Multi-Objective Optimization of Cancer Chemotherapy Using Swarm Intelligence

Andrei Petrovski[1], John McCall[1] and Bhavani Sudha[1]

Abstract. Cancer chemotherapy aims at achieving a number of treatment goals, not all of which are commensurate. Because of this, the problem of chemotherapy optimization necessitates the use of multi-objective optimization methods. The techniques based on swarm intelligence have certain features that make them applicable and effective in addressing multiple treatment objectives of cancer chemotherapy. This paper demonstrates the adaptive capabilities of particle swarm optimization (PSO) that enables this bio-inspired metaheuristic to carry out an efficient search for both effective and versatile chemotherapy treatments.

Keywords. Swarm intelligence, multi-objective optimization, medicine, metaheuristics.

1. INTRODUCTION

Chemotherapeutic cancer treatment necessitates making complex, and often life-critical, decisions about the best way to administer cytotoxic (i.e. destroying the cells) drugs. Typically patient information is incomplete and noisy, the range of treatment options is subject to complex constraints and the aims of treatment are multi-objective. For these reasons, medical decision support is a rich application area for evolutionary algorithms and related techniques [5].

The authors have developed a decision support system, the Oncology Workbench [4], the primary aims of which are to help oncologists to design, evaluate and optimise multi-drug chemotherapy treatments of cancer. The optimization engine of the Oncology Workbench (OWCH) is capable of determining better sequencing and dosing of the specified cytotoxic drugs, and recommends the optimised treatment(s) to the oncologists for consideration.

The current version of the optimization facility has several limitations though. Firstly, the search for better treatment schedules is conducted using genetic algorithms. Our previous studies [6] and [4] showed that alternative evolutionary techniques of computational optimization, the particle swarm optimization (PSO) in particular, can facilitate more efficient optimization of chemotherapeutic treatments. The better efficiency is achieved through finding feasible (i.e. satisfying all cancer chemotherapy constraints) treatment schedules much faster.

Secondly, the current version of the OWCH decision support system implements a single-objective optimization of cancer chemotherapy. The treatment schedules are selected on the basis of the tumour reduction criterion – that is, how effectively they reduce the overall tumour burden during the treatment period. Although tumour reduction, and ideally elimination, remains the primary objective of cancer chemotherapy, its achievement is not always possible due to various medical constraints. It is essential therefore to provide oncologists with the option to change the objective of treatment optimization from finding the best curative treatment to developing a good controlling (palliative) treatment capable of extending the patient survival time (PST) or the quality of patients' lives.

The authors have successfully explored the possibility of multi-objective optimization of cancer chemotherapy in their previous work [7], but confined their study to the domain of genetic algorithms only. Given inferior performance of this optimization technique in comparison with particle swarm optimization (PSO) and estimation of distribution algorithms (EDAs), it seems logical to revive the multi-objective approach to cancer chemotherapy optimization in the context of more efficient heuristics with the purpose to enhance the optimization facility of the developed decision support system for oncologists.

The remainder of this paper is organized as follows. In Section 2 we explain the objectives of cancer chemotherapy optimization, highlighting their conflicting nature that excludes the possibility of converting the problem from multiple to single objective optimization using weighting or any other techniques. Section 3 discusses how the multi-objectivity of chemotherapy optimization can be addressed by the PSO method. In Section 4 we present the experimental results that are followed by a short concluding Section 5, where we outline the direction of future work.

2. MILTI-OBJECTIVE NATURE OF CHEMOTHERAPY OPTIMIZATION

Chemotherapy is the use of cytotoxic drugs to control or eliminate cancer. The drugs are administered to the body by a variety of routes with the aim to create a certain concentration in the bloodstream that will act to systematically kill cells [1]. This means that both cancerous and normal cells will suffer due to the effect of these drugs.

The intention of the drug administration is to eradicate the tumour or, at least, to control the proliferation of cancerous cells. However, the treatment has toxic side-effects on the rest of the body. The success of chemotherapy therefore depends crucially on maintaining sufficient damage to the tumour while effectively managing the toxic side-effects.

