make the algorithm less effective. For example, in algorithms using strictly feasible reference individuals like Genocop III [45], [33], after a change some reference individuals may become infeasible. Similarly, in some "decoder" methods the reference lists for ordinal representations (e.g. the ordered lists of cities (TSP [46])/ ordered lists of knapsack items (KSP [47]) / order lists of tasks (scheduling [48])) might no longer be in order after a change because the cities/items/tasks have changed their values. Another example can be found in dynamic/adaptive penalty methods (e.g. [43] [44]) where the penalty parameters learnt by the methods might no longer be suitable because the balance between feasible and infeasible solutions has changed.

---

[7]Of course, in penalty methods, if change detections are made on the original fitness instead of on the penalised fitness, the increase of penalty values will not have any impact on detecting changes. However, in this case, change detection might suffer from another problem: changes in constraint functions will go undetected unless additional improvements are made to detect constraint changes explicitly.

2) The issue of outdated strategy: The CH strategies themselves can also be outdated when solving DCOPs. This might occur when the CH strategies have problem-dependent parameters, whose values might be tailored to work best in only one (class of) stationary environment, to solve a DCOP. In such cases, if the parameters are fine-tuned for the problem before change, the algorithm might only work well until a change occurs. Typical examples are penalty methods with pre-defined penalty factors and/or other pre-defined parameters that control how the penalty is defined. Other examples are some combinatorial repair methods, methods with special operators, or decoder methods. Detailed reviews are in [37] and [49]).

Strategy-being-outdated might also occur with many adaptive CH strategies that are not problem-dependent because these strategies rely on some specific assumptions that are only true in stationary problems.

Typical examples are self-adaptive fitness formulation [50] and stochastic ranking [38]. The general approach of these strategies is to balance feasibility/infeasibility based on the performance of the current population, assuming that the population always reflects a "memory" of information about the search space and the convergence process. This assumption is not true in dynamic environments. When a change occurs, the search space might change its shape and consequently the "memory" of the population no longer reflects the property of the new search space but only a small area where the population currently is. This disadvantage has been observed in [31] for the case of the state-of-the-art SRES.

Another type of CH strategies relying on outdated assumptions are dynamic/adaptive methods that use the running time value (e.g. the number of generations so far) to balance feasibility and infeasibility. CH strategies of this type [42], [43], [51], [52], [39] assume that the population will eventually converge to the good regions and hence they handle constraints by increasingly rejecting more infeasible solutions when time goes by, or by reducing the mutation step size when time goes by, to increase the convergence speed to good regions. In DCOPs, because after a change good feasible regions might no longer be good or feasible, if the CH strategy still imposes its previous balancing mechanism to increase convergence speed, the algorithm could end up converging to the wrong place and will not be able to track the moving optima. This disadvantage has been experimentally confirmed in [31] for the case of the state-of-the-art SMES.

C. Possible suggestions to improve current constraint handling strategies in solving DCOPs

The discussions in the two previous subsections show that, to handle constraints effectively in DCOPs, a CH strategy might need to satisfy the requirements below:

1) Make sure that the goal of CH does not conflict with the goal of DO. Particularly:
   a) Allow diversified individuals to be distributed in the whole search space.
   b) Do not reject diversified individuals even if they do not contribute to CH.

c) Pay special attention whenever changes are detected by monitoring the fitness values of current individuals (it is necessary to check to see if a drop in performance is really caused by an environmental change).

2) Make sure that the algorithm is updated whenever a change occurs. Particularly:
   a) Problem knowledge needs to be updated.
   b) The CH strategy might also need to be updated whenever a change occurs.

An algorithm needs to handle both environmental dynamics and constraints effectively to work well in DCOPs. This means that a "good" algorithm for DCOPs needs to satisfy not only the requirements for CH above but also the four requirements for DO identified in Section IV-D.

D. Experimental analyses

An experimental analysis was carried out to test the performance of the *repair method*, a representative CH strategy, on the G24 benchmark set. The purpose is to answer three questions: (i) what is the usefulness of the repair method in solving DCOPs; (ii) whether the hypothesis about the difficulties of DCOPs toward CH strategies, as mentioned in Section V-B, is true; (iii) if the hypothesis is true, would these difficulties affect the performance of CH strategies (in particular the repair method) in solving DCOPs. These results would help gain more understanding about how to design better algorithms to solve DCOPs.

1) Chosen constraint handling technique for the analysis: For this analysis the *repair method* [33] was chosen because it is representative, simple, easy to implement, problem-independent and is designed specifically for the continuous domain.

Repair-based methods, however, also have one disadvantage: they may require a considerable number of feasibility checks to find a feasible individual. As a result, repair-based methods might not be suitable for solving problems with very expensive constraint functions and problems with very small feasible areas.

2) Repair algorithms & the method in Genocop III [33]:
   a) General ideas: The idea of repairing is, if it is possible to map (repair) an infeasible solution to a feasible solution, then instead of searching the best feasible solution directly, it might be possible to look for an individual that can potentially produce the best repaired solution. The better the repaired solution, the higher the fitness value of an individual. In certain cases, the feasible solution created by the repair process can also be used to replace some of the search individuals.

