PATIENT JOURNEY OPTIMIZATION USING A MULTI-AGENT APPROACH

Chung-ho Choi, William K. Cheung, Jiming Liu
Department of Computer Science, Hong Kong Baptist University, Hong Kong

Ian T. Cheung
Quality and Safety Division, Hospital Authority, Hong Kong

Keywords: Multi-agent systems, Patient journey optimization, Scheduling algorithms, Health information systems.

Abstract: With the increasing expectation from patients and the regulations enacted by the government, exploring ways to shorten patient journey has caught increasing attention. Patient journey optimization typically involves the coordination of treatment scheduling at multiple medical units. This decentralized nature of the problem makes conventional centralized operation research methods hard to be applied and motivates the use of the multi-agent approach. In this paper, we focus on cancer patient treatment. We model patients and medical units as autonomous agents which interact locally via a bidding process and a coordination process for patient journey optimization. With reference to a dataset containing more than five thousand cancer patient journeys, the effectiveness of the proposed algorithm under different settings of implementation has been evaluated via experimental simulations.

1 INTRODUCTION

Shortening the length of patient journey with unnecessary waiting time avoided is always an expectation of patients. Also, it is one of the key performance indicators for evaluating the effectiveness of healthcare providers. In order to raise the standard for healthcare services, careful application of scheduling algorithms is crucial (Vermeulen et al., 2006). While conventional operations research methods have been found effective for scheduling-related problems (Vissers et al., 2005; Patrick and Puterman, 2008), most of them take the centralized approach and are not designed to be applied to decentralized situations (Vermeulen et al., 2006; Decker et al., 2000).

The multi-agent approach is characterized by emphasizing on local interaction and self-organization of different entities being modeled. These properties make it especially suitable for tackling complex tasks with a lot of stakeholders involved and the environment is dynamic (Czap and Becker, 2003). Multi-agent methods have found applications in a variety of problem domains, such as airport resource scheduling (Mao et al., 2007), load allocation in transportation logistics (Robu et al., 2008), supply chain management (Wang et al., 2008), etc. Recently, it has also been applied to patient scheduling in (Paulussen et al., 2003; Vermeulen et al., 2006) with some initial success demonstrated. And yet, there are limitations. Paulussen *et al.* (Paulussen et al., 2003) assumed that a health state can be accurately quantified as utility measure for guiding the scheduling process. Vermeulen *et al.* (Vermeulen et al., 2006) assumed that temporal constraints between treatment operations need not be considered during the scheduling process.

The objective of this study is to explore the extent to which patient journey can be improved by better coordinating and mobilizing resources distributed at different medical units. In particular, we formulate the scheduling problem according to the cancer treatment practice in Hong Kong. We propose the use of a multi-agent approach in which autonomous agents interact with each other to arrive an effective overall schedule with reduced waiting times. To evaluate the effectiveness of the proposed approach, we have carried out simulations with the proposed approach given different settings of the environment. To initialize the simulation environment, we made use of a patient identity anonymized data set collected by Hospi-

Choi C., K. Cheung W., Liu J. and T. Cheung I. (2010).
PATIENT JOURNEY OPTIMIZATION USING A MULTI-AGENT APPROACH.
In *Proceedings of the Third International Conference on Health Informatics*, pages 181-186
DOI: 10.5220/0002693901810186

tal Authority in Hong Kong which contains 5819 cancer patient journeys with the admission period spanning over 6 months.

The rest of the paper is organized as follows. The patient scheduling problem formulation is described in Section 2. Sections 3 and 4 present the details of the proposed agent-based scheduling algorithm. Section 5 presents some preliminary experimental results and Section 6 concludes the paper.

2 PROBLEM FORMULATION

In this section, we first briefly describe the establishment of the cancer clusters in Hong Kong. Then, we formulate the patient scheduling problem for cancer treatment as an optimization problem.

