Improving Outpatient Non-Oncology Infusion through Centralization and Scheduling Heuristics

by

Adam Ryan Marshall
B.S.M.E, University of Oklahoma, Norman, 2004

Submitted to the Department of Mechanical Engineering and the MIT Sloan School of Management in partial fulfillment of the requirements for the degrees of
Master of Science in Mechanical Engineering
and
Master of Business Administration
in conjunction with the Leaders for Global Operations Program at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2016

Author .... Signature redacted ....
Department of Mechanical Engineering and the MIT Sloan School of Management
May 6, 2016

Certified by .... Signature redacted ....
Retsef Levi, Thesis Supervisor
J. Spencer Standish Professor of Management, Professor of Operations Management,
MIT Sloan School of Management

Certified by .... Signature redacted ....
David Simchi Levi, Thesis Supervisor
Professor of ~~Engineering~~ Systems, Institute ~~for Data,~~ Systems, and Society,
Department of Civil Engineering

Certified by .... Signature redacted ....
Brian Anthony, Thesis Reader
Director, Master of Engineering in Manufacturing Program,
Department of Mechanical Engineering

Approved by .... Signature redacted ....
Rohan Abeyaratne
Chairman, Committee on Graduate Students, Department of Mechanical Engineering

Approved by .... Signature redacted ....
Maura Herson
Director, MBA Program, MIT Sloan School of Management

THIS PAGE INTENTIONALLY LEFT BLANK

Improving Outpatient Non-Oncology Infusion through Centralization and Scheduling Heuristics

by

Adam Ryan Marshall

Submitted to the MIT Sloan School of Management and the Department of Mechanical Engineering on May 6, 2016, in partial fulfillment of the requirements for the degrees of
Master of Business Administration
and
Master of Science in Mechanical Engineering

Abstract

The use of highly effective intravenously infused specialty drugs has increased significantly over the past two decades as they have led to dramatic improvements in patients' quality-of-life. At Massachusetts General Hospital, these drugs are administered in ten independent outpatient clinics. While some clinics only need to offer sporadic treatments and have low utilization of resources, other clinics find patient access is severely limited due to high utilization, poor scheduling practices, and inadequate staffing. This thesis describes methods to increase patient access to infusion while improving resource utilization.

Underlying this improvement is a specially developed scheduling algorithm that smooths chair utilization while permitting flexible, multi-day scheduling. By employing the new scheduling algorithm, the recommended centralized infusion unit will be able to provide more expedient care, offer emergent appointments, avoid unnecessary hospital infusion admissions, and make more efficient use of clinical resources. Adding only two days of flexibility to appointments reduces resource requirements by up to 57%. Also, the day-to-day variability in patient volume is stabilized. Finally, the centralization of administrative resources ensures efficient prior authorization processing, leading to significant financial savings.

Thesis Supervisor: Retsef Levi
Title: J. Spencer Standish Professor of Management, Professor of Operations Management, MIT Sloan School of Management

Thesis Supervisor: David Simchi Levi
Title: Professor of Engineering Systems, Institute for Data, Systems, and Society, Department of Civil Engineering

Thesis Reader: Brian Anthony
Title: Director, Master of Engineering in Manufacturing Program, Department of Mechanical Engineering

Acknowledgments

The work described in this thesis would not have been possible without the support and contribution of the incredible team that comprises the MGH-MIT collaboration. Professor Retsef Levi and Dr. John Stone, thank you for your guidance throughout this process, for laying the groundwork that made this project successful, and for identifying challenges in health care that, when resolved have a meaningful impact upon patient's lives. Dr. Peter Dunn and Bethany Daily, thank you for your support, and for providing an environment within a complex organization where students can both contribute and learn. Kimia Ghobadi and Cecilia Zenteno, you deserve much of the credit for the success of this project. Thank you for your tireless assistance, guidance, and technical support. Professors David Simchi-Levi and Brian Anthony, your technical input throughout the project was valuable and appreciated. All of my LGO classmates and fellow MGH interns, thanks for your friendship, collaboration, and for never leaving a question unanswered. I would also like to acknowledge the Leaders for Global Operations Program for supporting this work. Finally, I would not have been able to participate in LGO without the never ending support and encouragement provided to me by family, especially my wife Kristin and my daughter Whitney.

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Background

1.1.1 Massachusetts General Hospital

Massachusetts General Hospital (MGH), founded in 1861 in Boston, Massachusetts, is widely regarded as one of the leading health care institutions in the United States. A teaching hospital associated with Harvard Medical School, MGH is a pioneer in medical research, and is credited with many revolutionary surgical procedures, including the development of general anesthesia. With 999 beds, the hospital treats over 48,000 inpatients per year. Additionally, MGH serves over 1.5 million outpatients annually, in clinics located across several specialized treatment facilities[1].

For over ten years, MGH has collaborated with the Massachusetts Institute of Technology (MIT) Sloan School of Management, and starting in 2011 has functioned as a partner organization with the Leaders for Global Operations Program (LGO). This research partnership has led to many operational improvements at the hospital, especially in the areas of surgical inpatient flow and bed management[2][3][4][5]. Based upon that success, projects within the MGH-MIT collaboration began to apply the same rigorous analytical and operational research methodologies to outpatient care[6]. Recently, efforts have focused on the improvement of oncology infusion scheduling, with the aim of balancing resource allocation within the MGH Cancer Center[7]. This thesis describes the second MGH-MIT collaborative project targeting outpatient infusion care: the improvement of non-oncology drug infusion.

1.1.2 Non-Oncology Drug Infusion

Non-oncology infusion (infusion) refers to the intravenous administration of drugs to treat conditions other than cancer. Many infusion therapies are life-saving or life-altering treatments, and often involve the use of highly specialized biologic drugs that target conditions such as rheumatoid arthritis, Crohn's Disease, Multiple Sclerosis, and the side effects associated with organ transplants. Infusion appointments range anywhere from 30 minutes to eight hours, with two hour appointments being the most common. Patients may only need one treatment, but often must return to the infusion clinic at regular intervals, sometimes for life. While most appointments are scheduled well in advance, some are urgent and require expedient access to resources.

In 2014, 2,463 patients received infusion requiring more than 18,412 scheduled hours of in-clinic time. At MGH, more than two dozen infusion types are administered in at least ten outpatient specialty clinics with a total of 29 infusion chairs and no fewer than 15 nurses. The MGH infusion clinics are not organizationally aligned, and are scattered across multiple medical services. While some clinics may only need to offer sporadic treatments and have low utilization of resources, other clinics find patient access is severely limited due to high utilization, poor scheduling practices, and inadequate staffing.

1.2 Project Overview

This project[1] investigated operational issues existing within the non-oncology infusion clinics with the goal of developing interventions to (i) improve patient access to treatment and quality of care; (ii) reduce the patient wait time for first infusion appointments; (iii) provide options for same day and emergent treatment (iv) optimally use outpatient and inpatient clinical and administrative resources.

[1]This project was carried out within the framework of an IRB-approved study MIT Protocol #12010014856, "MGH-MIT Collaboration: Surgical Inpatient Flow" Principal Investigator: Retsef Levi. MGH Protocol #2011P001124, "MGH-MIT Collaboration: Inpatient and Ambulatory Patient Flow and Capacity Optimization" Principal Investigator: Peter Dunn, MD.

The project was undertaken in three phases. Initial efforts targeted the development of a detailed understanding of the current state of infusion, discussed in Chapter 4. The second phase included modeling and simulation of potential future state options, described in Chapter 5. Finally, in phase three, recommendations for future clinic organization, location, and administrative resources were conveyed (see Chapter 6).

Phase I: Current State

The current state analysis identified and quantified the resources invested in the institution to provide non-oncology infusion and the operational challenges observed in the clinics. These include patient access issues, inefficient use of resources, and avoidable inpatient infusion that consumes significant amounts of already scarce inpatient beds. Data indicate that patients often wait weeks for their first infusion appointment (mean 51.3 days). Analysis shows this delay is composed of two major contributing factors, one administrative and one scheduling. Administrative delay is related to the significant burden associated with the approval of insurance mandated *prior authorization* paperwork. Scheduling delay is caused by the in-ability to fit a patient into the clinic's constrained schedule. It is also of note that over time, first appointment lead time delay is increasing; from 2012 to 2014 the delay increased by 12.25%. These access challenges are further reflected in the fact that no clinic has urgent (same day) infusion capacity.

While any future system must address both scheduling and administrative delay, the administrative challenge associated with prior authorization is of particular concern. Significant numbers of resources are allocated to prior authorization. While no clinic has a dedicated focal, many have part time staff assigned to authorization processing and related tasks. An estimated number of 15-20 staff spend at least part of their day working on gaining approval for infusion from insurance providers.

Resource allocation and utilization were also studied in this phase. The clinics were found to have widely varying resources that were not appropriately matched with demand. Clinic size ranges from a one chair, part time operation (Dermatology) to an eight chair clinic staffed by three nurses (Neurology). In total, 15 nurses treat patients in 29 infusion chairs. Yet, while each of these clinics described the workload as "busy," analysis of appointment data indicated poor physical utilization of chair resources (range: 18% to 49%, mean: 37%[2]). This discrepancy indicates that nurse staffing, and not necessarily space, is the binding constraint, and that, at the clinic level, there exists a mismatch between supply and demand.

Phase I concluded with a detailed study of avoidable inpatient infusion hospital admission. Avoidable admission occurs when infusion patients are treated (infused) using inpatient resources (hospital beds), not for clinical reasons, but due to access challenges or unavailable appropriate outpatient capacity. For example, if a Rhematology patient requires infusion within three days, but the clinic is fully booked, the patient may be admitted to the hospital in order to receive timely treatment. The analysis suggests that there are at least 497 and up to 1007 inpatient bed days per year consumed by avoidable infusion admission. This is equivalent to 1.4 to 2.8 beds per day, on average. These treatments come at considerable expense to both the patient and the hospital.

Phase II: Future State Models

Phase II of the project shifted into the development of future state models aimed at addressing the operational challenges noted in the infusion clinics. Centralization of resources (pooling) was determined to be a viable means to address the three issues noted. Specifically, centralization should provide more efficient use of resources, which, in turn, frees up capacity and improves patient access.

Two options for centralization at the clinical level, physical and multi-location, were proposed and studied. Physical centralization describes, as the name indicates, the combination of all clinics into a single location. Multi-location centralization involves the use of two or more separate clinics linked by a unified staffing and administrative unit. Key to ensuring efficient use of resources in both of these scenarios is the use of a specialized scheduling algorithm that aims to minimize resource (chair) requirements across both single day (intraday) and multi-day windows (interday). Although heavily modified and improved for this effort, the algorithm is based upon the MGH-MIT Cancer Center project's max(max-min) scheduling heuristic[7]. To further refine the scheduling process, bin packing heuristics were also applied to select between appointments of equal impact (fit). The resulting modeling

[2]The utilization calculation is discussed in detail in Chapter 4.

effort includes twelve physical scenarios, six multi-location scenarios, and twelve sensitivity studies.

Modeling results show that centralization, combined with the improved scheduling processes, is a viable way to improve resource utilization and thus increase patient access to treatment. The addition of only the basic intraday scheduling algorithm, which reschedules appointments within the same day and seeks to minimize resources, reduces the peak chair requirement from 23 to 12, a 48% improvement[3]. Significant reduction in day-to-day variability of the patient census is also observed, with the coefficient of variation changing from a value of 0.74 to 0.26. Figure 1-1 shows a comparison in chair requirements between the current state reference model, and selected intraday and interday scheduling scenarios[4].