Generally, a repair process can be described in three steps:
1) If a newly created individual $\mathbf{s}$ (can be feasible or infeasible) needs repair, use a heuristic $repair()$ to repair $\mathbf{s}$, mapping $\mathbf{s}$ to a new, feasible individual $\mathbf{z}$.
2) The objective value $f(\mathbf{z})$ of $\mathbf{z}$ is used as input to calculate the fitness value of $\mathbf{s}$, $\mathrm{eval}(\mathbf{s}) = h(f(\mathbf{z}))$ where h is the mapping from objective values to fitness.
3) If the repair approach is Lamarckian, replace one or more search individuals by $\mathbf{z}$

In the *repair method* [45], [33] chosen for this experiment, the $repair$ () heuristic is as follows:

1) The population is divided into two sub-populations: a search population S containing normally-evolving individuals, which can be fully feasible or only linearly feasible, and a reference population R containing only fully feasible individuals.

2) During the search process, while each individual **r** in R is evaluated using their objective function as usual, each individual **s** in S is considered to be repaired based on an individual from R. Details of the repair routine can be found in Algorithm 1.

It is important to note there are two possible variants of deciding whether a search individual **s** needs to be repaired in Genocop III (step 2 above). In the first variant [45], a search individual **s** is repaired *only* if **s** is infeasible. In the second and latest variant [33], the implementation shows that search individuals are repaired regardless of their feasibility.

In all experiments in this paper, the second variant was implemented. From now on, unless stated otherwise the term *repair method* will be used to refer to the continuous-based repair approach proposed in [33].

b) Feasibility/infeasibility balancing strategy and problem knowledge in the repair method: The repair method and other repair approaches have the ability to adaptively balance feasibility and infeasibility. This balance is achieved by accepting both infeasible and feasible individuals, provided that they can produce good repaired solutions and by updating the fitness values of search individuals with those of the mapped, feasible solutions. This way the repair method ensures that infeasible solutions are accepted and they cannot have better fitness values than the best feasible solution available.

The strategy above needs certain problem information, which is provided by the reference population R and the search population S. R is an essential source of information to direct the algorithm toward promising feasible regions (during the repair process (*Repair* routine, Algorithm 1), newly repaired solutions are always generated in the directions toward reference individuals). R also provides the balancing strategy with information about the best feasible solution available (via their fitness values) so that the strategy can make sure that no infeasible individual can have better fitness values than this best feasible solution.

The search population S is also an essential source of problem information. It helps indicate which point in the search space would lead to potentially promising feasible regions (via repair). In the selection phase the balancing strategy then uses this information to select those individuals that would potentially lead to the most promising regions.

c) How can the characteristics of DCOPs affect the repair method?: The repair method suffers from the problem of outdated information, which in turn makes the feasibility/infeasibility balancing strategy outdated.

The first type of information might become outdated when a change occurs is the fitness values of search individuals. Because the fitness of a search individual is always based on the objective value of the corresponding mapped feasible solution, it is assumed that the search population always offers

a "memory" of good areas in the search space and directions toward these good areas. The higher the fitness value of an individual, the better the feasible region achieved by repairing this individual.

In a dynamic environment, the memory, or fitness values of search individuals, can become outdated right after a change if the objective values of the corresponding repaired solutions change. Particularly, the high fitness values of existing individuals might no longer lead to good repaired solutions and vice versa. Worse, search individuals with high-but-outdated fitnessh values might incorrectly bias the selection process, which makes the search process less effective.

The second type of information that might become outdated when a change occurs is the set of reference individuals that are used to repair all other search individuals. The key assumption that all reference individuals are feasible and are the best in the population is only true in stationary environments. In dynamic environments, after a change, some existing reference individuals might no longer remain the best in the population or might even become infeasible. These outdated reference individuals not only violate the assumption named above but might also wrongly bias the search and drive more individuals *away* from the good regions, making the search process less effective.

In the following experiments an analysis was made to see if the above hypotheses are correct and how significant their effects are.

3) Experimental settings:

a) Tested algorithms: In this experiment, the repair method was integrated with a basic GA. The integrated version is called GA+Repair and is described in Algorithm 2. This integration makes it possible to analyse the strengths and weaknesses of the repair strategy because the only difference between GA and GA+Repair is the repair operator and hence any difference in performance would be caused by the repair operator[8]. In addition, because all other tested strategies are integrated with a basic GA, it is natural to integrate the repair method with the GA[9] to compare it with these strategies.

Even though the GA+Repair is a simplified version of Genocop III, both algorithms have very similar behaviours when solving different groups of DCOPs. This similarity suggests the result tested with GA+Repair can be generalised to other approaches that use the repair method. For detailed results of Genocop III's performance in the G24 benchmark set and a comparison of its performance with other existing and new algorithms, readers are referred to the study in [31].

b) Parameter settings: The tested algorithms use the same parameter settings as the previously tested GA, RIGA, and HyperM except that the population now is divided into a

---

[8] It is more difficult to analyse the effect of the repair strategy in the original Genocop III because this algorithm implements multiple CH strategies (beside the repair operator, there are ten other specialised operators to handle linear constraints)

[9] It should be noted that while Genocop III allows 25% of the repaired individuals to replace individuals in the population (Lamarckian evolution), in GA+Repair none of the repaired individuals is used to replace the original individuals (Baldwinian evolution). The reason is that in [31] it was found that Lamarckian evolution does not significantly increase/decrease the performance of Genocop III in solving DCOPs.