[1] School of Computing, The Robert Gordon University, AB25 3UE, UK.
Email: {ap, jm, bs}@comp.rgu.ac.uk

Designing chemotherapeutic treatments is a complex optimal control problem that involves mathematical modelling of tumour growth, satisfaction of several medical constraints, and resolving conflicting objectives of treatment optimization. We will address all these aspects of cancer chemotherapy in the following subsections.

2.1 Mathematical Modelling of Tumour Growth and Reduction

Various models can be used to simulate the response of a tumour to chemotherapy. The most popular is the Gompertz growth model with a linear cell-loss effect [RASC'99]. The model takes the following form:

$$\frac{dN}{dt} = N(t)\cdot\left[\lambda \ln\left(\frac{\Theta}{N(t)}\right) - \right.$$
$$\left. - \sum_{j=1}^{d} \kappa_j \sum_{i=1}^{n} C_{ij}\left\{H(t-t_i) - H(t-t_{i+1})\right\}\right] \qquad (1)$$

where $N(t)$ represents the number of tumour cells at time t; λ and Θ are the parameters of tumour growth, $H(t)$ is the Heaviside step function; C_{ij} denote the concentrations of anti-cancer drugs exceeding the threshold concentration level below which the drugs have no therapeutic effect; and κ_j are the quantities representing the efficacy of anti-cancer drugs.

The first term in equation (1) represents the growth of an untreated tumour, whereas the second term models the effects of cytotoxic drugs leading to tumour reduction. The pharmacodynamics (PD) of drugs is taken into account by estimating their concentration above the therapeutic threshold that is equal to the minimum drug concentration necessary to kill any cancerous cells. The pharmacokinetics (PK) of treatment is addressed by means of finding drug concentrations at the tumour site, eliminating the necessity of tracking the drug dissipation throughout the body.

Although more detailed models of a tumour response to treatment have been introduced and studied in the context of computational optimization ([10], [14] and [15]), they rely on the correct estimation of a considerable number of PD/PK parameters, which significantly complicates their use in practice. Furthermore, the substitution $u(t) = \ln(\Theta/N(t))$ yields an analytical solution to Equation (1) that can be used evaluating the quality of treatment schedules found by computational search heuristics [3].

2.2 Constraints of Cancer Chemotherapy

The adverse effects of cancer chemotherapy stem from the systemic nature of this treatment: drugs are delivered via the bloodstream and therefore affect all body tissues. Since most anti-cancer drugs are highly toxic, they inevitably cause damage to sensitive tissues elsewhere in the body. In order to limit this damage, toxicity constraints need to be placed on the amount of drug applied at any time interval, on the cumulative drug dosage over the treatment period, and on the damage caused to various sensitive tissues. In addition to toxicity constraints, the tumour size (i.e. the number of cancerous cells) must be maintained below a lethal level during the whole treatment period for obvious reasons [7].

Cytotoxic drugs are usually delivered according to a discrete dosage program in which there are n doses given at times $t_1, t_2, \ldots t_n$ [3]. In the case of multi-drug chemotherapy, each dose is a cocktail of d drugs characterised by their concentration levels $C_{ij}, i \in \overline{1,n},\ j \in \overline{1,d}$ in the bloodplasma. If we denote a set of these concentrations as a vector $\mathbf{c} = (C_{ij})$, then the constraints of chemotherapeutic treatment of cancer can be represented in the following general form [7]:

1. Maximum instantaneous dose $C_{\max}$ for each drug acting as a single agent:

$$g_1(\mathbf{c}) = \left\{ C_{\max\, j} - C_{ij} \geq 0 \,\vdots\, \forall i \in \overline{1,n}, \forall j \in \overline{1,d} \right\} \qquad (2)$$

2. Maximum cumulative C_{cum} dose for drug acting as a single agent:

$$g_2(\mathbf{c}) = \left\{ C_{\text{cum}\, j} - \sum_{i=1}^{n} C_{ij} \geq 0 \,\vdots\, \forall j \in \overline{1,d} \right\} \qquad (3)$$

