Operations Research for Health Care 3 (2014) 116–128

Contents lists available at ScienceDirect

Operations Research for Health Care

journal homepage: www.elsevier.com/locate/orhc

A generic modelling framework to evaluate network blood management policies: The Canadian Blood Services experience

John T. Blake *, Matthew Hardy

Dalhousie University, Department of Industrial Engineering, Halifax, Nova Scotia, Canada
Canadian Blood Services, Medical, Scientific, and Research Affairs, Ottawa, Ontario, Canada

ARTICLE INFO

Article history:
Received 31 October 2013
Accepted 13 May 2014
Available online 2 June 2014

Keywords:
Blood inventory
Simulation modelling
Network distribution
Response-surface methods
Re-useable simulations

ABSTRACT

Canadian Blood Services produces and distributes approximately 850,000 units of red cells annually. These units are distributed through ten production and/or distribution sites. Each distribution site acts as a regional hub serving between 20 and 110 hospital customers. Distribution sites hold a target inventory that is based on an integer number of median days demand on hand. In this paper, we report on the development and use of a simulation based methodology to evaluate network inventory policies for regional blood distribution sites in Canada. A generic framework was developed to represent each of the ten different regional networks. The modelling approach was validated by comparing model results against data from two networks. Once validated, ten instances were developed. For each model instance, a set of experiments was conducted, from which response surfaces were created. Non-linear optimization methods were applied to identify optimal supplier/consumer inventory policies using the response surfaces. We conclude that a generic modelling framework can be useful for regional blood supply chains, but suggest that at least four instances are necessary to recoup the efforts of building a reusable model.

1. Introduction

Canadian Blood Services is a not-for-profit, charitable organization whose mission is to manage the supply of blood and blood products in Canada. It operates in all Canadian provinces and territories, except for the Province of Quebec, which has its own blood supplier. Canadian Blood Services produces and distributes approximately 850,000 units of red blood cells (RBC) annually. These units are distributed through ten sites (nine production and distribution facilities and one distribution-only hub). Because both the supply of, and demand for, blood is stochastic, it is necessary to hold inventory to accommodate uncertainty. Canadian Blood Services has an established set of inventory targets and thresholds for all of its products, including RBC, measured in median days demand on hand (DOH). The targets are based on weekday demand data, averaged over a 25 week time horizon (see Table 1).

Threshold levels for RBC have been established through long experience and, while likely conservative, are simple to follow and well accepted by those within the organization and by its hospital customers. However, neither has the optimality of the threshold levels been established, nor is it well understood how the system would respond if inventory thresholds were changed.

The annual cost of operations for Canadian Blood Services in 2011–2012 was $1B (CAN), of which approximately 50% was related to the collection of transfusable products, the vast majority of which were whole blood from which RBCs are produced. Like most health care organizations, Canadian Blood Services is challenged by its funders to be an efficient provider. As a result, Canadian Blood Services has in recent years undergone a number of structural changes to streamline its operations; further changes are expected in the future (see Table 2).

In the Canadian Blood Services' environment, where network topology is in flux, but safety remains an overarching concern, it was deemed necessary to have a tool to evaluate inventory policy at a regional distribution sites (DS) in response to changes in the environment. However, the amount of inventory that should be held at a DS depends on the network in which it operates; collection rates, the amount of inventory held by consumers, the number of customer sites and their volume of demand, and transport distance, routes, and schedules all influence the need for inventory at a particular DS. Therefore, it was hypothesized that recommendations for inventory levels should take into account local conditions. However, building ten unique models, each representing a single distribution centre, would be prohibitively

* Correspondence to: Department of Industrial Engineering, Dalhousie University, 5269 Morris Street, Halifax, Nova Scotia, Canada B3H 4R2. Tel.: +1 902 494 6068.
E-mail address: john.blake@dal.ca (J.T. Blake).

http://dx.doi.org/10.1016/j.orhc.2014.05.002
2211-6923/

Table 1
Canadian Blood Services inventory thresholds, measured in median days inventory on hand. Centres strive to maintain their inventory within the "optimal" level.

Level	Days on hand (DOH)
Emergency	0.5–1.5
Serious	1.5–3.0
Minimum	3.0–5.0
Optimal	5.0–8.0
Maximum	8.0–10.0
Excess	>10.0

expensive. In this paper we describe the process to design, develop, and test a tool to evaluate inventory and inventory ordering policies for Canadian Blood Services distribution sites through the development of a generic modelling framework. We illustrate its application through an example case study to optimize supplier and consumer inventories in one network using response surface methods. Finally, we provide observations regarding the overhead requirements and potential benefits of adopting a reusable modelling framework.