Algorithm 1 *routine* Repair(Indiv **s**)

1) Randomly pick an individual $\mathbf{r} \in R$
2) Generate individual $\mathbf{z}$ in the segment between $\mathbf{s}$ and $\mathbf{r}$
   a) $a = U(0,1)$
   b) $\mathbf{z} = a.\mathbf{s} + (1-a).\mathbf{r}$
   c) While $\mathbf{z}$ is infeasible, back to step 2a
   d) If a feasible z is not found after 100 trials, $z = r$ and $\mathrm{eval}(\mathbf{z}) = \mathrm{eval}(\mathbf{r})$
3) a) Evaluate $\mathbf{z}$
   b) If ($f(\mathbf{z})$ better than $f(\mathbf{r})$): $\mathbf{r} = \mathbf{z}; \mathrm{eval}(\mathbf{r}) = f(\mathbf{z})$
   c) Update the fitness value of $\mathbf{s}$: $\mathrm{eval}(\mathbf{s}) = f(\mathbf{z})$
4) Return the individual $\mathbf{s}$

search population and a reference population (see Table VI), as implemented in the original Genocop III [33].

c) Performance measures: Three different measures were used. The first measure, which is the modified version of the *off-line error* measure (see Section IV-B4), was used to evaluate/compare the general performance of the GA+Repair. Similar to the previous experiment, using this measure the average performance of GA+Repair was also summarised in each major group of problems (see results in Figure 2) and the effect of each problem characteristic on GA+Repair was analysed in 21 test cases shown in Table V of Section III (see results in Figure 3 and Figure 4).

The second and third measures were specifically proposed for this experiment. The second measure, named *feasible reference individuals*, was used to analyse the behaviour of the *repair method* when some reference individuals become outdated due to environmental changes (see Figure 5). The third measure, named *feasible individuals in each disconnected region*, was used to analyse the ability of repair methods to balance feasibility and infeasibility on problems with optima switching between disconnected feasible regions (see Figure 6). Details of these two measures will be described later.

4) The impact of outdated information/strategy on the performance of the repair method :

a) Overall observation of performance in groups of problems: In the group of *stationary constrained problems* (fF, fC), the results in Figure 2 show that, as expected, a specialised CH technique such as the repair method in GA+Repair performs significantly better than methods not designed for handling constraints like the existing DO algorithms. In *stationary unconstrained group* (fF, noC), also, as expected, the repair method in GA+Repair is no longer particularly useful. Figure 2 shows that GA+Repair performs worse than all other methods in dynamic, unconstrained problems (dF, noC).

In the groups of *DCOPs* (fF+dC, dF+fC, dF+dC), things are different. As can be seen in Figure 2, in DCOPs the difference between GA+Repair and GA is no longer as significant as it is in the stationary constrained case, meaning that the performance of GA+Repair significantly decreases. This happens in all three cases of DCOPs where only the constraints are dynamic (fF, dC), where only the objective functions are dynamic (dF, fC) and where both constraints and objective functions are dynamic (dF, dC).