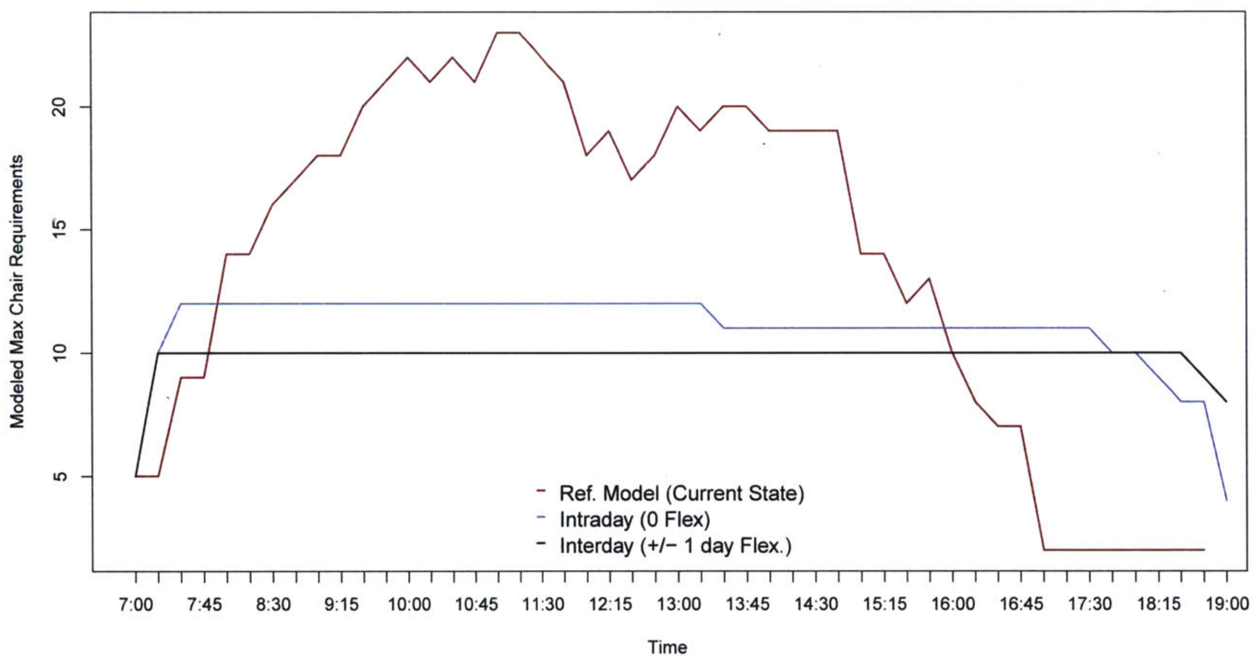

Figure 1-1: Chair Requirement Comparison; Reference Model, Intraday, Interday Scheduling

Most importantly, the modeling effort reveals that the addition of a small amount of flexibility through interday scheduling results in further resource savings and a solution robust to various sensitivity scenarios. The addition of flexibility, only one day prior to and

[3]Scenario P-F-0, detailed in Chapter 5.

[4]The figure shows a comparison between the reference model (current state), an intraday scenario (no flexibility) with the bin packing heuristic set to "first available," and an interday +/- one day flexibility scenario, also with a first available heuristic"

one day after an appointment's original date, yields an additional two chair reduction (ten total required). The day-to-day variability (coefficient of variation of census) is also improved from 0.26 to 0.21. Additional flexibility provides further, although somewhat diminishing, gains. The interday scheduling process is also proved to be robust, and projected resource requirements are valid in the presence of additional constraints such as fixed appointments and patient preference for particular times of day.

Thus, through the application of centralization and advanced scheduling techniques, the peak number of chairs used can be reduced from 23 to 10, a 57% reduction. Nurse staffing levels are lowered from a previous level of 15 to a new minimum level of five, a 67% reduction. Finally, through the improved utilization of capacity, emergent appointments can be accommodated and inpatient infusion significantly reduced or potentially eliminated.

Phase III: Recommendations

The successful modeling effort led to the final phase of the project, the recommendation for future state configuration. In the near term, the partial centralization of resources at two locations (an existing MGH clinic and MGH-West) will allow for improved access to care and growth. The final recommendation, however, is for the creation of new physical clinic on the MGH campus. The new clinic will work in conjunction with the resources at MGH-West, and both will operate with a single staffing unit and be able to accommodate all non-oncology infusion.

Based upon the conservative results of the intraday scenarios[5], the new main-campus clinic shall be composed of 16 chairs. Twelve of the included chairs are directly derived from the modeling results. An additional four chairs are added to account for growth and intraday disruption. The clinic is recommended to be staffed by six nurses, five of which are required per the assumed 2.5:1 patient to nurse ratio maximum. An additional nurse is strongly recommended to allow for breaks and to provide sufficient system slack. To be successful, both the interim and future state clinics must make use of the MGH-MIT scheduling algorithm.

The centralization of all administrative resources, namely prior authorization processing,

[5]Scenarios P-F-O through P-R-0, described in detail in Chapter 5

is also proposed. Following the model of the MGH Cancer Center, a small, three to four person team should be created to process all prior authorizations. The team is projected to be able to process most authorizations in five days or less, greatly improving the lead time for a new infusion appointment.

By pursuing these recommendations, MGH is expected to be able to provide improved access to patients, make more effective use of clinic and physical resources, and realize substantial financial benefit through the avoidance of unnecessary inpatient infusion.

1.3 Thesis Organization and Structure

Following this brief introduction to Massachusetts General Hospital and the infusion project, chapter two of the thesis continues with a literature review summarizing relevant research on infusion practices, administrative processes, scheduling algorithms, and bin-packing heuristics. Next, in chapter three, infusion patient flow, along with administration and scheduling processes are described in detail. Chapter four summarizes the major findings uncovered after analyzing the current state of infusion at MGH; the methods and data used to analyze the current state are also described. Future state modeling approaches, results, and the details of proposed scheduling algorithms are discussed in Chapter 5. Finally, the thesis concludes in Chapter 6 with recommendations for interventions and future clinic configuration to improve infusion processes at the hospital.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Literature Review

The following chapter summarizes selected research in areas relevant to major topics affecting this project. These include non-oncology infusion, the growth of specialty drugs, management practices for infusion clinics, and the impact of prior authorizations. Next, capacity planning and the effects of pooling in service operations is reviewed. Scheduling systems in the health care setting, and common problems that are encountered are also covered. Finally, the chapter concludes with a discussion of bin packing problems and their application to scheduling.

2.1 Infusion and Infusion Clinic Management

Several relevant works discussing the growth of *specialty pharmacy* and the management of infusion were identified [8][9][10][11][12][13]. Specialty pharmacy is a major component of the non-oncology infusion treatments administered in the MGH clinics, described in detail in both Chapters 3 and 4. In an ariticle of particular interest, Patel and Audet[10] describe the incredible growth of specialty pharmacy costs in recent years. Specialty pharmaceuticals, per their definition, are drugs developed to treat "complex, chronic, rare, and difficult to manage conditions." Although some of these drugs, often referred to as *biologics*, are oral or injectable medications, many are developed as intravenous infusion therapies. The authors state that costs of these treatments are increasing, and that in 2013 specialty drugs accounted for 29% of US health care pharmaceutical expenditure. Compared to normal saysmall molecule drugs, where spending only increased by 0.1%, spending on specialty biologics increased by up to 28% in 2013[10].

Demand for these therapies is also increasing, potentially fueled by the fact that by 2018 there will be an additional $13 billion dollar market for biosimilars, that is, biologic generic drugs, due to expiration of multiple patents. As these drugs are often expensive, and can cost upwards of $10,000 dollars per month, many health care payers have instituted mitigation strategies in order to control costs. Patel and Audet reference three major methods that governing bodies and insurers use to limit specialty pharmacy expense, in wake of the rising availability and demand: formulary restrictions, prior authorizations, and quantity limits. Formulary restrictions place constraints on the diagnosis for which a payer will reimburse when a drug is prescribed. Quantity limits place restrictions on the volume of drug that it is reimbursable. Of particular relevance to this project, however, is prior authorization, where payers require prescribing physicians to obtain pre-approval for administering the drug based upon "clinical need".

Freeman, et al.[14] also state that demand for infusion therapies, such as infliximab, has grown, and is expected to continue to grow and "outpace traditional small molecules." The authors indicate that because of the specialized nature of infusion therapies, and their reliance on a medical professional for administration, efficient coordinated care is required to control costs. They discuss a pilot project to study ways to reduce costs during the delivery of infliximab, a common infusion therapy, in southern Ohio Hospitals and Gastroenterology practices. The project involved developing a collaborative team of stakeholders, representing patients, insurers, pharmacy, and infusion providers. Providers, the authors state, often have a choice between administering infusion in their offices, at an infusion center, or in a hospital setting. In this pilot project, cost savings resulted through the avoidance of infusion in the hospital. In order to deliver these savings, however, an integrated care system across multiple stakeholders had to be created and managed.

Infusion clinics are often either specialized and practice-based or multi-specialty[1] centers offering a range of treatments. Foley and Dunne[11] describe considerations important to the

[1]Specialized, practice based centers are located within a clinic that focuses on a certain patient group, such as GI or Rheumatology. Multispecialty clinics accept a wider range of patients.

management of a clinic based infusion practice that are applicable to both of the common settings. Of importance, they claim, are the development of standard operating procedures for infusion, along with payer and reimbursement monitoring. Also discussed are means of staffing such a clinic; the authors state that an "experienced, part time nurse can treat up to three patients" at a time. Finally, the authors mention the criticality of proper scheduling techniques, and that appointments should be "staggered, as to not overwhelm the nursing staff at any point in the day."

While Foley and Dunne focus on clinic based infusion, over 40% of infusions in the United States are performed at multi-specialty centers, and that number is expected to stay stable or grow[13]. Foley and Dunne also often mention the patient benefits of practice-based infusion, namely "integrated patient care and treatment adherence." Ostrov[13], however, provides a counter argument to these claims, in a discussion of patient preference for infusion in a multi-specialty infusion center. The study, conducted at the Penn State Hershey Medical Center, was performed to understand patient preference and use it to properly structure infusion care. Through surveys, data on preference for non-oncology patients receiving treatment in an infusion center were obtained. Responses to a seven question Likert scale were collected from over 70 patients, nurses, and physicians. The data show that nurses and physicians assumed, like Foley and Dunne, that patients would prefer treatment in a non multi-specialty setting. The patient results, however, refuted this assumption.

The patients actually had high levels of satisfaction and preferred treatment in a multi-specialy center over other settings. The most important item of concern to patients, as revealed by the survey, are the "high quality, skilled nurses" responsible for administering their treatment. Infusion centers, the authors claim, are often convenient, and they provide ample time for patients to interface with nurses, who are responsible for administering the bulk of their care. This time spent with skilled health care providers yields a strong sense of satisfaction with the care and service provided. The authors admit that the study only had a limited number of participants, yet they state that the results are of interest to those considering "whether to have separate or unified infusion centers."

2.2 Prior Authorization and Administrative Burden

As discussed, in order to control the rapidly rising costs associated with specialty infusion drugs, payers have been proactive and implemented onerous pre-approval requirements, commonly referred to as *prior authorizations* or PA. The burden associated with these authorizations is present at MGH, and is referenced in Chapters 3 and 4 of this work. The significant administrative burden associated with PA requirements in the US health care system is rigorously documented in popular publications[15][16][17], but only rarely discussed and quantified in academic literature[18]. Bendix[17] discusses prior authorizations, and describes the "fighting with payers" as a "predicament". Authorizations, the author claims, in 2011, were responsible for $69 billion dollars in cost to physicians. This cost is derived from both lost productivity and decreased revenue. Further estimates claim that more than 860 million hours are spent annually by US physicians on authorizations, and that the average physician may spend 20 hours per week on authorization related tasks.

It is not expected, however, that the number of authorizations required will decrease. As such, the author claims, clinics and physicians must be proactive to reduce the associated challenges. Recommended methods to do so include the use of software and technology to track authorizations, and building expertise among staff in authorization requirements and procedures.