3. Maximum permissible size $N_{\max}$ of the tumour:

$$g_3(\mathbf{c}) = \left\{ N_{\max} - N(t_i) \geq 0 \,\vdots\, \forall i \in \overline{1,n} \right\} \qquad (4)$$

4. Restriction on the toxic side-effects of multi-drug chemotherapy:

$$g_4(\mathbf{c}) = \left\{ C_{\text{s-eff}\ k} - \sum_{j=1}^{d} \eta_{kj} C_{ij} \geq 0 \,\vdots\, \forall i \in \overline{1,n}, \forall k \in \overline{1,m} \right\} \qquad (5)$$

The factors η_{kj} in the last constraint represent the risk of damaging the k^{th} organ or tissue (such as heart, bone marrow, lung etc.) by administering the j^{th} drug.

It is also possible to introduce constraints associated with the drug delivery schedule. Standard chemotherapies often use fixed dosages and drug combinations at equally spaced intervals throughout the treatment period. The main reasons for such a delivery mode are that it is easier to organise patients to present for treatment at fixed times and that fixed dosages and combinations are more easily communicated and administered [4].

2.3 Objectives of Chemotherapy Treatments

The main goals of chemotherapeutic treatments are:

- *Cure*: curative treatments attempt to eradicate the tumour. In our studies we define eradication to mean a reduction of the tumour from an initial size of around 10^9 cells (minimum detectable tumour size) to below 10^3 cells.
- *Control*: if cure is not possible, the goal is to control the disease, i.e. stop the cancer from growing and spreading. One commonly used performance measure of controlling treatments is the patient survival time (PST) that needs to be extended.
- *Palliation*: if the cancer is at an advanced stage, even controlling the disease, let alone curing it, might become impossible. In such cases chemotherapy drugs may be used to relieve symptoms caused by the cancer, thereby improving the quality of life, even though the drugs may not lengthen the patient's life.

For some people, chemotherapy is the only treatment mode used in an attempt to cure, control, or palliate the cancer. At the same time, chemotherapy may be given along with other treatments – it may be used as *neoadjuvant* therapy (before surgery or radiation in order to shrink the tumour, for example), or as *adjuvant* therapy (after surgery or radiation to prevent the growth of stray cancer cells remaining in the body) [ACS].

Using the notation from the previous subsections, the objective of a curative treatment can be mathematically formulated as follows [3]:

$$\underset{\mathbf{c}}{\text{maximise}} \quad f_1(\mathbf{c}) = \int_{t_1}^{t_n} \ln\left(\frac{\Theta}{N(\tau)}\right) d\tau \qquad (\mathbf{6})$$

subject to the state equation (1) and the constraints (2)-(5).

If we denote the PST as T, then the treatment objective of extending the patient survival time (PST) can be represented as:

$$\underset{\mathbf{c}}{\text{maximise}} \quad f_2(\mathbf{c}) = \int_{t_1}^{T} d\tau = T \qquad (\mathbf{7})$$

again subject to (1)-(5).

The main objective of a palliative treatment is to improve the quality of life of the patients for whom neither curative nor controlling treatments can be found. As such, the last objective is qualitative in nature and it is not a straightforward task to express it in rigorous mathematical terms. One crude approach to improving the quality of patients' life is to minimise the amount of drugs delivered simultaneously in the attempt to weaken toxic side-effects caused by chemotherapy.

As can be concluded from the discussion above, cancer chemotherapy is a complex treatment mode that might be used for pursuing a variety of treatment objectives. What makes chemotherapy an interesting domain for applying multi-objective optimization techniques is that not all of these objectives are mutually exclusive. In fact, oncologists are often faced with unpredicted reactions from the patients to well-known and widely-used chemotherapy schedules. In such situations the treatment objective will have to be dynamically adjusted, or alternative treatment strategies will have to be developed that yield satisfactory outcomes with respect to more than one treatment objective. In the next section we will look at how swarm intelligence can assist oncologists in finding such strategies.