Algorithm 2 GA+Repair

Note: It is assumed that the problem is maximisation
1) *Initialise*:
   a) Randomly initialise m individuals in search pop S
   b) Initialise n individuals in the reference population R
      i) Randomly generate points until a feasible $\mathbf{r}$ is found
      ii) Update fitness: $\mathrm{eval}(\mathbf{r}) = f(\mathbf{r})$ & add r to R
      iii) Repeat step 1(b)i until n individuals are found
2) *Search*: For $i = 1 : m$
   a) $p_1 = U(0,1); p_2 = U(0,1)$
   b) *Crossover*: **If** $(p_1 < P_{Xover})$
      i) Use nonlinear ranking selection to choose a pair of parents from S
      ii) Crossover an offspring $\mathbf{s}$ from the chosen parents
      iii) Evaluate $\mathbf{s}$ and repair $\mathbf{s}$ using $Repair(\mathbf{s})$
      iv) Use nonlinear ranking selection to replace one of the worst individuals in S by $\mathbf{s}$
   c) *Mutation*: **If** $(p_2 < P_{Mutate})$
      i) Use nonlinear ranking-selection to choose a parent from S
      ii) Mutate an offspring $\mathbf{s}$ from the chosen parent
      iii) Evaluate $\mathbf{s}$ and repair $\mathbf{s}$ using $Repair(\mathbf{s})$
      iv) Use nonlinear ranking selection to replace one of the worst individuals in S by $\mathbf{s}$
   d) *Otherwise*: **If** $(p_1 \geq P_{Xover})$ **and** $(p_2 \geq P_{Mutate})$
      i) Use nonlinear ranking-selection to choose an individual $\mathbf{s}$ from S
      ii) If $\mathbf{s}$ has not been evaluated since last generation, evaluate $\mathbf{s}$
      iii) Repair $\mathbf{s}$ using the routine $Repair(\mathbf{s})$
      iv) Using nonlinear ranking selection to replace one of the worst individuals in S by $\mathbf{s}$
3) Evolve the reference population after each 100 evaluations: For $i = 1 : n$
   a) *Crossover*: **If** $(U(0,1) < P_{Xover})$
      i) Use nonlinear ranking-selection to choose a pair of parents from R
      ii) Crossover an offspring $\mathbf{r}$ from the parents
      iii) **If** $\mathbf{r}$ is feasible
         A) Evaluate $\mathbf{r}$ and $\mathbf{x}$,the better of the two parents
         B) If $f(\mathbf{r})$ better than $f(\mathbf{x})$ then $\mathbf{x} = \mathbf{r}$ and fitness value $\mathrm{eval}(\mathbf{x}) = f(\mathbf{r})$
   b) *Mutation*: **If** $(U(0,1) < P_{Mutation})$
      i) Nonlinear ranking-selection to choose a parent $\mathbf{x}$ from R
      ii) Mutate an offspring $\mathbf{r}$ from $\mathbf{x}$
      iii) **If** $\mathbf{r}$ is feasible
         A) Evaluate $\mathbf{r}$ and $\mathbf{x}$
         B) If $f(\mathbf{r})$ better than $f(\mathbf{x})$ then $\mathbf{x} = \mathbf{r}$ and fitness value $\mathrm{eval}(\mathbf{x}) = f(\mathbf{r})$
4) Return to step 2

Details of the *impact of dynamic objective functions* on the repair method can be seen in pair-wise comparisons in pairs 9 and 14 of Figure 3 where GA+Repair is tested in pairs of almost identical constrained problems except that one has a fixed and the other has a dynamic objective function. As can be seen in these plots, the performance of GA+Repair significantly decreased in case the objective function is dynamic. The difference in performance of GA+Repair between the two problems of each pair is significantly larger than that of GA and existing DO algorithms, meaning that the presence of dynamic objective functions has a much greater impact on the repair method than on GA and existing DO methods.

Details of the *impact of dynamic constraints* on the repair method can be seen in the pair-wise comparisons in plot i of Figure 3 and plot a of Figure 4 (pairs of almost identical problems except that one has fixed and the other has dynamic constraints). Similar to the previous case, the results also show that the performance of GA+Repair is significantly decreased in case the constraints are dynamic and the presence of dynamic constraints has a much greater impact on the repair method than on existing DO methods.

However, although the presence of environmental dynamics does significantly decrease the performance of GA+Repair, in Figure 2 it is interesting to see that the algorithm still performs better than existing DO algorithms in DCOPs (only that the difference become significantly smaller compared to the static constrained case). This shows the repair method has some characteristics that make it promising for solving DCOPs.

Another interesting, and somewhat counter-intuitive observation in our experiment is the presence of constraints do not make the problems more difficult to solve by GA+Repair. Instead, the presence of constraints always helps GA+Repair work better. Evidence can be found in the pair-wise comparison in pairs 1, 5, 6, 11, 12, 13 of Figure 3 and in pair 21 of Figure 4 where GA+Repair always performs better on the problem with constraints than on the problem without constraints. The experiment also shows that GA+Repair performs better where there is an infeasible barrier separating two feasible regions. Moreover, the larger the barrier, the better the performance of GA+Repair (see pairs 17, 18 in Figure 4).

The experimental results confirm dynamics do have a significant effect on the performance of the repair method. Following is a further analysis to investigate if this effect is indeed caused by the outdated problem information (reference individuals and search individuals) and by the outdated strategy as suspected by our hypothesis.

b) Analyse the behaviours of outdated reference individuals: As mentioned earlier, outdated information might be caused by reference individuals having their objective values changed or even become infeasible after a change. The previously proposed measure *feasible reference individuals* was used to test if the algorithm is able to update the reference individuals properly. If the algorithm is able to update the reference individuals properly, it should be able to maintain a reference population of all feasible individuals during the search process.

The most suitable environments to test this behaviour of the repair method are DCOPs with dynamic constraints where


Fig. 5. This figure shows how GA+Repair maintains feasible reference individuals in problems with moving infeasible regions. The total number of reference individuals is five. The plot in the figures shows, among these five reference individuals, how many are actually feasible during the search process.

after each change the previous best feasible solutions are hidden by the moving infeasible region. They are G24_4, G24_5 (dF, dC) and G24_7 (fF, dC). As discussed earlier (see Figure 2), in both groups the performance of GA+Repair decreases significantly compared to the case where the constraints are fixed (fF,fC). In these problems, if one or more reference individuals do become infeasible, there should be a drop in the total number of feasible reference individuals and this is one of the reasons making the repair method less effective.