Leinss et al.[18] are more descriptive in their recommendations for controlling authorization burden. They describe a "comprehensive medication prior authorization service" implemented at Froedert Hospital in Milwaukee, Wisconsin. The service, the authors claim, resulted in significant savings due to a decrease in lost charges and an increase in reimbursement. The new service also has helped to facilitate "closer working relationships" between the various clinics and administrative organizations.

Key to the success of this initiative, Leinss states, was the "expertise and experience" brought to the program through the hiring of a full time Medicare billing and coding specialist. The specialist focused on pilot projects in clinics with high numbers of reimbursement write-offs due to specialty medications. Example clinics provided included Gastroenterology, Rheumatology, and Neurology. The "sheer volume" of specialty medications used by these clinics led to significant numbers of write-offs, which were reduced through the expertise and dedicated attention of the authorization focal. As the initiative expanded the centralized authorization staff grew, yet in only six months, the project was able to decrease write-offs by over $6.2 million dollars, a claimed 20:1 return on investment.

Additionally, Leinss claims, as the authorization staff continued to work with various hospital based clinics, they were able to assist in the remedy of other operational inefficiencies. Staff began to track medication inventory, and ensured that it is available for incoming patients. Also, the PA team has been successful at enrolling patients in co-payment and medication assistance programs, which are both of help to patients and increase hospital revenue.

2.3 Capacity Planning and Resource Pooling in Service Operations

One of the major outcomes of this project is the development of the required resources for a proposed MGH centralized infusion clinic, a summary of which is contained in Chapter 6. Methods for planning capacity in service operations, and the benefits of resource pooling (centralization), especially in healthcare are well documented in academic literature [19][20][21][22]. In his well known discussion of capacity planning for service operations, Sasser [23] discusses two strategies for the manager: (i) chasing demand or (ii) providing level capacity. When a manager chases demand, historical data is used to forecast resource requirements in the short term. Chase strategies are most suited to low skill, low compensation services. In contrast, level capacity, where the level of resources is set in advance and held steady, is most aligned for services where high labor skill is required, error rate is needed to be low, and the amount of job discretion is high. He also mentions several strategies to manage demand and supply in a service. Of particular relevance to this project is the suggestion of sharing capacity, a type of centralization.

Centralization of resources in health care is common, and even the anecdotal example given by Sasser[23] in his article is of a hospital sharing equipment with a nearby institution. As mentioned, multitude of successful examples of resource pooling (centralization) in health care exist. The prior authorization centralization discussed previously[18] is one such example. Additional cases include improved pancreatic cancer patient outcomes in Scotland[20], and a 1.2% decrease in colon cancer mortality due to the high volumes offered by a surgical center of excellence[21]. The benefits of centralization are typically attributed to the "reduction in variability due to the portfolio effect" [24].

However, when centralizing resources in a health care setting, diligence must be taken to ensure that the outcome is beneficial to the patient and the system. Vanberkel et al.[25] point out that there are trade offs that must be made, and that sometimes a decentralized system may make sense for patients and managers. The authors construct a complex queuing model to evaluate de-centralization. The model takes into account clinic load, size, appointment lengths, appointment variability, and several other factors. In summary, the authors claim un-pooled resources may make sense for highly specialized patient groups. If choosing to de-centralize, they advise managers to be cognizant that access may decrease unless service time in the de-pooled clinic is also decreased.

2.4 Appointment Scheduling in Health care

As discussed in Chapter 5 of this work, in addition to addressing the administrative issues associated with infusion, this project seeks to develop a robust appointment scheduling process. While health care scheduling is similar to problems in general production planning, there are specific concerns and nuances that must be addressed. The literature on appointment scheduling in health care is rich [26][27][28]. Gupta and Denton[26] provides an excellent overview of appointment scheduling methods and challenges in the health care setting. Appointment scheduling systems, the authors state, "lie at the intersection of efficiency and timely access to health services." Successful scheduling systems are described as those that provide access to emergent treatment (same day appointments), without impacting previously scheduled patients.

A "critical bottleneck" of many appointment systems, Gupta claims, is a lack of the application of operations research techniques. Other challenges include questions on how to account for no-show patients, cancellations, and means of dealing with arrival variability. Additionally, many appointment systems lack any consideration for patient preference, due to the inherent complexities associated with quantification.

The work of Rieb[7] presents an appointment system for health care that includes both operations research techniques and the incorporation of patient preference in an outpatient Cancer Center infusion unit. Rieb's work, which this thesis leverages and extends, uses linear programming to identify an optimal scheduling solution for historical appointment data. Based on the results, she then proposes a scheduling heuristic (a practical, rules based approach) which she evaluates via simulation with that same existing data.

Her results show that when the heuristic method is applied, it performs nearly as well as the optimal offline solution, and the average infusion resource requirements (chairs) are reduced by up to one-third. Finally, the work also ensures that the proposed scheduling system is robust, not only from an operations research standpoint, but also from a patient perspective. The scheduling process includes the ability to incorporate patient preference for appointment time and specific nurse availability.

Rieb's Cancer Center project also demonstrates an additional key concept of relevance to this work. By combining a scheduling heuristic and historical appointment data, a model of a proposed clinical configuration is created. The use of this model can lead to relevant, reasonable insights on the expected real world performance of the appointment scheduling system. That approach is pursued in this work, and it yields similar insights and successes.

2.5 Bin-Packing and Scheduling

This project also set out to improve the scheduling process implemented in the aforementioned Cancer Center project. One method to achieve this goal was through the application of modified *bin packing* heuristics to the scheduling algorithm. Bin packing is a process by which pieces are fit into a minimum number of spaces, or bins. The details of the coupling of the infusion center scheduling heuristics and bin packing are discussed in Chapter 5.

Bin packing problems are common in many industries including manufacturing, where solutions are often sought to efficiently utilize raw materials, a process called trim loss minimization or nesting[29]. Additional applications include production scheduling, where the processes can be used to minimize production make-span[30], and in other general scheduling problems[31][32][33][34]. In health care, related methods are often used to book operating rooms[35].

In an overview of bin packing and scheduling, Renault[30] states that bin packing algorithms are typically on-line processes, that is they receive their input "one piece at a time" and must make "irreversible decisions on the processing on the current piece" with no information on future requests. A commonly used bin packing heuristic is the first fit method, where an arriving piece is placed into the first "available" bin.

Although no non-oncology infusion specific scheduling algorithms were noted in academic literature, Tanaka evaluates the use of the first fit and multiple other bin packing algorithms in an oncology infusion setting. The goal of his scheduling project was the reduction in-clinic patient wait time. [34]. Although relevant to this project, his work differs in that he uses bin-packing to schedule directly to specific infusion chairs, rather than to smooth overall resource requirements. Similar to the MGH Cancer Center Project, Tanaka employs a simulation and modeling approach to test the validity of his selected bin packing scheduling algorithms. He tests multiple bin packing options, including the aforementioned first fit, along with several other specially modified heuristics. His work finds that by applying modified bin packing to his constrained scheduling problem, infusion patient wait time can be minimized.

Chapter 3

Non-Oncology Infusion at Massachusetts General Hospital

Infusion is administered in several outpatient clinics across the MGH campus. Each clinic is affiliated with a department or division of the hospital specializing in the care of certain diseases or conditions. For example, the Rheumatology Infusion Clinic, physically located within the Rheumatology Unit, is part of the Rheumatology, Allergy, and Immunology Division. Combined, the clinics serve over 13,000 appointments annually. These appointments account for nearly 23,000 hours of in-clinic infusion procedures administered to nearly 4,000 unique patients. This chapter describes the resources and business processes that exist in a typical infusion clinic, and summarizes the specific resources at the MGH clinics included as part of this study. Chapter 4 delves further into appointment data analysis and the issues uncovered in the current state of the system. The information described in both chapters was collected via interviews with clinical and administrative experts, shadowing at each of the MGH infusion clinics, and appointment data analysis.

3.1 Methods and Approach

3.1.1 Shadowing and Interviews

Each of the ten MGH infusion clinics was visited, and subject matter experts were interviewed and observed. In total, approximately 40 hours of interviews and shadowing were conducted. At a minimum, interviews were held with key staff members, including infusion nurses, prior authorization focals, and appointment schedulers. Often, the nurse manager, practice manager, and physician director of each clinic were also interviewed. Appendix E contains a sample questionnaire used at each of the interviews. The roles of each of these staff members is discussed later in this chapter.

In addition to the MGH infusion clinics, an electronic request for data on multi-specialty infusion clinics at other hospitals was submitted to the University Health Consortium (UHC), an affiliation of university associated research hospitals. Of the 11 hospitals that responded to the UHC inquiry, four were selected and interviewed. These interviews were conducted similarly to the internal MGH discussions, and the data act as a baseline to compare MGH to the best practices of other leading hospitals.

3.1.2 Data Sources

To support detailed current state analysis and future state modeling, appointment data was extracted from the hospital's electronic records systems. From 2009 until July 2014 the appointment record system in use for the outpatient clinics was IDX. After July 2014, the clinics began a transition to a new system, EPIC. Data was extracted from each system. Inclusive of both IDX and EPIC records, more than one million appointment records were extracted and parsed. For detailed modeling, only the period from July 2014 to July 2015 was used, a data set comprised of 14,932 infusion appointments across all the included clinics[1]. Appendix A includes a complete list of the EPIC infusion codes used in this study. Additional databases were also employed for minor supporting studies, including inpatient records (EPIC, EPSi), pharmacy coding systems (SunQuest), and administrative records

[1] The included clinics, discussed below, are all hospital based. Infusion at external, hospital-associated clinics (private practices) is not included.

(Huron).

It is also of note that much of the appointment data, prior to July 2014, is held in an older system named IDX. In order to analyze this legacy data and provide meaningful results, additional axillary information, such as appointment length and drug names, had to be gathered and integrated into a new, comprehensive non-oncology appointment data-set (time frame 2009-July 2014). Due to the unstructured nature of the IDX data, along with the lack of standardization in data entry and terminology across clinics, this integration required substantial effort. Post IDX appointment data, extracted from EPIC starting in July 2014, is substantially more structured and can be analyzed with little-to-no augmentation.

3.2 Patient Volume and Demand

From 2009 to 2014, the number of unique patients receiving infusion at MGH grew from 1836 to 2463, an increase of 34%. The chair hours required to fulfill these infusion appointments also increased, growing from 15,106 to 18,412, an increase of 20% over this same five year period. Figure 3-1 depicts this growth. It is also of note that the length of infusion appointments, the driver of the chair hours required, changes over time. As new drugs come onto the market and others are more widely used, the length of time necessary to infusion a drug may change. Figure 3-2 shows a standard box plot of the schedule appointment time distributions, by infusion clinic, from July 2014 to July 2015. The average appointment length across all clinics was 1.72 hours, the standard deviation 1.43 hours, and the range 0.25 to 8.0 hours. Appointment scheduled times are inclusive of the setup and clean-up procedures described in chapter 3. Interviews and shadowing also revealed that infusion appointments rarely over-run their scheduled time interval.