2.1 Cancer Patient Treatment - A Hong Kong Scenario

In Hong Kong, there are seven cancer clusters established and we denote them as $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, ..., \mathcal{C}_7\}$. Currently, on-demand information exchange among the clusters for patient scheduling is not yet extensively used. That is why it is common for cancer patients to be scheduled to receive treatments at only one cancer cluster, even though some of the treatments could be provided earlier by other clusters.

Generally speaking, once a patient is suspected to have cancer, the doctor will specify the patient a treatment plan which contains a sequence of treatment operations. We denote the set of treatment operations as $\Gamma = \{$radiotherapy, surgery, chemotherapy$\}$.

To carry out the treatment operations, medical resources are needed. We denote the set of medical resources (or units) as $A = \{$radiotherapy unit, operation unit, chemotherapy unit$\}$. We assume that one treatment operation can only be performed at one medical unit of the corresponding type. A patient journey is defined as the duration from the date of admission to the date of the last treatment operation completed.

2.2 Formulation

Let $K := A \times \mathcal{C}$ be the cartesian product of A and $\mathcal{C}$ giving the complete set of medical units, $M := K \rightarrow \Gamma$ be an one-to-one mapping between K and Γ specifying the treatment type of the medical units, and P be the set of cancer patients being scheduled.

Also, given a patient i, let N^i_Γ denote the number of treatment operations needed, D^i_0 denote the admission date, D^i_j denote the date of the j^{th} treatment operation where $1 \leq j \leq N^i_\Gamma$, $V^i_j \in K$ be the unit at which the j^{th} treatment operation is performed, $Tr^i_j \in \Gamma$ be the type of treatment for the j^{th} operation, C_k be the daily capacity (i.e. number of patients that could be treated) of medical unit $k \in K$, T_t be the duration (in days) of treatment type $t \in \Gamma$, and Z be the set of dates on which patient scheduling is being considered.

With the assumption that all the patients are being treated equally in terms of urgency, the scheduling problem can be formulated as:

$$\min_D \sum_{i=1}^{|P|} \sum_{j=1}^{N^i_\Gamma - 1} (|D^i_j - D^i_{j+1}|) \quad (1)$$

with the following constraints to be satisfied:

$$D^i_{j+1} > D^i_j + T_{Tr^i_j} \quad (2)$$

$$\forall d \in Z \quad \left|\{i : D^i_j = d \wedge V^i_j = k \wedge Tr^i_j = M(k)\}\right| \leq C_k \quad (3)$$

$$D^i_j > D^i_0 > 0 \quad (4)$$

The objective function in (1) is to minimize the time lags between treatment operations for cancer patients. Constraint (2) ensures the temporal constraints between treatment operations are not violated, constraint (3) is used to ensure all medical units are operating within their capacities. Constraint (4) ensures that patients would only be scheduled to receive treatment operations after their admissions.

3 SCHEDULING FRAMEWORK

Theoretically, patient waiting times could be minimized by optimizing (1). However, it is impractical to do so as it is hard to assume that a cancer cluster is willing to share its real-time resource allocation related data (e.g., C_k) with other clusters due to both technical and managerial reasons.

Hence, in this section, we propose the use of a multi-agent approach for our domain. Particularly, in our proposed framework, there are two types of agents, namely *patient agents* and *resource agents*.

3.1 Patient Agent

A patient agent is used to represent one cancer patient and is denoted as P_i with $i = 1, 2, ..., |P|$. It stores the patient's treatment plan. As it is common that some treatment operations have to be performed in prior to another, the set of treatment operations to be received by a patient has to satisfy certain temporal constraints. Hence, each patient agent P_i maintains an ordered set $Tr^i = \{Tr^i_1, Tr^i_2, ...Tr^i_{N^i_\Gamma}\}$ as its treatment plan.

3.2 Resource Agent

A resource agent is used to manage a specific medical unit. Here, we denote R_{ab} as a resource agent representing medical unit $a \in A$ at cancer cluster $b \in \mathcal{C}$. Each resource agent has full access to the schedule of the medical unit it represents, but not the others.