3. MULTI-OBJECTIVE OPTIMIZATION USING SWARM INTELLIGENCE

Swarm intelligence (SI) is based on the principles underlying the behaviour of natural systems consisting of many "agents", and exploiting local communication form, highly distributed control, and emergent strategies [12]. In contrast to centralized programs, based on the global goal, the SI algorithms use the synergy of the individual efforts of a population of agents that are not aware of the global objective to be reached. This creates a framework conducive for discovering versatile strategies highly adaptable to changing environments.

The driving force of the SI algorithms is the interactions between the agents themselves as well as between the agents and the environment. The capability of an individual agent is fairly limited; however, when a large number of agents are brought together, constructive behaviour can emerge that leads to effective addressing of the objectives not known *a priori*.

Computational modelling of swarms has resulted in numerous successful applications aimed at global function optimization [12], finding optimal routes and schedules, and at locating multiple optima [10]. In the next section we will look at how the SI algorithms – the Particle Swarm Optimization (PSO) method – can be adapted for solving multi-objective optimization problems.

3.1 Multi-objective Particle Swarm Optimization (MOPSO)

In general, a multi-objective optimization problem (MOP) consists of n decision variables comprising a decision vector $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \Omega \subset \Re^n$, m constraints $g_1(\mathbf{x}), g_2(\mathbf{x}), \ldots, g_m(\mathbf{x})$, and k objectives expressed as (non)linear criteria or objective functions $f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})$. Brought together, the multiple objectives define the evaluation function $F\big(f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})\big): \Omega \to \Lambda \subset \Re^k$, which, if some of the objectives are in conflict, places a partial, rather than normal, ordering on the search space Ω. In order to mathematically define this partial ordering, a notion of Pareto dominance is introduced in the objective space Λ.

The specificity of multi-objective optimization is to find a set of non-dominated decision vectors rather than the global optimum, which might not even exist. Pareto-optimal decision vectors cannot be improved in any objective without causing deterioration of at least one other objective. Such decision vectors comprise the Pareto-optimal set, $P^* \subset \Omega$, in the search space.

The mapping of the Pareto-optimal set to the objective function space gives rise to the Pareto front PF^*. The Pareto front can be non-convex and non-connected; nonetheless, if it is known, or at least approximated reasonably well, the decision maker will be able to select a solution via a choice of acceptable objective performance and, as a result of this, the problem of multi-objective optimization can be resolved.

Solving multi-objective optimization problems therefore involves a decision maker (DM), who brings additional knowledge about the problem, on the basis of which a particular solution from the Pareto set is chosen. Depending on the stage at which the DM is consulted, multi-objective optimization is categorized into *a priori*, *a posteriori*, and interactive optimization [11].

In *a priori* optimization the DM is consulted at an early stage with the purpose to set the optimization priorities. If *a posteriori* optimization is used, the DM is presented with a set of Pareto optimal solutions, and the task is to select the most suitable one. The interactive optimization combines both approaches and provides dynamic feedback from the DM that is used to guide the optimization process.

Many studies have shown that the process of multi-objective optimization can be effectively implemented by evolutionary algorithms. A number of these algorithms specifically aimed at solving multi-objective optimization problems include the non-dominated sorting GA II (NSGA II), the niched Pareto GA (NPGA), the Pareto-archived evolutionary strategy (PAES), the Pareto-envelope based selection algorithm (PESA), the strength Pareto evolutionary algorithm 2 (SPEA2), the incrementing multi-objective optimization evolutionary algorithm (IMOEA), and the like [11].

A relatively new addition to the class of evolutionary algorithms used in the context of multi-objective optimization is the approach based on particle swarms. PSO seems particularly suitable for multi-objective optimization mainly because of the high convergence speed [2]. Furthermore, a number of techniques have been proposed to improve the efficiency of multi-objective PSO (MOPSO) [13].