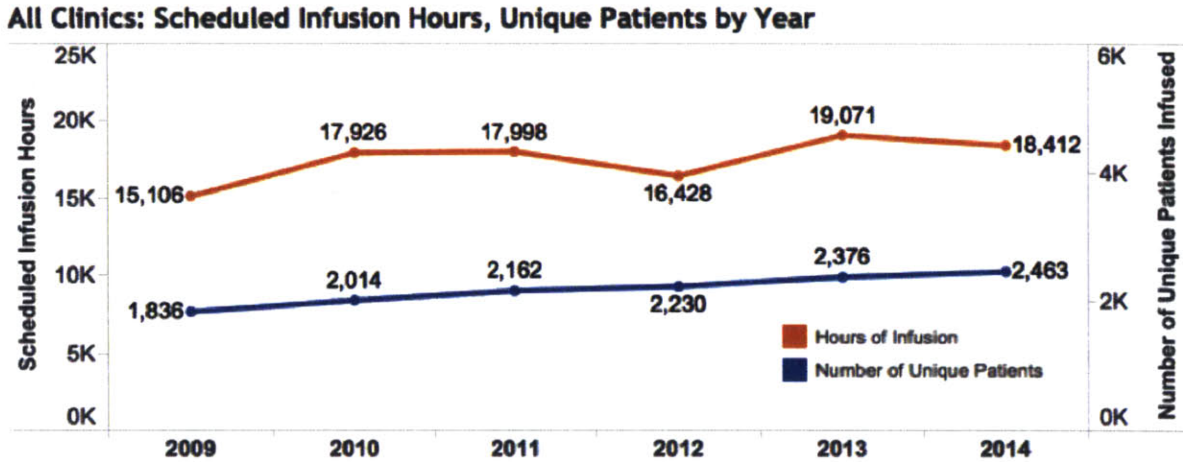

Figure 3-1: Infusion Demand, 2009-2014

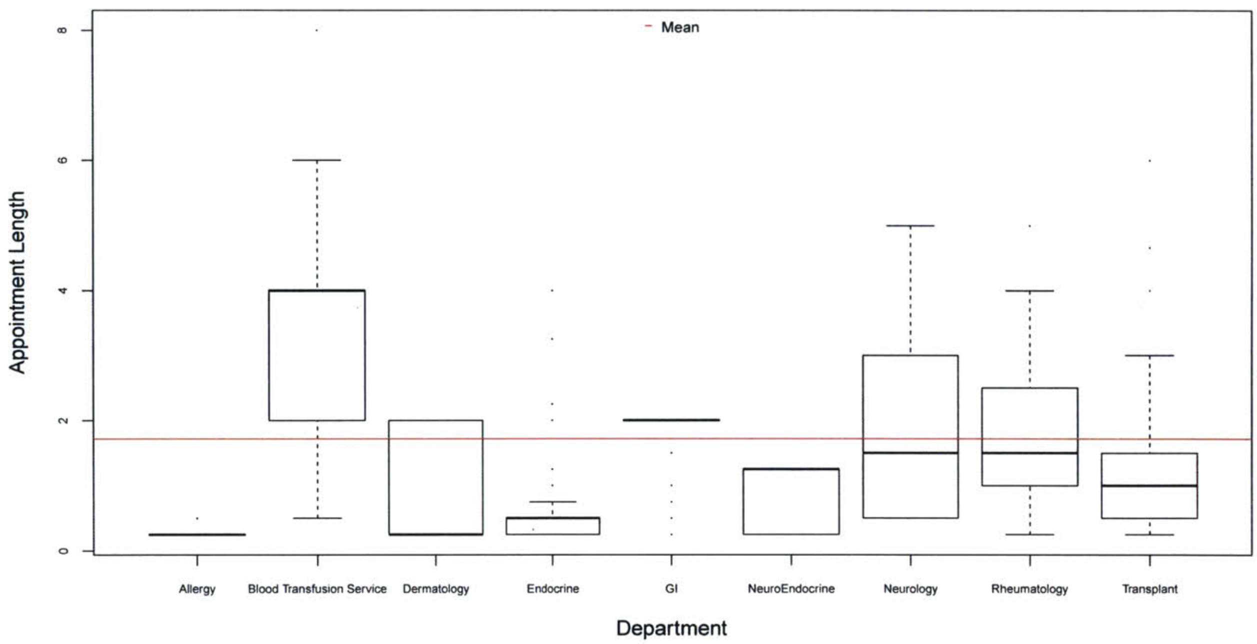

Figure 3-2: Infusion Appointment Length Distribution, July 2014-July 2015

The demand analysis includes all appointments that occurred in the time period of interest, along with any appointments that were "no-show," that is, instances the patient did not arrive or contact the clinic ahead of time to cancel. It is assumed that no-show appointment slots are not possible to be filled, and thus have the same impact upon clinical resources as a normal, completed appointment. Cancellations are excluded, as it is assumed that those appointments are rescheduled during period of study. The complete list of infusion appointments, organized by clinic, is available in Appendix A.

3.3 Drugs, Treatments and Appointment Types

Infusion Drugs and Treatments

The infusion clinics at MGH administer dozens of drugs and other treatments. The three most common appointments types are (i) IVIg (Intravenous Immunoglobin), (ii) Infliximab, and (iii) Rituximab. Although these treatments encapsulate the majority of infusion (as measured by hours of scheduled appointments), there are many additional drugs and treatments demanded in lower appointment volumes. More than seventy different appointment type codes are used to describe infusion at the ten MGH clinics. Appendix A contains a complete listing of these appointment codes, many of which are descriptive of the drug or process name. Note that in this work, infusion appointments may be referred to as providing a drug, treatment, or therapy, interchangeably.

The top ten infusion types in 2014, by hours scheduled, are shown in figure 3-3. More than 76% of scheduled infusion hours are dedicated to three drugs: IVIG, Remicade (infliximab), and Rituxan (rituximab). The remaining quarter of the scheduled hours of infusion is comprised of a long tail composed of nearly two dozen different treatments. Figure 3-4 expands upon 3-3 and shows all of the infusion appointment types[2].

It is also of note that many of the infusion drugs are administered in multiple clinics. IVIG, the most common treatment, is administered in two clinics (Neurology and GI). Infliximab is given in three clinics (Endocrine, GI, Rheumatology), while rituximab is available in two (Neurology and Rheumatology). Rheumatology and Neurology have the widest range of available infusion types, with 14 and eight different drugs, respectively.

[2]Transfusions are not included in this analysis. Transplant, Dermatology, and Allergy are not included due how those clinics encode appointment data.