3.3 Scheduling Algorithm

We adopt a two-phase scheduling algorithm similar to what being proposed in (Paulussen et al., 2003; Vermeulen et al., 2006). For each newly admitted patient, a treatment plan is first designed and then the corresponding treatment operations are initially scheduled (initial assignment phase). Then, a timeslot-swapping process is enforced for shortening the patient journey (rescheduling phase). Here we assume that any two patient agents are willing to exchange their timeslots as far as none of their schedules is worsen (as suggested in (Vermeulen et al., 2006)) and none of the temporal constraints as specified in *Eq.*(2) is violated.[1]

4 AGENT COORDINATION

In this section, more details about the scheduling algorithm are given, including 1) how the patient agents interact with the resource agents, and 2) how some "unnecessary" swappings can be rejected so as to further improve the scheduling optimality.

4.1 A Bidding Process for Agent Matchmaking

Figure 1 shows our proposed framework. As what have been introduced earlier, there are two types of agents, namely patient agents and resource agents. In order to show clearly the coordination between agents, we further categorize patient agents into *initiating patient agents* and *target patient agents*. Initiating patient agents P_I are those patient agents who initiate a request for timeslot exchange while target patient agents P_G are those who are willing to participate in the exchanging process.

In order to shorten its patient journey, an initiating patient agent P_I would first send out a request for rescheduling to the corresponding resource agents

[1]This assumption may imply that some policy-wise incentive has to be in place so that different medical units are willing to share their resources in this manner, which however is not the main focus of our study.

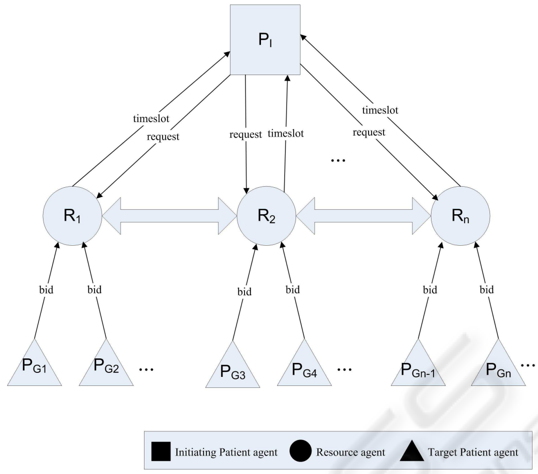

Figure 1: The proposed agent coordination framework for patient scheduling.

R_{ab}. The request includes the earliest possible start date (*EPS*) and the latest possible start date (*LPS*) of its associated treatment operation. In order not to violate the temporal constraints between treatment operations, the *EPS* can be defined as:

$$EPS^I_j = D^I_{j-1} + T_{Tr^I_{j-1}} + \delta_1. \tag{5}$$

Note that δ_1 denotes how many days a patient should be admitted (if needed) before a treatment operation to be carried out. In our experiment, we set to be 1. In practice, this value could be designated by healthcare providers. With a similar argument, *LPS* is defined as:

$$LPS^I_j = D^I_j - 1. \tag{6}$$

Once a resource agent receives a request with *EPS* and *LPS*, it will first check whether there are available timeslots released by deceased patients. If yes, the released timeslot will be assigned to the initiating patient agent. If not, the resource agent will then pass the request to those patient agents (target patient agents, P_G) which reserved resources of the same type in the period from *EPS* to *LPS*. Those target patient agents who have received the request will submit a bid to the resource agent in response.

There are several factors needed to be considered in computing the bid value.

- The target patient agent should not have its last treatment operation to be exchanged; otherwise, it would end up with a lengthened journey.
- As it is impractical to reschedule a patient's treatment operation without prior notification, we assume that the exchange of timeslots would not be considered if the patient will have less than a

week's time of notification.[2]

- The target patient agent has to ensure that the temporal constraints between its treatment operations would not be violated after the exchange.

Taking into account the above considerations, the bid value submitted by P_G is formulated as:

$$Bid^G = (D^G_{j_t} - EPS^I_{j_i}) + Last + Noti + Temp, \quad (7)$$

where $Last$, $Noti$ and $Temp$ are three binary variables. $Last = 0$ if the j_tth operation is not the last one for P_G, or ∞ otherwise. $Noti = 0$ if a week's time of notification for the target patient agent to be notified, or ∞ otherwise. $Temp = 0$ if there are no temporal constraints violated, or ∞ otherwise.