The *plot of feasible reference individuals* of GA+Repair is given in Figure 5. The figure shows, in all cases the original repair method was not able to keep all reference individuals feasible during the search. The number of feasible reference individuals drops to a very low level when a change occurs and most of the time the number of feasible reference individuals is much lower than five.

The results confirm the hypothesis that after a change, the population of reference individuals has become outdated due to the moving infeasible regions.

c) Analyse the behaviours of the outdated balancing strategy: In Section V-D2, it was suspected that individuals being outdated can also have a negative impact on the balancing strategy, which balances feasibility and infeasibility of the repair method. To test if the algorithm is still able to balance feasibility/infeasibility properly in dynamic environments, the proposed measure: *feasible individuals in each disconnected region* was used to monitor the number of feasible individuals in each disconnected feasible region and the ratio of feasibility/infeasibility. The balancing mechanism should be able to manage a good distribution of individuals so that the better feasible regions should have more feasible individuals if it works well in the DCOP case.

The most suitable environments to test this behaviour are DCOPs with two disconnected feasible regions where the global optimum keeps switching from one region to another after each change or after some consecutive changes. They are G24_1, G24_2, G24_3b, G24_4, G24_5, G24_6a, G24_6c, G24_6d, and G24_8b. All these problems belong to the group *SwO* in Figure 2, where the performance of existing CH algorithms significantly decreases compared to the stationary constrained case (fF, fC). On such SwO problems, if the balancing mechanisms work well, at each change period the algorithm should be able to focus most feasible individuals

on the region where the global optimum is currently in while still maintaining the same ratio of feasibility/infeasibility for diversity purpose.

The *plot of feasible individuals in each disconnected region* for GA+Repair is given in Figure 6. The figure shows that in all cases except G24_3b, the repair method was not able to focus most feasible individuals on the region where the global optimum is currently in. Instead, the majority of feasible individuals still remained in one single region (region 2). The number of individuals in the other region (region 1) remained low regardless of where the global optimum is. This is due to that, although the global optimum has switched to region 1, many individuals in region 2 were not updated and still have the outdated fitness values which might be even higher than the new global optimum value. These outdated individuals incorrectly attract a large number of individuals to the old feasible region. These results show that, due to its outdated strategy, the algorithm was not able to follow the switching optimum well[10].
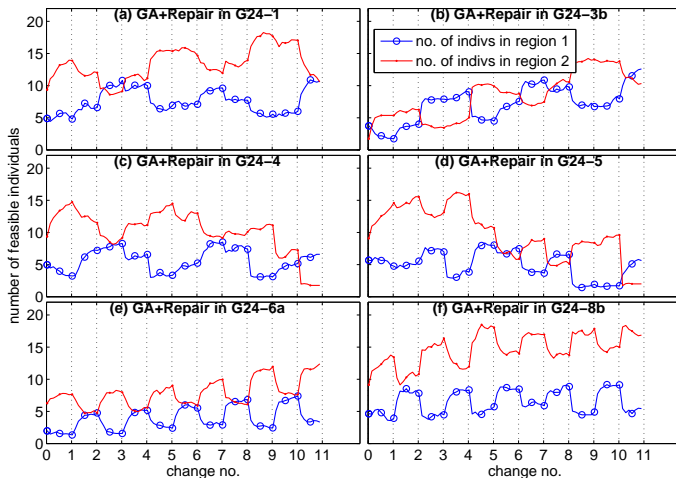

Fig. 6. This figure shows how the balance strategy of GA+Repair distributes its feasible individuals in disconnected feasible regions. The problems tested in this figure are those with global optima switching between two disconnected feasible regions. [12]

VI. Conclusion & future research

In this paper we have identified some special and not well studied characteristics of DCOPs that might cause significant challenges to existing DO and CH strategies. Although these characteristics are common in real-world applications, in the continuous domain they have not been considered in most existing DO studies and they have not been captured in existing continuous DO benchmark problems.

A set of dynamic constrained benchmark problems for simulating the characteristics of DCOPs have been proposed to help close this gap. To help researchers assess algorithm

[10]Note that in the G24 set, individuals being outdated might not always be totally harmful because the changes in many problems are cyclic and hence the outdated individuals might actually play the role of memory elements to recall the previous good solutions. However, it is not clear how beneficial such memory elements could be, because the experiments show that GA+Repair still becomes less effective in the presence of environmental dynamics.

performance in DCOPs, seven new measures have been proposed to evaluate the performance / analyse the behaviours of algorithms on dynamic unconstrained/constrained problems and one existing measure has also been modified to make it usable in DCOPs.

Using the newly proposed benchmark problems and measures, some literature reviews and detailed experimental analyses have been carried out to investigate the strengths and weaknesses of existing DO strategies (GA/RIGA/HyperM) and CH strategies (repair methods) in solving DCOPs. The experimental analyses reveal some interesting findings about the ability of existing algorithms in solving DCOPs. These findings can be categorised as follows.

First, three interesting findings about the performance of existing DO strategies in DCOPs have been identified: (a) the use of elitism might have a positive impact on the performance of existing diversity-maintaining strategies and but might have a negative impact on the performance of diversity-introducing strategies if they are not used with diversity-maintaining strategies; (b) the presence of infeasible areas has a negative impact on the performance of diversity-introducing/maintaining strategies; and (c) the presence of switching optima (between disconnected regions) has a negative impact on the performance of DO strategies if they are combined with penalty functions.