A detailed description of a MOPSO algorithm is provided in [2]. We have adopted this algorithm for addressing the problem of cancer chemotherapy optimization given multiple treatment objectives. The implementation issues are discussed in the next section.

3.2 Applying MOPSO for Optimising Chemotherapeutic Treatments

In our previous work [6] we explored the local- and global-best PSO algorithms on a single-objective optimization problem of finding a curative cancer chemotherapy treatment. In our implementation, each particle represents a candidate treatment schedule $\mathbf{c} = (C_{ij}), i \in \overline{1,n},\ j \in \overline{1,d}$, encoded as a string of integers. The representation space $\mathbf{I}$ (a discretized version of Ω) can then be expressed as a Cartesian product:

$$\mathbf{I} = C_1^1 \times C_1^2 \times \ldots \times C_1^d \times C_2^1 \times C_2^2 \times \ldots \times C_2^d \times \ldots \times C_n^1 \times C_n^2 \times \ldots \times C_n^d \quad \textbf{(8)}$$

of the sets C_i^j, each of which represents an integer value of concentration units for the i^{th} dose of the j^{th} drug. The magnitude of the concentration unit is calculated as a certain portion of the maximum instantaneous dose $C_{\max j}$ for each cytotoxic drug used.

The PSO algorithm is initialised with a population of random candidate solutions, conceptualised as particles. These particles are flown through the hyperspace Ω of solutions to the chemotherapy optimization problem. The position of each particle $\bar{c}_i^{k+1}$ at iteration $k+1$ corresponds to a treatment schedule of cytotoxic drugs and is determined by the following formula:

$$\bar{c}_i^{k+1} = \bar{c}_i^{k} + \bar{v}_i^{k} \quad \textbf{(9)}$$

where is $\bar{v}_i^{k}$ a randomised velocity vector assigned to each particle in a swarm. The velocity vector drives the optimization process and reflects the 'socially exchanged' information. Each particle in the swarm is attracted towards the locations representing best chemotherapeutic treatments found by the particle itself, its neighbours, and/or the entire population. This is achieved by defining the velocity vector in (9) for each particle as:

$$\bar{v}_i^{k} = w \cdot \bar{v}_i^{k-1} + b_1 \cdot r_1 \cdot \left(\bar{c}_i^{*} - \bar{c}_i^{k-1}\right) + b_2 \cdot r_2 \cdot \left(\bar{c}_i^{**} - \bar{c}_i^{k-1}\right) \quad \textbf{(10)}$$

where:

- w is the inertia coefficient the value of which is randomly generated from the range $[0.5,1]$;
- b_1 and b_2 are empirical coefficients used to improve PSO performance – in our experiments $b_1 = b_2 = 4$;
- r_1 and r_2 are random numbers is the range $[0,1]$;
- $\bar{c}_i^{*}$ and $\bar{c}_i^{**}$ are the best locations in Ω found by the particle i and the entire population respectively;
- $\bar{v}_i^{k-1}$ is the value of particle i velocity at previous iteration of the algorithm; the values $\bar{v}_i^{0}$ are initialised at random from the range $[0,2]$, i.e. $\bar{v}_i^{0} \in [0,2],\ i \in \overline{1,50}$. Particle velocities calculated from (10) have the lower and upper bounds, $|v_i| \le 1$, that reduce the particles' oscillation.

The search heuristic of the PSO algorithms, defined by equations (9)-(10), tend to focus exploration around remembered locations in the search space – a particular advantage when optimal solutions are confined to the boundary of a feasible region (wherein all constraints are satisfied)[4]. Multi-objective PSO approximates this boundary using the archive of non-dominated solutions, which in our implementation has the maximum size of

500 chemotherapy schedules that perform well with respect to the multiple treatment objectives described in Section 2.3.

The program implementing the MOPSO algorithms is written in Java. The termination criterion was chosen to be 250,000 fitness function evaluations. With the population of 50 particles, this implies evolving 5,000 generations of particle swarms. The results of our experiments are presented in the next section we present.