Among all the target patient agents, the one with the lowest bid value will be accepted and the timeslot swapping between P_I and P_G will be confirmed. If two bids are found to be numerically identical, the resource agent will select one at random.

4.2 A Coordination Process for Rejecting Unnecessary Swappings

A timeslot swapping confirmed as described in the previous section does not necessarily lead to a shortened patient journey. To illustrate that, suppose there is a patient agent with 3 treatment operations to be rescheduled. In case the last treatment operation could not be rescheduled to be performed earlier, any rescheduling of the first 2 are useless as the duration of the whole journey remains unchanged (see Figure 2(a)). As another example, even a shortened patient journey can be achieved, rescheduling of the first 2 treatment operations could also be useless if the rescheduling of the last one cannot be benefited from the rescheduling of the first two (see Figure 2(b)).

In order to eliminate these useless swappings, a resource agent after identifying the most optimal bid among the target patient agents will not notify the initiating patient agent immediately. Instead, it will pass the bid to the resource agent which is responsible for the succeeding treatment operation of the initiating patient agent. Having received such a bid, the resource agent could derive a new EPS, denoted as ${}^{(new)}EPS^I_{j_i+1}$. Clearly, unnecessary swappings occur if that resource agent could not find a bid such that ${}^{(new)}EPS^I_{j_i+1} \leq D^G_{j'_t} \leq EPS^I_{j_i+1}$, where $Tr^I_{j_i+1} = Tr^G_{j'_t}$. In that case, the resource agent will notify its antecedent to discard the bid such that the corresponding timeslots would not be exchanged.

[2]In general, the time of notification can be adjusted according to the real situation.

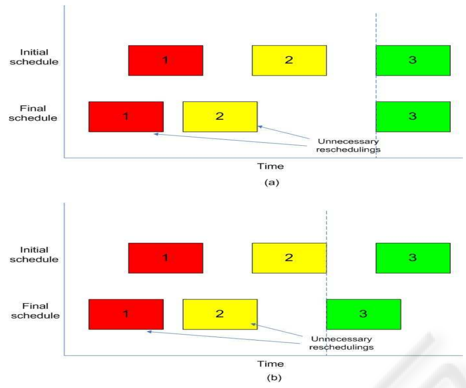

Figure 2: Unnecessary reschedulings.

5 EXPERIMENTAL VALIDATION

To evaluate the effectiveness of the proposed multi-agent approach, we first obtained a dataset containing 5819 cancer patients who were treated at the 7 cancer clusters in Hong Kong within an admission period of 6 months. The average length of the patient journey among all cancer clusters is 90.7 days. Based on the dataset, we had carried two groups of simulation.

For the first group of simulation, we made use of the scheduled treatment plans in the dataset for the initial assignment and studied to what extent the multi-agent approach can improve the patient journey as a whole.

For the second group of simulation, we aim to make the simulation more flexible by making use of only the statistics of the scheduled treatment plans and the capacities of the medical units as revealed in the dataset. The initial assignment strategy and the rescheduling strategy can both be specified by the user for evaluation. Also, we tried to increase the capacity value as revealed in the dataset by percentage so as to see to what extent the patient journey can be improved with more resources injected.

5.1 Simulation with Initial Assignment and Unit Capacity Fixed (Exp. A)

Four different settings had been used to demonstrate various aspects of the proposed algorithm:

Setting 1: Patient agents are willing to exchange timeslots with others whenever there is a Pareto improvement.

Setting 2: It is assumed that only 20% of the patients of each cancer cluster are allowed to undergo timeslot swapping.

Setting 3: It is assumed that only swappings between two nearby cancer clusters are allowed.

Setting 4: Timeslots released by deceased patients are allocated to the patient agents who have the longest patient journeys at a time point.