Second, it has been found that even if CH strategies can be combined with DO strategies, there might be two types of difficulties in applying existing CH strategies to solving DCOPs: (a) difficulties in handling dynamics, in particularly maintaining diversity and detecting changes; and (b) difficulties in handling constraints, which are caused by outdated CH strategies and problem-knowledge.

Third, some counter-intuitive behaviours were observed: the presence of constraints and dynamics in DCOPs might not always make the problems harder to solve. For example, the presence of constraints helps algorithms using the repair method like GA+Repair work better in the tested problems.

Finally, based on the findings about the strengths and weaknesses of some existing DO and CH strategies, a list of possible requirements that DO and CH algorithms should meet to solve DCOPs effectively have been suggested. This list of requirements can be used as a guideline to design new algorithms to solve DCOPs in future research.

The results in this paper raise some open questions for future research. One direction is to develop new algorithms specialised in solving DCOPs based on our suggested list of requirements. Another direction is to investigate the impact of DCOPs' characteristics on other state-of-the-art CH and DO strategies. We are interested in investigating the performance of other adaptive feasibility/infeasibility balancing strategies, e.g. [38], [53] in DCOPs. We also plan to study the situations where the presence of constraints and dynamics would make it easier for certain classes of algorithms to solve DCOPs.

The research has some limitations to be improved in future research: memory-based approaches have not been considered in our analysis; the algorithms and methods used to analyse representative DO and CH strategies are very basic to keep the analysis at a manageable level; the analysis has been

tested only in a benchmark problems with unimodal objective functions; the types of changes are limited to linear and sinuous/cyclic changes and there was no consideration of hard/soft constraints. It would be interesting to extend the analysis on the multimodal set of benchmark problems in [18] and apply other types of changes such as random, chaotic and non-linear changes.

REFERENCES

[1] Y. Jin and J. Branke, "Evolutionary optimization in uncertain environments—a survey," *IEEE Transactions on Evolutionary Computation*, vol. 9, no. 3, pp. 303–317, 2005.

[2] M. Andrews and A. L. Tuson, "Dynamic optimisation: A practitioner requirements study," in *Proceedings of the The 24th Annual Workshop of the UK Planning and Scheduling Special Interest Group (PlanSIG 2005)*, London, UK, 2005.

[3] Y. Wang and M. Wineberg, "Estimation of evolvability genetic algorithm and dynamic environments," *Genetic Programming and Evolvable Machines*, vol. 7, no. 4, pp. 355–382, 2006.

[4] D. M. Prata, E. L. Lima, and J. C. Pinto, "Simultaneous data reconciliation and parameter estimation in bulk polypropylene polymerizations in real time," *Macromolecular Symposia*, vol. 243, no. 1, pp. 91–103, 2006.

[5] L. Araujo and J. J. Merelo, "A genetic algorithm for dynamic modelling and prediction of activity in document streams," in *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*. New York, NY, USA: ACM, 2007, pp. 1896–1903.

[6] P. Tawdross, S. K. Lakshmanan, and A. Konig, "Intrinsic evolution of predictable behavior evolvable hardware in dynamic environment," in *HIS '06: Proceedings of the Sixth International Conference on Hybrid Intelligent Systems*. Washington, DC, USA: IEEE Computer Society, 2006, p. 60.

[7] M. Rocha, J. Neves, and A. Veloso, "Evolutionary algorithms for static and dynamic optimization of fed-batch fermentation processes," in *Adaptive and Natural Computing Algorithms*, B. Ribeiro *et al.*, Eds. Springer, 2005, pp. 288–291.

[8] K. Deb, U. B. Rao, and S. Karthik, "Dynamic multi-objective optimization and decision-making using modified NSGA-II: A case study on hydro-thermal power scheduling," in *Evolutionary Multi-Criterion Optimization, 4th International Conference, EMO 2007, Matsushima, Japan, March 5-8, 2007, Proceedings*, ser. Lecture Notes in Computer Science, vol. 4403. Springer, 2007, pp. 803–817.

[9] P. Ioannou, A. Chassiakos, H. Jula, and R. Unglaub, "Dynamic optimization of cargo movement by trucks in metropolitan areas with adjacent ports," METRANS Transportation Center, University of Southern California, Los Angeles, CA 90089, USA., Tech. Rep., 2002. [Online]. Available: www.metrans.org/research/final/00-15_Final.htm

[10] K. Mertens, T. Holvoet, and Y. Berbers, "The DynCOAA algorithm for dynamic constraint optimization problems," in *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. New York, NY, USA: ACM, 2006, pp. 1421–1423.

[11] J. M. Thompson and K. A. Dowsland, "A robust simulated annealing based examination timetabling system," *Computers and Operations Research*, vol. 25, no. 7-8, pp. 637–648, 1998.

[12] U. Aickelin and K. Dowsland, "Exploiting problem structure in a genetic algorithm approach to a nurse rostering problem," *Journal of Scheduling*, vol. 3, pp. 139–153, 2000.