2. Literature review

Blood products are perishable and thus subject to natural wastage. Accordingly, ordering decisions for blood products must be conditioned not only on the amount of stock on hand, but also its age. When the shelf-life of a product is one period, the perishable inventory problem reduces to that of the well-known news-vendor problem. When the shelf-life is more than a single period, the problem is much more complicated. The difficulty of solving the general *m*-period problem has led to the development of a vast literature describing solutions to approximate problems [1].

Beliën and Forcé [2] note publication peaks in the 1970s and the 2000s and suggest that the literature has historically been oriented towards ordering policies for a single hospital or single supplier. See [3] or [4] for an excellent review of early work in blood management models and [2] for a more recent survey article. Literature on network planning for a blood supply chain has historically been less well developed. Brodheim and his co-authors describe a number of studies to set inventory levels within a regional blood distribution network [5], [6]. Hesse et al. [7] describe an application of inventory management techniques to platelets in a system in which a centralized blood bank supplies 35 client hospitals. Yegul [8] describes the development of a custom model to evaluate inventory policies within a regional blood network, but describes only a single instance of his model. Lang [9] uses simulation based optimization to set inventory for a two-echelon system consisting of a single supplier and seven hospitals in which transhipment is allowed between hospitals. Blake and Hardy [10] describe the use of a simulation model of a regional network to evaluate changes to its logistics system. Simonetti et al. [11] develop a model of all US blood producers and consumers and use it to test the impact of a shorter shelf-life for red blood cells, but are able only to consider suppliers and consumers in aggregate. Blake et al. [12], also evaluating the impact of a shorter shelf-life for RBC, present a network model for the blood supply chain in Quebec, including multiple suppliers and multiple consumers, but again discuss only a single instance of their model.

While there is a substantial literature on red blood cell inventory modelling, the literature has focused primarily on single supplier/single consumer systems. A number of papers report network planning models, but examples of such models are less common. While there are numerous papers reporting models for general supply chain management, there are few papers describing the development of a reusable modelling approach to a class of regional blood inventory problems and limited evidence of implementation of these models.

Steele et al. [13] define a reusable simulation as one that allows a single model to be applied to multiple systems in a given domain. They describe two types of generic models: composable and configurable. A composable model is one which consists of generic components that are assembled in a unique fashion to represent a specific problem instance. A configurable model is one in with a generic model is created and configured at run-time to represent a specific problem instance. Robinson et al. [14] note that reusability is the "holy grail" of modelling, and covers a spectrum ranging from simple code scavenging to reuse of entire models. Fletcher and Worthington [15] define a spectrum for reusable models that includes four levels of abstraction (i.e. how general the model is) and transportability (i.e. how many sites the model could be used at).

The prevalence of reusable simulation models varies by area of application [16]. Reusable simulations are common in defence [17] and aerospace industries [18], where high costs, long lead times, and a complex operational environment have led to standards for intra-operability and a requirement that vendors create simulation objects to accompany equipment proposals. Other settings where reusable models have been reported include automated material handling systems design [19], semi-conductor manufacturing [16], airport screening [20] and modular housing construction [21]. Generic models, generic modelling frameworks, and even modelling ontologies can be found in general supply chain management, where they have been used for a variety of system designs as well as training purposes [22].

Examples of generic or re-useable simulation models in health care are less evident than in other settings. Mahdavi et al. [23], reviewing the literature in operational planning models for health services, note that, of 4000 recent papers, only about 100 could be classified as re-useable. Simulation based studies were observed to have been created primarily to represent single units or departments within a hospital for the purposes of comparing different operational plans. Less than 10% of the reusable models in the Mahdavi et al. review provided evidence of implementation.

Re-use is typically driven by the cost of model development. Pidd, in [14], notes that re-usability is expensive and that, to be effective, the costs of model development must be spread over a large number of instances. Thus, development expenses are likely borne by one agency, while the benefits accrue to another. This, it is argued, it has made it difficult to develop reusable models in the health care industry. As a result of the lack of generic representations and the cost and expertise required to create custom models, simulation based methods have not permeated health care, despite the obvious need for, and benefit of, such methods; Brailsford et al. [24] note that, despite the thousands of simulation studies appearing in the literature over the past 40 years, few report implementation or sustained adoption.

Fletcher and Worthington [15] argue that to make models portable it is necessary to reduce the level of abstraction dealt with in the simulation; reusability, accordingly, centres on establishing the appropriate balance between model portability and level of abstraction. Gunal [25] notes that the complexity necessary to make a model "properly parametrised" and thus transferrable between hospitals may actually present a barrier to adoption by end users. He argues for a structured approach to model development, software representation, and identification of appropriate model scope and transferability. Bell et al. [26] suggest the development of inter-operability for simulation objects through a formal engineered ontology [26]. Mustafee and Taylor [27] describe the potential for combining commercial simulation packages with grid computing infrastructure. Grid computing, they argue, would support re-use by enabling complex models to be executed quickly, facilitate re-use of model components, and foster collaboration and joint project development. However, further research is required