Given the four aforementioned settings, Figure 3a shows the average length of the patient journeys associated with the seven cancer clusters in Hong Kong.

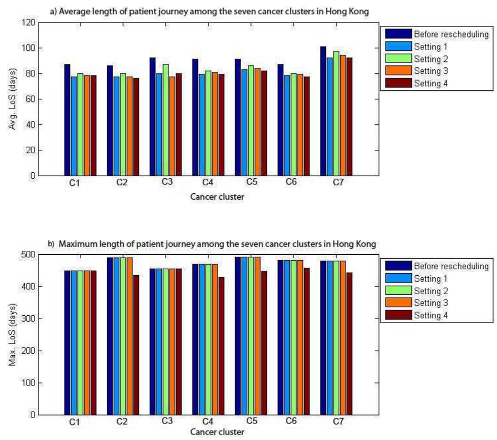

Figure 3: a) Average length of patient journey; b) Maximum length of patient journey among the seven cancer clusters in Hong Kong under 4 different settings in Exp. A.

The experimental results obtained show that, on average, the average length of journey can be reduced by 9.8 days for those 5819 cancer patients if no restriction is imposed on the exchange of timeslots whenever there is a Pareto improvement (Setting 1). Given only 20% of patients per cancer cluster are allowed for timeslot exchange (Setting 2), we found that the average length of journey could still be reduced by an average of 6.1 days. With the geographical restriction on allowing only swappings between nearby cancer clusters (Setting 3), the average length of journey can also be reduced by 9.3 days.

However, it should also be noted that according to Figure 3b, the maximum length of journey remains unchanged. The reason is obvious as no one is willing to swap with those with the longest length of journey. Reductions on the maximum length of journey can only be observed for Setting 4 where the released timeslots due to deceased patients are allocated to those with the longest journey.

5.2 Simulation Revealing the Effects of Varying Initial Assignment and Unit Capacity (Exp. B)

Contrary to the group of simulation previously presented, we tried to simulate also the initial assignment process using again the treatment plans of the 5819 cancer patients. However, all the treatment operations were not scheduled according to the data but simulated based on the statistics of the inter-operation duration obtained from the dataset. In particular, during the initial assignment phase, patient agents would be assigned an initial schedule one by one based on their admission orders. For each treatment operation, each patient agent would be assigned with the next earliest available timeslot. However, it is obvious that there should be a minimum time lag between two subsequent treatment operations due to different medical reasons. We first used the average inter-operation duration computed based on the dataset and then performed the simulation. The average length of journey computed right after the initial assignment was found to be 83.3 days. The improvements obtained due to the different settings are not very much different from those obtained in the previous section. Due to the page limit, we would not show them here.

According to the dataset, the minimum days between any two treatment operations was found to be one only. This implies that treatment operation sometimes could be started one day after another if the resource is available. We had also tried to set the minimum days to one in our initial assignment phase and compared with the results obtained before. According to Figure 4, while we observed some improvements in performance, the enhancement however is not very significant. Hence, setting some reasonable time lags between treatment operations does not have too big an impact on lengthening the patient journey.

For all the results presented so far, it is assumed that the capacity of each medical unit is fixed. To study the cost-effectiveness of increasing the units' capacities for patient journey optimization, we increased the capacity by the same percentage for all the medical units. According to Figure 5, it can be observed that when all the resource capacities are increased incrementally (10%, 20%, 30%), the reduction on average length of patient journey will then drop accordingly. In fact, such drop could be attributed to the fact that when the resource capacities are increased, patients would then be scheduled with less idle times between treatment operations; and hence with less chance to exchange timeslots with others.

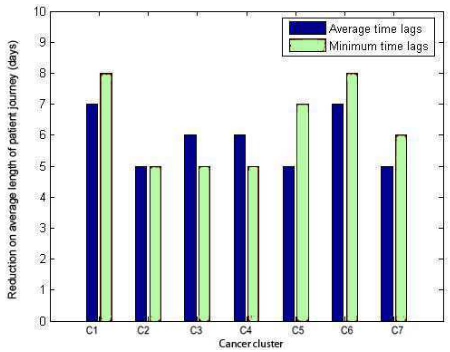

Figure 4: Reduction on average length of patient journey (Setting 4) by varying the time lags between treatment operations.