4. EXPERIMENTAL RESULTS

In order to provide a quantitative assessment of the MOPSO performance, three measures are often taken into consideration – the distribution of particles, their spread across the Pareto optimal front, and the ability to attain the global tradeoff [Liu, Gou, et al].

For the problem of multi-objective chemotherapy optimization neither the global tradeoff is known, nor can the true Pareto front be determined precisely. Therefore, as far as the performance measures are concerned, we are limited to only qualitative analysis of the particles' distribution and to the speed with which MOPSO can find feasible solutions to the multi-objective problem of chemotherapy optimization. Also, we can compare the quality of the solutions found by MOPSO with those discovered by a single-objective PSO algorithm that has been previously proposed by the authors [6].

Figure 1shows the spread of non-dominated solutions stored in the MOPSO archive at the end of a typical experimental run. The x-axis represents the quality of solutions with respect to the curative objective of chemotherapy treatment (6), whereas the y-axis shows the quality measure of palliative treatment (in both cases larger values are preferable because the objectives need to be maximized).

As can be seen from the figure, the non-dominated solutions are well-spread, providing a good incentive for the particles to fully explore the solution space Ω in the vicinity of the Pareto set. This is particularly important in the context of MOPSO because for each particle the global best reference $\overline{c}_i^{**}$ is selected from the set of non-dominated solutions based on the closest Euclidian distance.

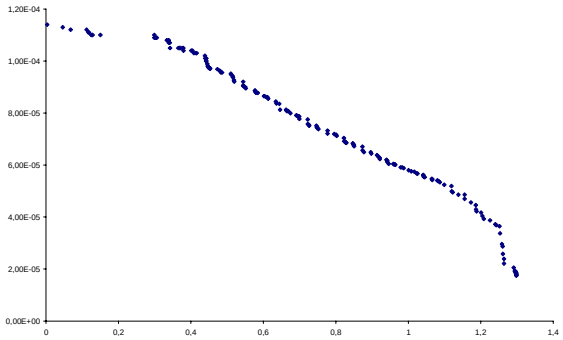

Figure 1. A set of non-dominated solutions found by the MOPSO algorithm

Secondly, the quality of the MOPSO solutions can be evaluated by comparing them with the best solutions found by a single-objective PSO algorithm in [PPSN'04]. The analysis shows that the performance index of the best MOPSO solution (Treatment **A**) with respect to the treatment objective (6) is 3.328 (indicating the log value of the ratio of the initial tumour size N_0 and its size at the end of the treatment interval), whereas the single-objective PSO (Treatment **B**) yields the value of 3.330 for the same objective.

Figure 2 shows the effects of both treatments on the same tumour and on the patient survival time. It demonstrates that the Treatment A found by MOPSO reduces the tumour almost as effectively as Treatment B, but given its reasonable performance with respect to the palliative treatment objective, causes substantially weaker toxic side-effects, improving thereby the quality of the patient life.

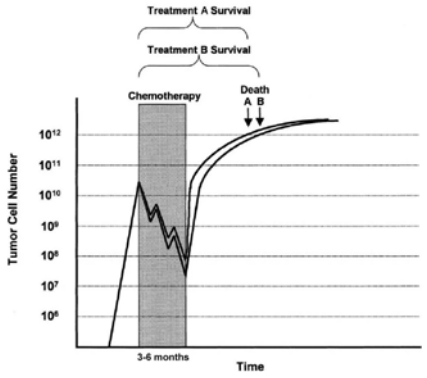

Figure 2. Multi- vs. single-objective PSO on a commensurate treatment objective

Thirdly, our experimental results have shown, that the MOPSO algorithm finds a feasible solution within 25 generations, indicating its superior search efficiency in comparison with other evolutionary algorithms – genetic algorithms in particular [6].