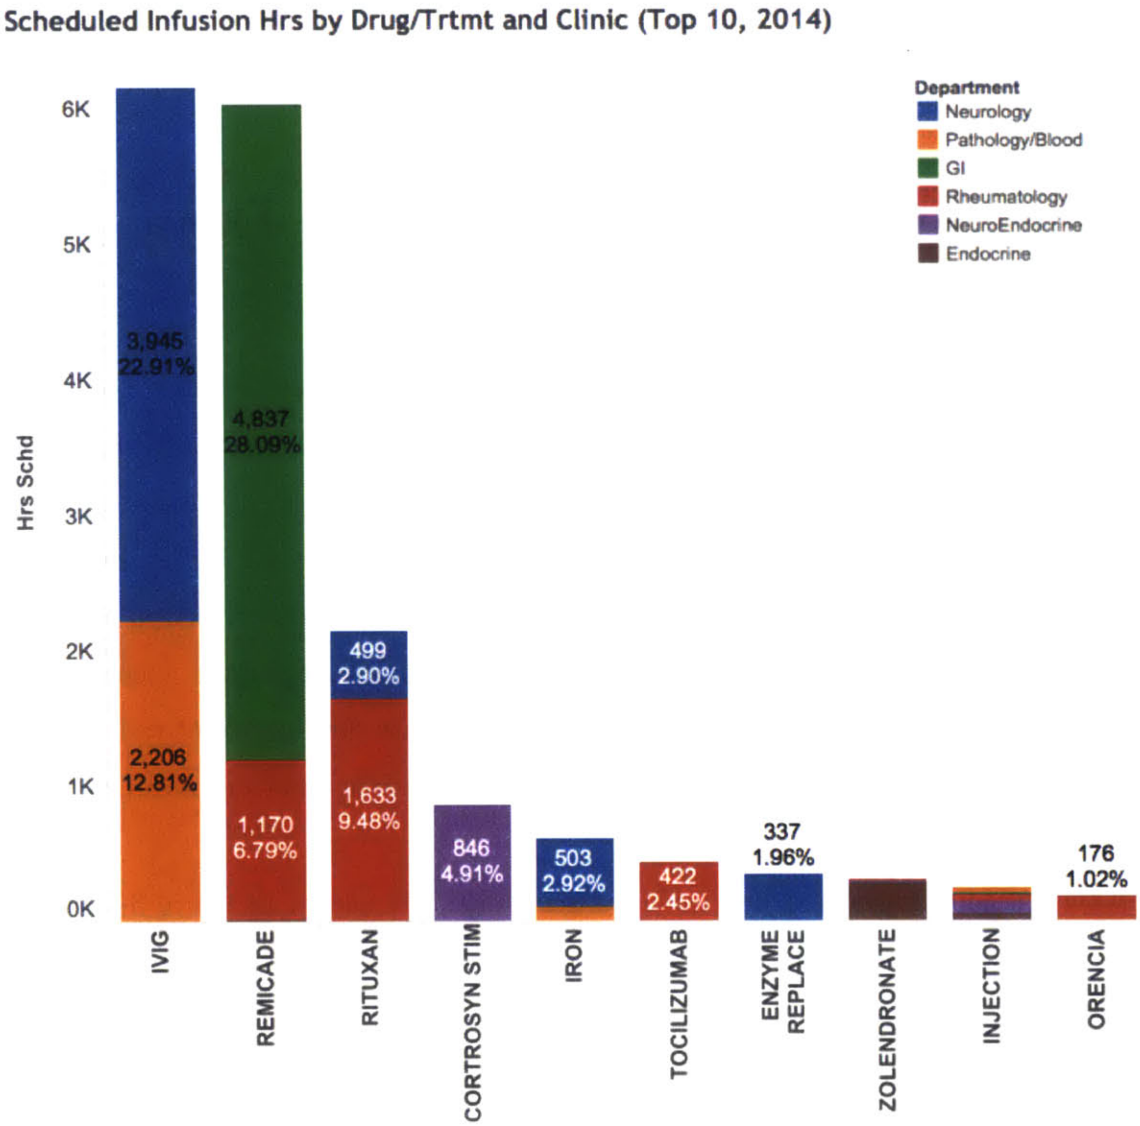

Figure 3-3: Top Ten Infusion Appointment Types, by Hours Scheduled, 2014

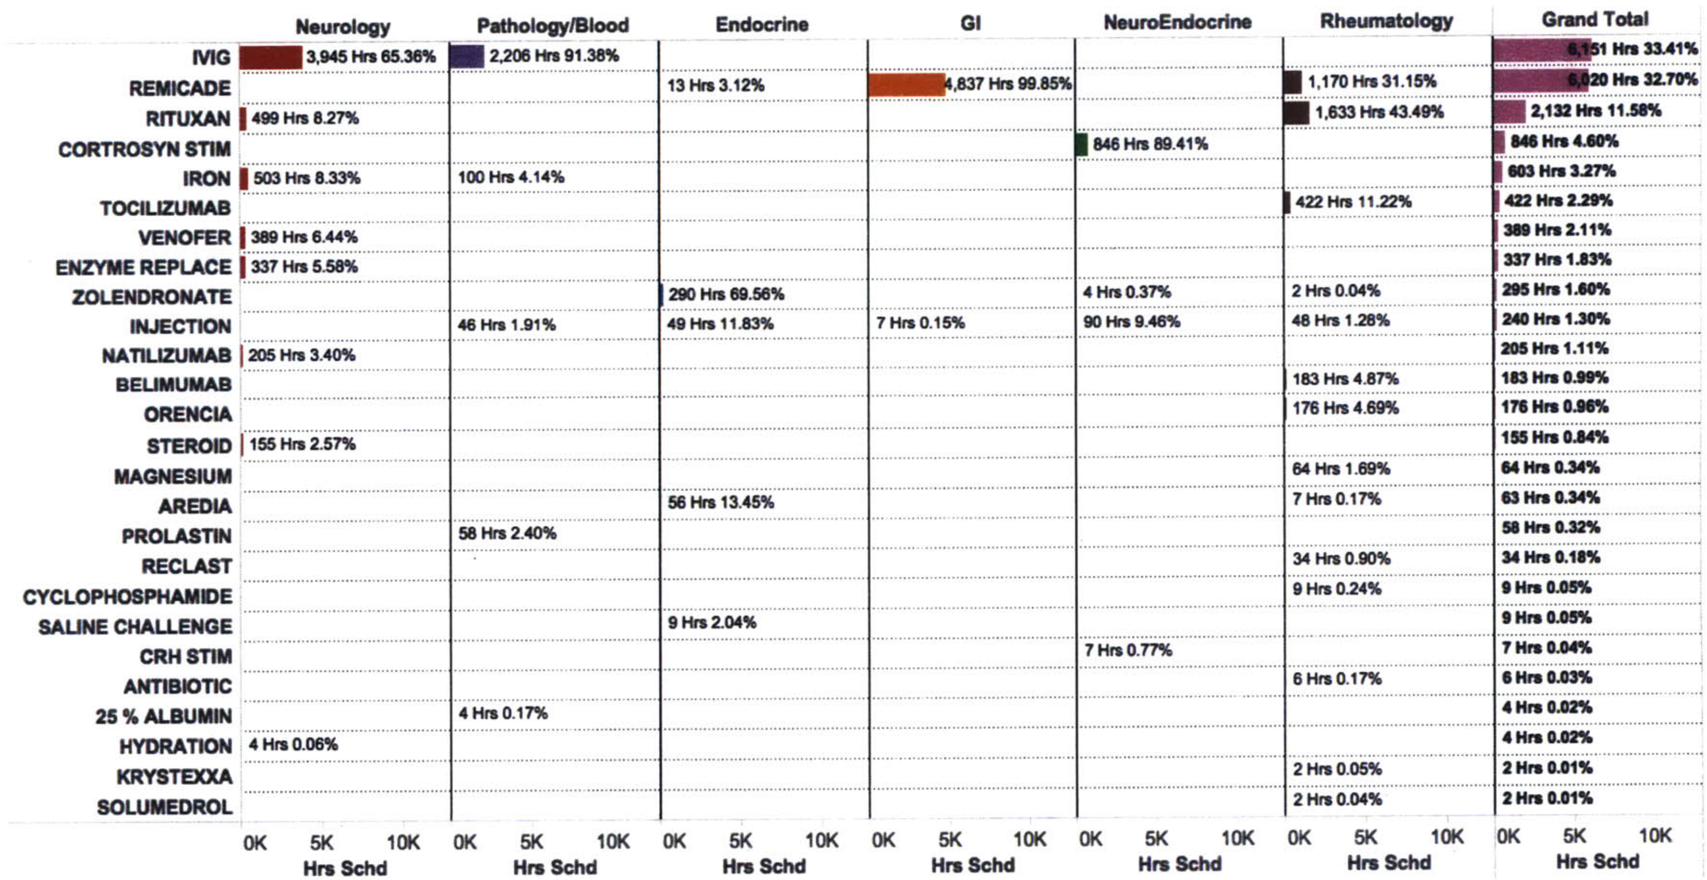

Figure 3-4: Infusion Drug Hours Scheduled, by Clinic, 2014

3.3.1 Appointment Types

The term "infusion" is used broadly to describe all of the appointments that occur in an infusion clinic. There are, however, several additional types of infusion appointments and other treatments that occur in the clinics.

- **Regular Infusions:** typical infusion appointments with standard preparatory processes. The patient arrives, vitals are checked, an IV is started, and the infusion is administered.

- **First Infusions:** Due to the potential for adverse reactions to drugs, first infusions are often scheduled for a longer period of time. This allows the drug to be administered more slowly, and provides the nurses more time to monitor the patient for any issues.

- **Injections:** Some patients receive an injection of a drug rather than an infusion. Often self injected at the patient's home, the first injection is sometimes required to be administered in a medical setting. Additionally, some patients prefer to have a nurse perform an injection rather than self-administer.

- **Diagnostic Tests:** Infusion resources (chairs and nurses) are sometimes used to perform medical diagnostic tests. Drugs may or may not be infused in a diagnostic test.

- **Transfusions:** Similar to an infusion appointment except that blood or blood products are infused (transfused).

- **Referrals:** Although uncommon, some MGH clinics have instituted ad-hoc processes that allow them to refer patients to other clinics for infusion. Referrals are most often made when a clinic lacks the resources to infuse a patient, of if the referring clinic is too busy to serve the patient in an expedient manner. Sometimes, referrals are made to external, non-MGH affiliated clinics.

- **Urgent Infusion:** Occasionally, patients will need to start infusion therapy within a short time frame, often one to three days.

- **Miscellaneous:** There are several other types of appointments in a typical infusion clinic including patient education and IV port cleaning.

- **Coupled Appointments:** Occasionally, patients are scheduled to see their medical provider before or after an infusion appointment. This is common (and often required) in oncology settings. In non-oncology infusion, coupled appointments are often employed for patient convenience, but they are neither common nor strictly necessary.

3.4 MGH Infusion Clinic Overview

The MGH Infusion clinics can be typically described by four key characteristics: (i) the number of chairs physically available for infusion, (ii) the number of nurses caring for patients in those chairs, (iii) the types of drugs and treatments administered, and (iv) the roles and responsibilities of administrative support staff.

3.4.1 Infusion Chairs

Patients receive infusion treatment in chairs similar to the rendering in figure 3-5. The chairs may be in private, as depicted in the figure, semi-private with curtains or other temporary dividers, or open to a communal space. Due to the length of the infusion appointments, chairs are designed to recline or convert fully into a flat bed. As outpatient infusion recipients are ambulatory, it is possible for them to move about the clinic, with infusion pump attached, in order to use the restroom or socialize with other patients.

3.4.2 Infusion Nurses

Infusion nurses are the providers of care to patients in an infusion center. It is common for nurses to be the only medical staff present in an infusion clinic, as physicians are often only available on-call, on an as-needed basis. Physicians often provide organizational and clinic oversight, but nurses manage and operate the infusion clinic. Infusion nurses are specially trained in the administration of intravenous drugs and, due to the similarities with oncology infusion and the use of chemotherapy drugs for non-oncological purposes, some are oncology certified. Nurses are typically not dedicated to patients in a one on one ratio; other than during the initial appointment set-up processes that include checking vital signs and starting

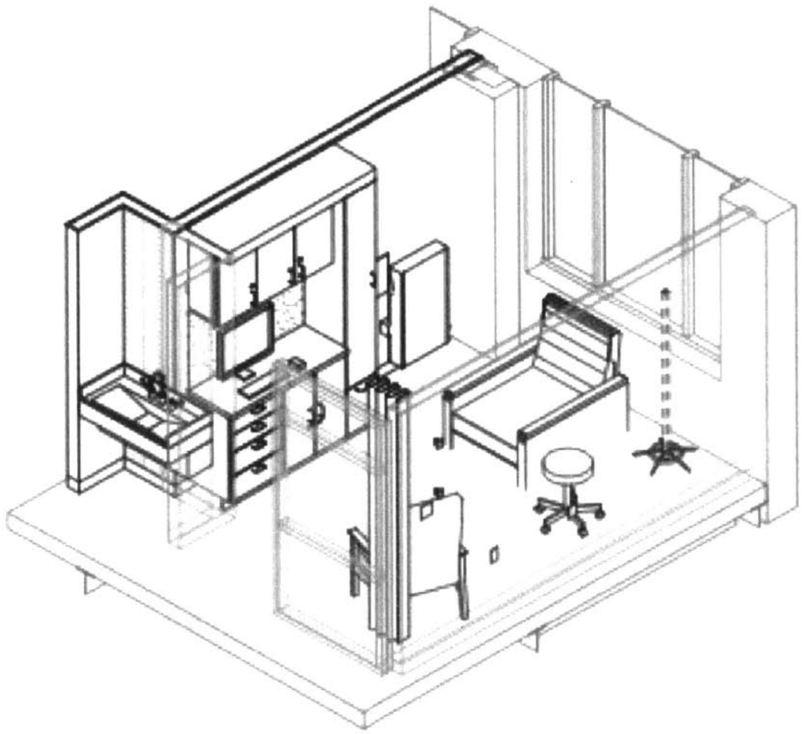

Figure 3-5: Typical Private Infusion Room[36].

the IV. After these setup processes are completed and infusion commences, nurses can care for patients in multiple chairs.

3.4.3 MGH Infusion Clinics

The ten MGH infusion clinics, listed below, range in size from a single chair open one day per week (Dermatology) to a large, eight chair unit staffed by several nurses (Neurology). A brief description of the various infusion units and their resources is given below and summarized in Table 3.1.

- **Neurology:** An eight chair infusion center, staffed by three nurses. Open Monday-Friday, 7:00am to 5:00pm. The most common infusion therapy administered is intravenous immunoglobulin (IVIg).

- **Gastroenterology (GI):** A five chair infusion center, staffed by two nurses. Open

Monday-Friday, 7:30am to 4:30pm. The most common infusion therapy administered is infliximaub.

- **Rheumatology:** A three chair infusion center staffed by a single nurse. Open Monday-Friday, 7:00am to 5:00pm. The most common infusion therapy administered is rituximab.

- **Blood Transfusion Service:** A combined infusion and blood transfusion clinic, with eight total chairs and beds, five of which are available for infusion. The clinic is staffed by five nurses, and open Monday to Friday, 7:00am to 3:30pm. The most common non-blood transfusion treatment is IVIg.

- **Endocrine:** A clinic with two chairs staffed by a single nurse limited to 32 hours per week. One of the two chairs is scheduled for only half of each working day. The clinic is open Monday to Thursday, 8:00am to 4:00pm. The clinic uses the chairs for both infusion therapies and diagnostic testing. The most common infusion administered is zolendronate.

- **Neuro-Endocrine:** A clinic with two chairs staffed by a single nurse limited to 32 hours per week. Open Tuesday to Friday, 9:00am to 3:00pm. The clinic uses the chairs for both infusion therapies and diagnostic testing. The most common use of the chairs is for cortosyn stim testing.

- **Transplant:** A clinic with two chairs staffed by a single nurse. Open Monday to Friday, 8:00am to 4:00pm. The most common treatment is rituximab.

- **Dermatology:** A clinic with a single chair, staffed by one nurse. Open one day per week from 9:00am to 5:00pm. The most common treatment is rituximab.

- **Allergy:** The clinic has demand for infusions but no capacity. Patients are referred to other clinics, most often Neurology, for infusion.

- **Infectious Disease:** The clinic has demand for infusions (antibiotics) but no capacity. Patients are referred to other clinics or admitted to the hospital for infusion.

- **Nephrology:** The clinic is an off campus, seven chair physician managed practice. Patients may be refereed by the managing physician, or by other internal, hospital managed practices[3]. The clinic performs mostly rituximab infusions.

- **Other Clinics:** Outpatient infusion capacity also exists in the Pediatric Gastroenterology Clinic (Pedi-GI), the Massachusetts Eye and Ear Infirmary (MEEI), the MGH Urgent Care clinic, and at at MGH-West[4]. These clinics are not part of this study. [5]

Table 3.1: MGH Infusion Resource Summary

Clinic	Nurses	Chairs	Chair-hours/Week
Neurology	3	8	400
Gastroenterology	2	5	225
Rheumatology	1	3	150
Blood Transfusion Service	5	6	255
Transplant	1	2	80
Dermatology	1	1	8
Endocrine	1	2	48
Neuro-Endocrine	1	2	48
Allergy	NA	NA	NA
Infectious Disease	NA	NA	NA
MGH Total	**15**	**29**	**1214**

Physical chair hours per week, H_c, can be calculated using 3.1 given below:

$$H_c = C_{clinic} \times O_{hours}, \tag{3.1}$$

where C_{clinic} is the number of chairs in the clinic, and O_{hours} is the total number of hours per week the clinic is open. While the clinics each added capacity and evolved, they did not standardize processes. None of the clinics have standardized the inter-clinic referral

[3]All of the MGH infusion clinics discussed are hospital managed practices. The financial structure of reimbursement for a physician managed clinic is different than a hospital managed practice, and the Nephrology clinic claims to offer patients a lower cost for treatment. Patients that are referred to this clinic from the MGH system are likely to receive ongoing treatment at Nephrology, resulting in significant lost revenue for the hospital

[4]An eight chair infusion unit is planned for the MGH-West campus in Waltham, Massachusetts. The clinic is affiliated with the Rheumatology division of MGH.