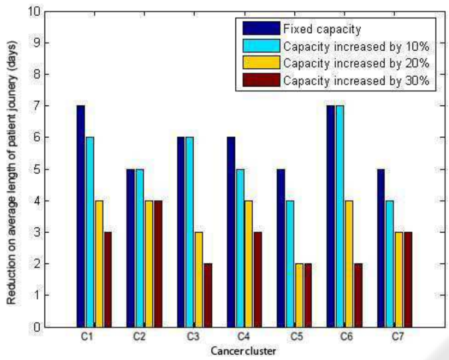

Figure 5: Reduction on average length of patient journey (Setting 4) by varying capacities.

6 CONCLUSIONS

In this paper, a multi-agent approach was proposed for patient journey optimization. Particularly, by applying the approach, the shortening of a patient journey will not lengthen the journeys of the others. Also, all the temporal constraints among the treatment operations for each patient would not be violated during the scheduling process.

The effectiveness of the proposed approach has been demonstrated by applying it to a dataset containing 5819 scheduled treatment plans of cancer patients in Hong Kong. The effects of varying the initial assignment and the unit capacity on the overall reduction in length of patient journey are also studied.

Currently, since we are using a Pareto improvement approach, it is assumed that no single patient (agent) would get a lengthened schedule after swapping timeslots with another. In the future, we are going to see whether there would be a greater improvement in achieving a reduced average length of patient journey when the above assumption for individuals is relaxed.

ACKNOWLEDGEMENTS

This is to acknowledge Hospital Authority in Hong Kong for providing the dataset to support this study.

REFERENCES

Czap, H. and Becker, M. (2003). Multi-agent systems and microeconomic theory: A negotiation approach to solve scheduling problems in high dynamic environments. In *HICSS '03: Proceedings of the 36th Annual Hawaii International Conference on System Sciences (HICSS'03) - Track 3*, page 83.2, Washington, DC, USA. IEEE Computer Society.

Decker, K., Li, J., and Demazeau, Y. (2000). Coordinating mutually exclusive resources using gpgp. *Autonomous Agents and Multi-Agent Systems*, 3:200–0.

Mao, X., Mors, A., Roos, N., and Witteveen, C. (2007). Coordinating competitive agents in dynamic airport resource scheduling. In *MATES '07: Proceedings of the 5th German conference on Multiagent System Technologies*, pages 133–144, Berlin, Heidelberg. Springer-Verlag.

Patrick, J. and Puterman, M. (2008). Reducing wait times through operations research: optimizing the use of surge capacity. *Healthc Q*, 11(3):77–83.

Paulussen, T. O., Dept, I. S., Decker, K. S., Heinzl, A., and Jennings, N. R. (2003). Distributed patient scheduling in hospitals. In *Coordination and Agent Technology in Value Networks. GITO*, pages 1224–1232. Morgan Kaufmann.

Robu, V., Noot, H., La Poutré, H., and van Schijndel, W.-J. (2008). An interactive platform for auction-based allocation of loads in transportation logistics. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 3–10, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Vermeulen, I., Bohte, S., Somefun, K., and La Poutre, H. (2006). Improving patient activity schedules by multi-agent pareto appointment exchanging. In *CEC-EEE '06: Proceedings of the The 8th IEEE International Conference on E-Commerce Technology and The 3rd IEEE International Conference on Enterprise Computing, E-Commerce, and E-Services*, page 9, Washington, DC, USA. IEEE Computer Society.

Vissers, J., Bekkers, J., and Adan, I. (2005). Patient mix optimization in tactical cardiothoracic surgery planning: a case study. *IMA Journal of Management Mathematics*, 16.

Wang, M., Liu, J., Wang, H., Cheung, W. K., and Xie, X. (2008). On-demand e-supply chain integration: A multi-agent constraint-based approach. *Expert Systems with Applications*, 34(4):2683 – 2692.