[13] H. Kim, "Target exploration for disconnected feasible regions in enterprise-driven multilevel product design," *American Institute of Aeronautics and Astronautics Journal*, vol. 44, no. 1, pp. 67–77, 2006.

[14] C. A. Liu, "New dynamic constrained optimization pso algorithm," in *ICNC '08: Proceedings of the 2008 Fourth International Conference on Natural Computation*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 650–653.

[15] H. Richter, "Memory design for constrained dynamic optimization problems," in *Proceedings of the European Conference on the Applications of Evolutionary Computation 2010, EvoApplications 2010*, ser. Lecture Notes in Computer Science, vol. 6024. Springer, 2010, pp. 552–561.

[16] T. T. Nguyen, "Continuous Dynamic Optimisation Using Evolutionary Algorithms," Ph.D. dissertation, School of Computer Science, University of Birmingham, January 2011, http://etheses.bham.ac.uk/1296 and http://www.staff.ljmu.ac.uk/enrtngu1/theses/phd_thesis_nguyen.pdf.

[17] T. T. Nguyen and X. Yao, "Benchmarking and solving dynamic constrained problems," in *Proceedings of the IEEE Congress on Evolutionary Computation CEC2009*. IEEE Press, May 2009, pp. 690–697.

[18] T. T. Nguyen, "A proposed real-valued dynamic constrained benchmark set," School of Computer Science, Univesity of Birmingham, Tech. Rep., 2008. [Online]. Available: \url{http://www.cs.bham.ac.uk/~txn/Papers/DCOP_benchmark.pdf}

[19] T. T. Nguyen and X. Yao, "Dynamic time-linkage problem revisited," in *Proceedings of the 2009 European Workshops on Applications of Evolutionary Computation, EvoWorkshops 2009*, ser. Lecture Notes in Computer Science, M. Giacobini, P. Machado, A. Brabazon, J. McCormack *et al.*, Eds., vol. 5484, 2009, pp. 735–744.

[20] C. Floudas, P. Pardalos, C. Adjiman, W. Esposito, Z. Gumus, S. Harding, J. Klepeis, C. Meyer, and C. Schweiger, *Handbook of Test Problems in Local and Global Optimization*, ser. Noncovex Optimization and Its Applications. Kluwer Academic Publishers, 1999, vol. 33.

[21] H. G. Cobb, "An investigation into the use of hypermutation as an adaptive operator in genetic algorithms having continuouis, time-dependent nonstationary environments," Naval Research Laboratory, Washington, USA, Technical Report AIC-90-001, 1990.

[22] D. Parrott and X. Li, "Locating and tracking multiple dynamic optima by a particle swarm model using speciation." *IEEE Trans. Evolutionary Computation*, vol. 10, no. 4, pp. 440–458, 2006.

[23] I. Moser and T. Hendtlass, "A simple and efficient multi-component algorithm for solving dynamic function optimisation problems," in *Proceedings of the IEEE Congress on Evolutionary Computation CEC'07.*, 2007, pp. 252–259.

[24] J. J. Grefenstette, "Genetic algorithms for changing environments," in *Parallel Problem Solving from Nature 2*, R. Maenner and B. Manderick, Eds. North Holland, 1992, pp. 137–144.

[25] S. Yang and X. Yao, "Experimental study on population-based incremental learning algorithms for dynamic optimization problems," *Soft Computing - A Fusion of Foundations, Methodologies and Applications*, vol. 9, no. 11, pp. 815–834, 2005.

[26] T. Blackwell and J. Branke, "Multiswarms, exclusion, and anti-convergence in dynamic environments." *IEEE Trans. Evolutionary Computation*, vol. 10, no. 4, pp. 459–472, 2006.

[27] J. Branke, *Evolutionary Optimization in Dynamic Environments*. Kluwer, 2001.

[28] H. Richter, "Detecting change in dynamic fitness landscapes," in *Proceedings of the IEEE 2009 Congress on Evolutionary Computation, CEC2009*, 2009, pp. 1613–1620.

[29] H. Richter and S. Yang, "Learning behavior in abstract memory schemes for dynamic optimization problems," *Soft Computing*, vol. 13, no. 12, pp. 1163–1173, 2009.

[30] D. Ayvaz, H. Topcuoglu, and F. Gurgen, "A comparative study of evolutionary optimization techniques in dynamic environments," in *GECCO '06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*. New York, NY, USA: ACM, 2006, pp. 1397–1398.

[31] T. T. Nguyen and X. Yao, "Solving dynamic constrained optimisation problems using repair methods," *IEEE Transactions on Evolutionary Computation (submitted)*, 2010. [Online]. Available: \url{http://www.cs.bham.ac.uk/~txn/Papers/Nguyen_Yao_dRepairGA.pdf}

[32] K. A. Morales and C. Quezada, "A universal eclectic genetic algorithm for constrained optimization," in *Proceedings 6th European Congress on Intelligent & Soft Computing, EUFIT'98*, 1998, pp. 518–522.