[5]These clinics are excluded from study either due to organizational bounds (MEEI), or due to the specialization required for pediatric infusion (Pedi-GI). Urgent care is excluded as most of the appointments are on-demand hydration and do not fall into the bounds of non-oncology infusion

processes, nor do any reserve capacity for emergent infusion. Two clinics also have demand for infusion treatment, yet they do not have any capacity to serve patients. The disparities in capacity, staffing, and the lack of procedures are key to understanding the operational challenges observed by the clinics, a topic discussed in Chapter 4.

3.5 Administrative Staff and Processes

In addition to nurses, the other key human resource in an infusion clinic are the administrative staff. Staff roles include (i) patient scheduling and check in and (ii) prior authorization processing.

3.5.1 Patient Scheduling

Scheduling of infusion appointments is performed by either an administrative staff member responsible for management of the clinic's schedule, or, in many cases, by the infusion nurses. The scheduler's tasks include making contact with patients, offering appointment options, and suggesting alternatives if the patient is unable to attend or must cancel their appointment. Nurses that also schedule appointments often find their patient care activities disrupted by phone calls and emails. The scheduling systems, while electronic, require manual selection of appointment options and manual data entry of patient information.

3.5.2 Prior Authorizations

As discussed in chapter 2, prior authorizations are a significant driver of the administrative burden in an infusion clinic. In order to guarantee payment by the insurer, an approved prior authorization must be in place before the patient is infused. The infusion clinics often identify a focal responsible for obtaining prior authorization approval. The focal must be knowledgeable of infusion drugs, their uses, and in the nuanced, dynamic processes required by each insurer. Additionally, when a prior authorization is challenged by the insurer, the focal must provide supporting documentation and facilitate negotiation between physicians and the payers. If a patient is infused and prior authorization is not obtained or later rejected by the insurer, the hospital may be forced to write off a portion or all of the appointment and drug expenses[6]. None of the MGH infusion clinics have a full time dedicated prior authorization focal, although several have part time assistance. The responsibility for approvals is often assigned to infusion nurses, medical assistants, or receptionists.

3.6 Infusion Patient Flow

In order for a patient to be scheduled, approved by insurance, and treated, coordination between all of the key resources in the clinic is required. A basic overview of this process, along with a reference time line for each phase, is shown in figure 3-6. The figure depicts both first infusion flow and patients requiring multiple appointments. Multiple appointment patients will sometimes require a PA for every appointment, while some payers authorize for a specified length of time.

Administration

The physician's ordering of an infusion, prior authorization initiation, and prior authorization processing comprise the administration phase of patient flow. The physician will make the determination that a patient needs infusion, and request (typically electronically) that a prior authorization is obtained. This is a relatively quick task, and the prior authorization request is often processed within a few hours of the patient's physician appointment. The prior authorization focal, once notified by the physician, contacts the insurance company and completes any necessary paperwork. The manual nature of these processes, the lack of specialized administrative resources, and the fact that many authorizations require medical information, leads to this task often requiring several weeks for completion. After the authorization has been obtained, the patient can be scheduled into the clinic's appointment booking system. The scheduler contacts the patient (most often by phone) and selects a date and time for an appointment. Although the process of scheduling a patient may only take a few moments, the actual appointment is often not able to be scheduled until a time

[6]It is also of note that many of the drug manufacturers provide financial assistance programs to help patients pay for treatment. While MGH clinics do not promote these programs, several external infusion clinics interviewed as part of this study promote them.

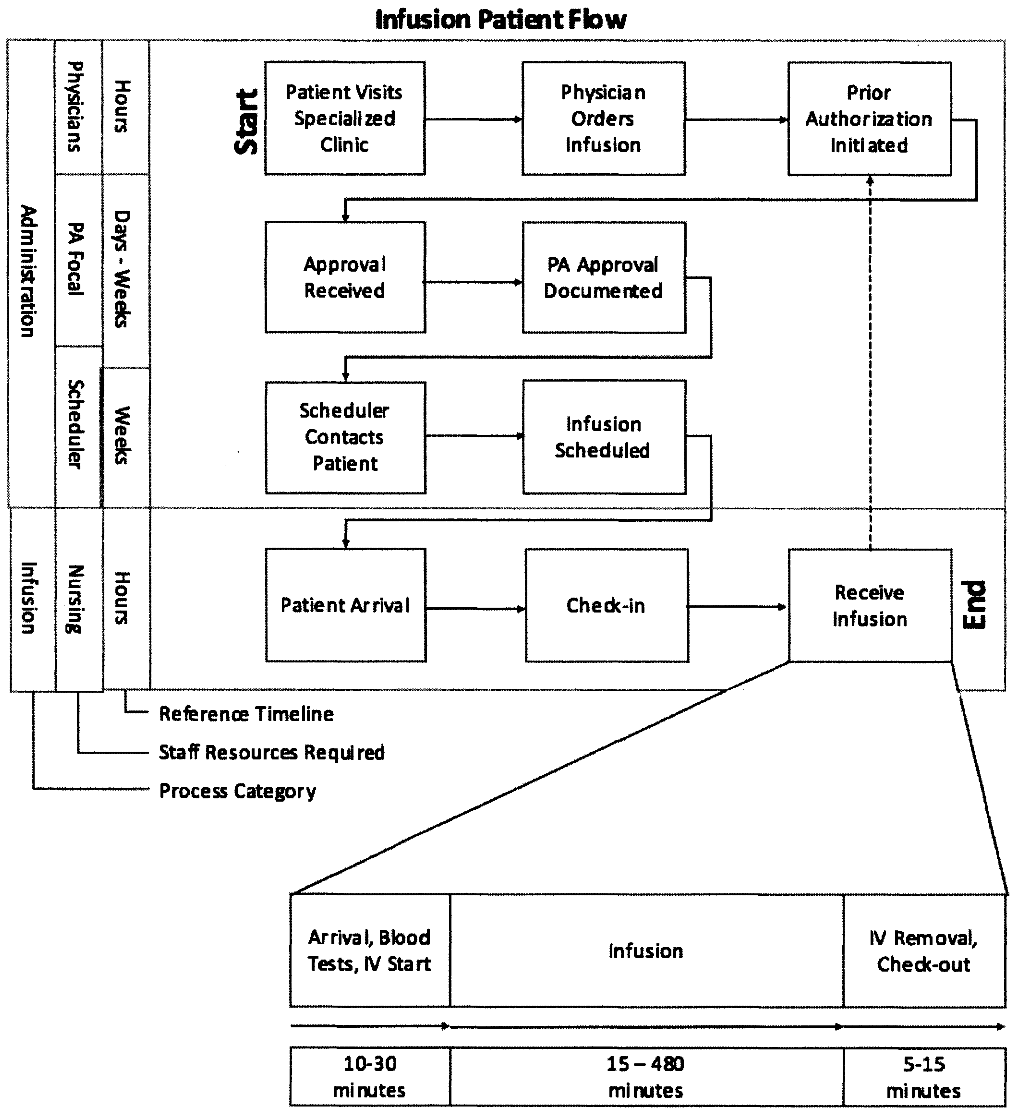

Figure 3-6: Basic overview of infusion patient flow, from physician appointment to treatment. Required resources and reference time line are indicated.

several days to weeks in the future.

Infusion

Upon arrival, the patient is checked-in and escorted to an available infusion chair. Depending upon the nature of the appointment, there are various setup procedures that need to be completed, including vital sign measurements and blood tests. The IV line must also be inserted, but all of these setup processes can usually be completed in less than thirty minutes.

The setup time is a parallel process to the mixing and delivery of the patient's drugs in the hospital's pharmacy. Drugs are typically requested the night before, and delays due to drug delivery issues are uncommon. After receiving an infusion, a patient may need to be rescheduled, and often an additional prior authorization is required. Authorizations may be valid for one infusion or for a time period, depending on the drug and payer. If a new authorizations is required, the request will be resubmitted, and the authorization reprocessed, as described in in Figure 3-6. If not, an additional appointment can be scheduled directly.

Chapter 4

Operational Challenges: Analysis and Findings

In order to understand the operational details of each of the infusion clinics, a detailed study of the current state of infusion procedures and processes was performed. The study included interviews, shadowing, appointment data collection, and data analysis. The results show that over time, as new specialty drugs have become increasingly prevalent and the demand for infusion treatment has grown (See Chapter 3), each of the clinics has independently created their own capacity to deliver outpatient infusion. This organic, tactical growth occurred without any over-arching centralized strategy or vision for the overall system. This has led to three categories of critical operational challenges that are discussed below. These include: long wait times for new infusion appointments, inefficient use of resources, and the occurrence of avoidable inpatient infusion.

4.1 Patient Access Challenges

Observation and analysis of the clinics revealed several significant operational issues, the first of which is patient access to treatment, in particular to the first infusion appointment. Interview data suggested that if a physician ordered an infusion for a patient, weeks would elapse before all of the administrative and scheduling processes were complete. As none of the clinics have urgent infusion capacity, additional delay is then caused by the congestion in the system forcing the first available appointment for the new patient to be at some point in the future. While this delay may be un-noticed by patients with long lead, elective procedures, those with urgent needs for treatment can be severely impacted.

In order to quantify this wait time, an analysis employing clinical appointment data was constructed. Using the process flow described in figure 3-6 as a guideline, a data filtering and search process was created. The analysis makes use of three key dates that are recorded for each appointment: (i) the date the first infusion appointment occurred, (ii) the date the first infusion appointment was entered into the scheduling system and (iii) the date of the patients latest non-infusion appointment in the clinic, prior to the first infusion. It is assumed that on the last non-infusion appointment the physician ordered the infusion to occur. Thus, the three dates can be used to construct a time line, as depicted in figure 4-1. The two components of this time line are PA/admin time and scheduling time. As previously discussed in chapter 3, administrative time encompasses the completion of prior authorization from the insurer. Scheduling time includes the delay that occurs starting when a patient could first be infused to when they are actually fit into the schedule.

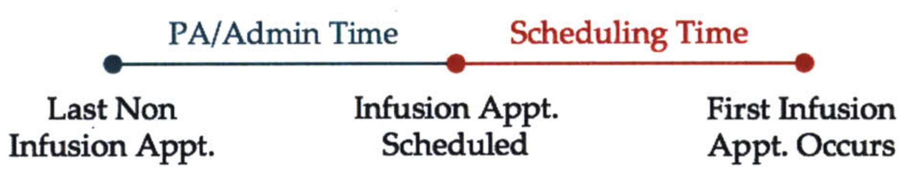

Figure 4-1: Appointment Lead Time Contributors

The results of this analysis, for new infusion patients, is shown in 4-2. The average administrative lead time, for all clinics (excluding Rheumatology) [1] was 44.2 days, the median was 28 days, and the standard deviation was 41.23 days. For scheduling time the average was 18.3 days, the median 10.7 days, and the standard deviation 21.5 days. The overall lead time (PA/admin time plus scheduling time), for all clinics, is 51.3 days, on average. Also of note is that when modeled in this manner, the average lead time is increasing. In 2012 the average was 45.7 days, increasing to 48.4 in 2013, and the aforementioned 51.3 days in 2014.

[1] Rheumatology is excluded from these calculations as they are the lone clinic that schedules patients prior to PA approval.