5. CONCLUSIONS AND FUTURE WORK

The optimization of cancer chemotherapy based on the swarm intelligence approach demonstrates high adaptive capabilities that can be effectively used for addressing multiple treatment objectives. The swarm's ability to retain information on the good solutions found by each particle and pass it round the whole population in an effective manner is a valuable asset in solving both single- and multi-objective chemotherapy optimization problems. The provision of a multi-objective optimization facility will expand the scope and enhance the quality of the decision support available to the oncologists from the Oncology Workbench system developed by the authors [4].

REFERENCES

[1] American Cancer Society:

http://www.cancer.org/docroot/home/index.asp

[2] Coello, C., Pulido, G., Lechuga, M.: Handling multiple objectives with particle swarm optimization. In *IEEE Transactions on Evolutionary Computation*, **8**, No. 3 (2004), pp. 256--79.

[3] Martin, R., Teo, K.: Optimal Control of Drug Administration in Cancer Chemotherapy. World Scientific, Singapore New Jersey London Hong Kong (1994).

[4] .McCall, J., Petrovski, A., Shakya, S.: Evolutionary Algorithms for Cancer Chemotherapy Optimization. In *Computational Intelligence in Bioinformatics* by Fogel D., Corne D., and Pan Y. (Eds.). IEEE Press (2008) pp. 265-96.

[5] Petrovski, A., McCall, J.: Smart problem solving environments for medical decision support. Proceeding of the *Workshops on Genetic and Evolutionary Computation (GECCO'05)*. ACM Press, New York (2005), pp.152-58.

[6] Petrovski, A., Sudha, B., McCall, J.: Optimising cancer chemotherapy using particle swarm optimization and genetic algorithms. Proceeding of the *8th International Conference on Parallel Problem Solving from Nature*, Lecture Notes in Computer Science, Volume **3242**. Springer, Berlin, Heidelberg (2004), pp. 633-41.

[7] Petrovski, A., McCall, J. A.: Multi-objective optimization of cancer chemotherapy using evolutionary algorithms. *Evolutionary Multi-Criterion Optimization*, Lecture Notes in Computer Science, Volume **1993**, Zitzler E., et al. (Eds.) Springer, Berlin, Heidelberg (2001), pp. 531-45.

[8] Petrovski, A., McCall, J.: Computational optimization of cancer chemotherapies using genetic algorithms. In: John, R., Birkhead, R. (eds.): *Soft Computing Techniques and Applications.* Series on Advances in Soft Computing. Physica-Verlag (1999), pp. 117-122.

[9] Liu, D., Tan, K., Goh, C., Ho, W.: A multi-objective memetic algorithm based on particle swarm optimization. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics*, **37**(1), (2007), pp. 42-50.

[10] Kay Chen Tan, Khor, E. F., Cai, J., Heng, C. M., Lee, T. H.: Automating the drug scheduling of cancer chemotherapy via evolutionary computation. Artificial Intelligence in Medicine **25**(2), (2002, pp. 169-185.

[11] Knowles, J., Corne, D.: Memetic algorithms for multi-objective optimization: Issues, methods and prospects. In W. Hart, N. Krasnogor and J.E. Smith (Eds.), *Recent Advances in Memetic Algorithms*. Springer: Studies in Fuzziness and Soft Computing, **166**, (2005). pp. 313—52.

[12] Olariu, S., Zomaya, A. (Eds.): *Handbook of Bioinspired Algorithms and Applications*. Chapman and Hall (2006).

[13] Reyes-Sierra, M., Coello, C.: A study of techniques to improve the efficiency of a multi-objective particle swarm optimizer. In S. Yang, Y.-S. Ong and Y. Jin (Eds.), *Evolutionary Computation in Dynamic and Uncertain Environments*. Springer (2007), pp. 269—96.

[14] Tse, S., Liang, Y., Leung, K., Lee, K., Mok T.: A memetic algorithm for multiple-drug cancer chemotherapy schedule optimization. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics* (2007), Vol. **37-1**, pp. 84-91.

[15] Villasana, M. and Ochoa, G.: Heuristic design of cancer chemotherapy. *IEEE Transactions of Evolutionary Computation* (2004), Vol. **8**, pp. 513-21.