[33] Z. Michalewicz, "The second version of Genocop III: a system which handles also nonlinear constraints," School of Computer Science, University of Adelaide, [Accessed February 2009]. [Online]. Available: {http://www.cs.adelaide.edu.au/~zbyszek/EvolSyst/gcopIII10.tar.Z}

[34] T. T. Nguyen and X. Yao, "Detailed experimental results of GA, RIGA, HyperM and GA+Repair on the G24 set of benchmark problems," School of Computer Science, University of Birmingham, Tech. Rep.,

2010, online. [Online]. Available: \url{http://www.cs.bham.ac.uk/~txn/Papers/DCOP_fulldata.pdf}

[35] H. Motulsky, *Intuitive Biostatistics*, 1st ed. Oxford University Press, October 1995.

[36] T. Back, D. B. Fogel, and Z. Michalewicz, Eds., *Handbook of Evolutionary Computation*. Bristol, UK, UK: IOP Publishing Ltd., 1997.

[37] C. A. Coello Coello, "Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art," *Computer Methods in Applied Mechanics and Engineering*, vol. 191, no. 11-12, pp. 1245–1287, January 2002.

[38] T. P. Runarsson and X. Yao, "Stochastic ranking for constrained evolutionary optimization," *IEEE Trans. Evolutionary Computation*, vol. 4, no. 3, pp. 284–294, 2000.

[39] E. Mezura-Montes and C. A. C. Coello, "A simple multimembered evolution strategy to solve constrained optimization problems," *IEEE Trans. Evolutionary Computation*, vol. 9, no. 1, pp. 1–17, 2005.

[40] S. Venkatraman and G. G. Yen, "A generic framework for constrained optimization using genetic algorithms," *IEEE Trans. Evolutionary Computation*, vol. 9, no. 4, pp. 424–435, 2005.

[41] T. P. Runarsson and X. Yao, "Search biases in constrained evolutionary optimization," *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 35, no. 2, pp. 233–243, 2005.

[42] J. Joines and C. Houck, "On the use of non-stationary penalty functions to solve nonlinear constrained optimization problems with gas," in *Proceedings of the first IEEE Conference on Evolutionary Computation*, D. Fogel, Ed. IEEE Press, 1994, pp. 579–584.

[43] A. B. Hadj-Alouane and J. C. Bean, "A genetic algorithm for the multiple-choice integer program," *Operations Research*, vol. 45, pp. 92–101, 1997.

[44] S. B. Hamida and A. Petrowski, "The need for improving the exploration operators for constrained optimization problems," in *Proceedings of the Congress on Evolutionary Computation 2000, CEC'00*, vol. 2, 2000, pp. 1176–1183.

[45] Z. Michalewicz and G. Nazhiyath, "Genocop III: A co-evolutionary algorithm for numerical optimization with nonlinear constraints," in *Proceedings of the Second IEEE International Conference on Evolutionary Computation*, D. B. Fogel, Ed. Piscataway, New Jersey: IEEE Press, 1995, pp. 647–651.

[46] J. J. Grefenstette, R. Gopal, B. Rosmaita, and D. van Gucht, "Genetic algorithm for the TSP," in *Proceedings of the First International Conference on Genetic Algorithms*, 1985, pp. 160–168.

[47] Z. Michalewicz, *Handbook of Evolutionary Computation*. Oxford University Press, 1997, ch. Constraint-Handling Techniques : Decoders., pp. C5.3:1–C5.3:3.

[48] G. Syswerda, "Schedule optimization using genetic algorithms," in *Handbook of Genetic Algorithms*, L. Davis, Ed. New York: Van Nostrand Reinhold, 1991, pp. 332–349.

[49] S. Salcedo-Sanz, "A survey of repair methods used as constraint handling techniques in evolutionary algorithms," *Computer Science Review*, vol. 3, no. 3, pp. 175 – 192, 2009.

[50] R. Farmani and J. A. Wright, "Self-adaptive fitness formulation for constrained optimization," *IEEE Trans. Evolutionary Computation*, vol. 7, no. 5, pp. 445–455, 2003.

[51] S. B. Hamida and M. Schoenauer, "ASCHEA: new results using adaptive segregational constraint handling," in *Proceedings of the 2002 Congress on Evolutionary Computation, 2002. CEC'02.*, vol. 1. Los Alamitos, CA, USA: IEEE Press, 2002, pp. 884–889.

[52] T. Takahama and S. Sakai, "Constrained optimization by applying the alpha; constrained method to the nonlinear simplex method with mutations," *IEEE Transactions on Evolutionary Computation*, vol. 9, no. 5, pp. 437 – 451, oct. 2005.

[53] Y. Wang, Z. Cai, Y. Zhou, and W. Zeng, "An adaptive tradeoff model for constrained evolutionary optimization," *IEEE Transactions on Evolutionary Computation*, vol. 12, no. 1, pp. 80 –92, feb. 2008.

[54] K. P. Williams, "Simple genetic algorithm (SGA) source code (in C)," [Accessed December 2008]. [Online]. Available: http://www.kenwilliams.org.uk/code/ga2.c