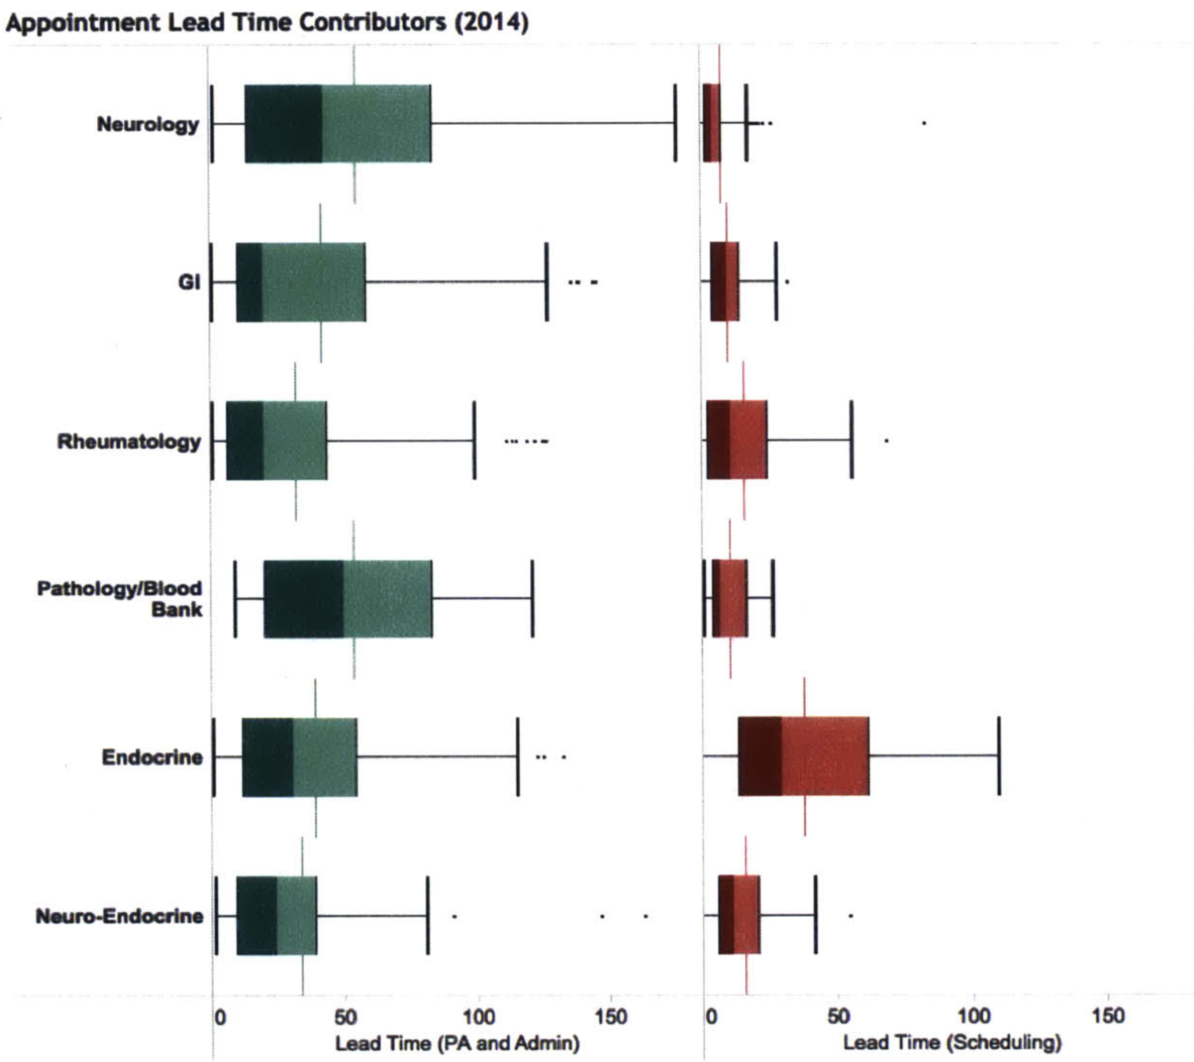

Figure 4-2: First Infusion Appointment Lead Time

These numbers may seem overly long, and it is difficult to picture that MGH would permit such lengthy delays for what are often life changing treatments. In order to validate these results, additional discussions with schedulers and prior authorization focals were conducted. Overall, very little data on administrative process turnaround time is tracked, and prior authorizations are no exceptions. Our additional subject matter expert interviews aligned with the analysis, as they indicated that authorizations can require between three and 30 days. Scheduling delay was more straightforward to validate; schedulers were asked to provide the date of the first available infusion appointment, and the answers given were in the 30-45 day range. Thus, although this calculation makes use of several underlying assumptions[2] it does prove to be realistic, and indicates that there are significant challenges associated with overcoming documentation hurdles and fitting patients into the infusion clinic's schedules.

The most important takeaway from this calculation, however, is that any intervention targeted at the alleviation of these access issues must address both drivers of new patient delay. The addition of capacity may solve scheduling constraints, but patients would still be forced to wait unnecessarily long due to administrative issues. Likewise, addressing only the administrative challenges would not lead to expedient appointment availability. The detailed means of addressing the scheduling challenges are addressed in Chapter 5. Recommendations on administrative changes are discussed in Chapter 6.

4.2 Inefficient Resource Utilization

The second operational challenge observed in the infusion clinics involves the use of physical resources: infusion chairs. As described in chapter 3, the clinics have differing numbers of chairs, and are open varying numbers of hours per week. The utilization of those chairs can be calculated using equation 4.1. In this calculation, I is the total number of infusion hours scheduled[3], and H is the total available hours of infusion in the same time period.

$$U_{Chairs} = I/H \tag{4.1}$$

Using the scheduled appointments from July 2014 to July 2015, and the available chair hours per week described in Table 3.1 average chair utilization can be calculated. Table 4.1 contains a summary of each clinic's utilization values. Average infusion chair utilization is only 32% across all the clinics, with the individual clinic's values ranging from 18% to 49%.

These values are especially concerning when coupled with the new infusion patient access challenges. The data show that infusion chair capacity is available, both intra-clinic and

[2]The major drawback of this approach to calculating lead time is that it does not discriminate between appointments that are purposely scheduled far in the future, and urgent appointments. Interviews with clinicians indicate that most, but not all first infusion appointments are not scheduled far in advance.

[3]Scheduled hours are all of the appointments in the EPIC data, for all clinics, coded as "COMPLETED" or "NO SHOW." Thus, all appointments that occupied chair space are included, as "NO SHOW" appointments are unable to be filled.

Table 4.1: Infusion Clinics Resources and Utilization

Clinic	**Nurses**	**Chairs**	**Hours**	**Average Utilization**
Neurology	3	8	400	30%
Gastroenterology	2	5	225	41%
Rheumatology	1	3	150	49%
Blood Transfusion Service	5	6	255	38%
Transplant	1	2	80	23%
Dermatology	1	1	8	–
Endocrine	1	2	48	18%
Neuro-Endocrine	1	2	48	43%
Allergy	–	–	–	–
Infectious Disease	–	–	–	–
Total	**15**	**29**	**1214**	37%

inter-clinic, yet patients are often forced to wait extended periods of time for an appointment. The results indicate that the access challenges are not driven so much by a lack of physical capacity, but instead by the staffing level of nurses within the clinic and suboptimal scheduling processes. As such, these data do not necessarily describe how "busy" the clinic may feel to the nurses staffing it at any given time.

It therefore becomes necessary to estimate the utilization of nurses, referred to here as nurse workload, $W_n(t)$. This value can be calculated for any time period using equation 4.2, where $P(t)$ is the patient census per time period, $n(t)$ is the nurse staffing level, and n_{max} is the maximum allowable patient to nurse ratio. n_{max} is a value obtained through discussion with nurses and managers at multiple infusion clinics, and validated during discussions with other hospitals[4].

$$W(t) = (P(t)/n(t))/n_{max} \tag{4.2}$$

The nurse workload results, using data from July 2014 to July 2015, are shown for two clinics, Neurology and Rheumatology, in figures 4-3 and 4-4. Examination of figure 4-3 reveals that nurse workload has high variation day to day, and sometimes exceeds the maximum advisable levels. For example, in the hours of 10:00-11:00am and 2:00-3:00pm, the max extent of the inter-quartile range exceeds 1.0, indicating that there were more than

[4]Based upon those discussions, n_{max} is set to a value of 2.5.

2.5 patients for every nurse in the clinic on these days. The average workload level, however, fluctuates throughout the day and rarely exceeds 0.4.[5] A dip in utilization is also shown during the lunch hour; this is indicative of staff availability dictating scheduling practices[6]

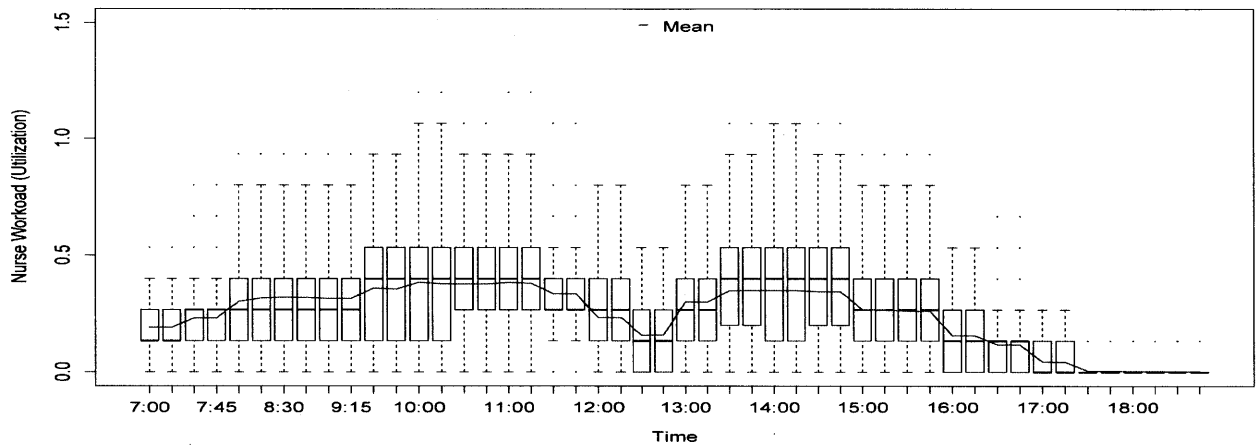

Figure 4-3: Nurse Load (Utilization), Neurology, July 2014 - July 2015

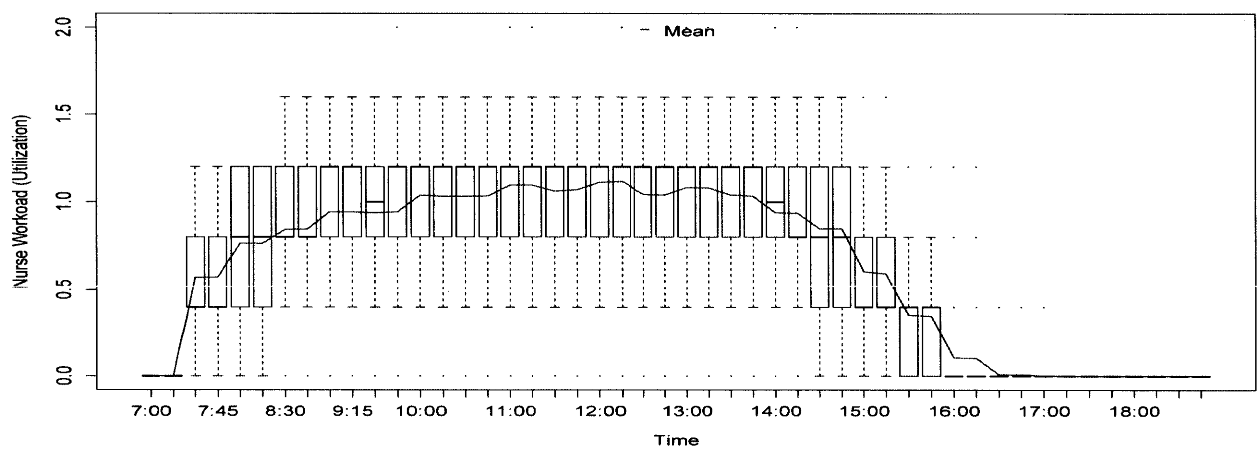

Figure 4-4: Nurse Load (Utilization), Rheumatology, July 2014 - July 2015

The Neurology workload values can be contrasted with those for Rheumatology, presented in figure 4-4. Staffed by only one nurse, the workload levels in the Rheumatology clinic are much more extreme than those observed in Neurology[7], and the mean workload value is

[5]Note that the average utilization is essentially zero after 5:00pm. This is due to the fact that in the period of study the clinic typically closes by 5pm.

[6]Shadowing and interviews revealed that in addition to treating patients, nurses were sometimes responsible for other tasks including scheduling, prior authorizations, answering the phone, and mixing medication. Thus utilization is likely higher than these calculations indicate.

[7]Like Neurology, low utilization values in Rheumatology after 5pm are because the clinic is typically closed after 5pm.

typically over 1.0. These results show the impact of staffing levels. This utilization value is highly variable, and it is not atypical for the single nurse on staff to be overloaded[8]. It is also of note that there is not the same lunchtime "dip" observed in Neurology. The clinic does not shut down during the lunch break, and it is common for the nurse to eat at a desk while treating patients.

Analysis of these two current state metrics indicate that although having enough physical capacity is important, staffing decisions and scheduling processes are critical to patient access and resource utilization.

4.3 Inpatient Infusion

The final operational challenge noted in the current state analysis is the occurrence of avoidable inpatient infusion. Several clinics indicated, during shadowing and interviews, that admitting a patient to the hospital for infusion was sometimes necessary. These admissions, in their view, were not due to medical necessity, but were in fact an additional symptom of the lack of outpatient access and the resource challenges discussed previously. As very few referral processes exist, some clinics lack physical capacity, and all clinics lack urgent capacity, inpatient infusion is sometimes seen as the only option to ensure expedient care.

Thus, these "avoidable" infusion admissions occur when patients urgently need to receive treatment, and although they could receive it an outpatient setting, they are admitted because no outpatient resources are available. Quantification of these admissions is exceedingly challenging. Admissions to the hospital are not categorized as avoidable in any electronic data system, and "avoidable" vs "necessary" involves a significant amount of clinical judgement.

In general, data on avoidable infusion admission was gathered using the following process. First, multiple searches of inpatient records using diagnosis, treatment, and pharmacy codes that potentially suggest infusion administration were conducted. The list of codes was

[8]During shadowing and interviews with Rheumatology, it was noted that due to the overloaded schedule, patients were sometimes referred to a non-MGH (the physician managed Nephrology clinic) for treatment. Detailed records are not kept, but a short study showed 20 patients were referred over a six month period. This represents a significant loss of revenue to the hospital

developed after discussions with subject matter experts from each of the infusion clinics. The list of patients was refined by limiting the length of stay (LOS) to three days or less. Next, clinical experts were asked to review the remaining data and indicate whether the admissions were avoidable or not. From these data a typical avoidable infusion admission length of stay was calculated. To further gather data on the volume of avoidable infusion admission, a special e-mail account was established for physicians and nurses to send notification of a potential avoidable admission cases. The e-mail data were used to further bolster and refine the estimates on avoidable infusion admission patient volume and typical length of stay. The details and methodology of the search process are described at length in Appendix B.

Table 4.2 summarizes the categories of admission, ranges of estimates, those appointments that were validated in hospital data, and the total estimated bed days per year taken up by avoidable infusion admissions. Review of this data reveals that significant numbers of bed days are being consumed by treatments that can, and should, be administered in the outpatient infusion clinics. Between 497 and 1007 bed days per year are consumed annually by avoidable admission, a value equivalent to between 1.3 and 2.8 hospital beds per day[9], mostly on hospital floors that are already the most highly utilized. Inpatient beds are a highly utilized resource, and even small reductions in utilization can lead to significant improvements in patient access and cost.

[9]Chemo Desensitization is reported here only to highlight that it is a type of avoidable admission related to infusion. Based upon discussions with clinical experts, the MGH Cancer Center is best suited to treat these patients.

Table 4.2: Avoidable Infusion Admission Summary Data

Admission Type (Clinic)	**Admits (Upper)**	**Admits (Lower)**	**Admits (Validated)**	**Length of Stay**	**Low Bed Days/yr**	**High Bed Days/yr**
Antibiotics, OPAT (Inf. Disease)	50.0	50.0	26.0	1.6	79.5	79.5
Antibiotics, Cellulitis (Inf. Disease)	161.0	484.0	646.0	1.6	254.4	764.7
Antibiotics, CF (Cystic Fibrosis)	16.0	16.0	16.0	4.0	64.0	64.0
Emergent Infusion (Rheumatology)	10.0	10.0	7.0	2.4	24.3	24.3
Emergent Transf. (BTS)	50.0	50.0	50.0	1.5	75.0	75.0
Chemo Desens (Allergy)	181.0	271.0	362.0	1.5	271.5	406.5
Total, excluding Chemo Desens	**287.0**	**610.0**	**745.0**	**11.1**	**497.2**	**1007.5**
Total	**468.0**	**881.0**	**1107.0**	**12.6**	**768.7**	**1414.0**

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Future State Modeling

5.1 Modeling Overview

The modelling effort described in this chapter was performed in alignment with the overall goal of this project, i.e., improving patient access to infusion treatment. In particular, models were constructed to evaluate the resource requirements and performance of a potential future centralized infusion unit at MGH. By applying a scheduling process to historical data, the resulting model can be analyzed to provide robust recommendations describing the configuration of a potential centralized infusion clinic. This chapter contains an overview of the modeling approach, the assumptions and technical details, as well as a detailed review of the corresponding results.

The operational challenges observed in the clinics (long wait for a new appointment, inefficient resource utilization, and avoidable inpatient infusion) are symptomatic of root causes that could be addressed through the pooling of resources and the standardization of processes. As described in Chapter 4, however, in order to ensure the entire system performs adequately, centralization and standardization must be coupled with effective scheduling techniques.

An outline of the modeling approach is presented in Figure 5-1. First, options for centralization were developed. These options, and the core assumptions key to all models, are detailed in Section 5.3. To ensure that the modeling effort is realistic and includes appropriate clinical and administrative practices, assumptions on the configuration of the proposed clinic were vetted with subject matter experts. These assumptions, the constraints placed on the scheduling process, and the historical appointment data used in the models, are described in Sections 5.3. The modeling effort results in a robust framework that provides a means to simulate the anticipated performance of different scenarios.

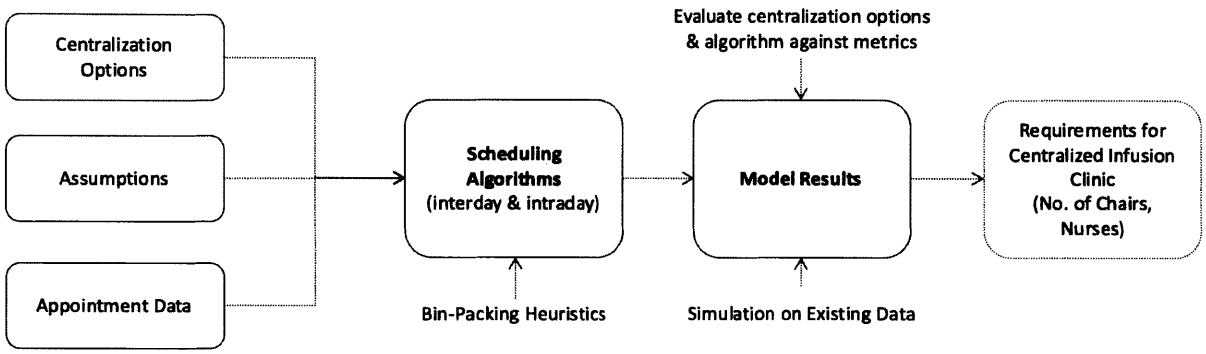

Figure 5-1: Model Development Process Overview

A customized, specialized on-line scheduling algorithm, which seeks to minimize resource requirements, is employed to schedule appointments in this modeling effort. The algorithm makes use of the same basic scheduling heuristic as the MGH-MIT Cancer Center project[7], but with some modification to improve performance and add details specific to the non-oncology setting. The algorithm is used to "reschedule" historical appointments, under specified constraints. The rescheduled data is considered reasonable representation (model), of a proposed, future clinic. When the new algorithm processes appointments, they are scheduled by several methods. The first method is *intraday*, where the appointments are constrained to occur within the same day they occurred in the historical data set. The second scheduling option is *interday*, a process by which appointments are permitted to float over a prescribed multi-day window. In order to test and potentially improve the performance of the scheduling process, variations of three bin-packing heuristics are also applied. A detailed overview of the technical specifics of the scheduling algorithm, including intraday and interday scheduling, along with a discussion on the selected bin-packing heuristics, is provided in Section 5.4.

As there are many permutations of centralization options and algorithm settings, each selected set of options is referred to as a "scenario". A review of the nomenclature and a summary of the selected scenarios is included in Section 5.3.1. After the data were processed for a given model, performance metrics were computed. Those metrics are described in Section 5.3. Results for all models are compared and contrasted in Section 5.5

5.2 Input Data Sources and Summary

Table 5.1 contains a summary of the appointment data used for all models. In total 13,211 appointments are scheduled in each model. These appointments comprise nearly 23,000 hours of infusion and represent 3,943 unique patients. All appointment data was extracted from the EPIC scheduling system. Appendix A contains a complete list of the included EPIC appointment type codes, which are descriptive of the drugs and treatments considered amenable for centralization and are thus included in each scenario.

Table 5.1: Model Data Summary (July 2014-July 2015)

Clinic	Appointments	Hours	Patients
Neuro-Endocrine	1033	895	568
Blood Transfusion Service	1693	6062	593
Neurology	1966	3712	438
GI	2731	5437	491
Endocrine	1423	706	1062
Rheumatology	2166	4698	539
Allergy	1415	354	150
Transplant	693	842	157
Dermatology	91	89	34
Total	**13211**	**22796**	**3943**

5.3 Core Assumptions and Constraints

The models are dependent upon several core assumptions:

- A single staffing unit will support each clinical scenario, whether physically centralized or multi-location. The unit will be composed of any necessary nursing staff and administrative staff.

- Centralized scheduling will be used to place appointments into the clinic(s) that comprise each scenario.

- Proposed clinics are assumed to be open 12 hours per day, from 7:00am (T_s) to 7:00pm (T_e), Monday through Friday[1]

- The demand from all MGH infusion clinics with existing physical capacity, as described in Chapter 3, plus the Allergy clinic, are included.

- All drugs and treatments that are amenable[2] to centralization are included. For a complete list of included appointment codes, see Appendix A

- Appointment data for amenable treatments, in the time period July 15, 2014 to July 15, 2015, are included in the model. "No Show" and "Completed" appointment types, as specified in EPIC data, are included. "Cancelled" appointments are assumed to be rescheduled within the period of interest and are not included.

Centralization Options

Two distinct categories of options for the centralization of resources are modeled. The first, *physical*, implies that all infusion nurses and chairs are combined into a single location on the MGH campus. The second option, referred to as *multi-location*, involves the use of two or more locations[3] linked by centralized nurse staffing and administrative personnel. Based upon direction from MGH leadership, only one multi-location scenario was considered. It makes use of one existing clinic on the MGH campus (Neurology), and the the newly established location at MGH-West, an 8 chair clinic briefly discussed in chapter 3.

[1] Opening hours are based upon the MGH Cancer Center and interviews with centralized infusion clinics.

[2] Amendable to centralization is determined by clinical judgement and based upon interviews with centralized clinics. Some diagnostic tests and specialized blood treatments are excluded.

[3] A temporary multi-location solution was requested to be evaluated by MGH leadership due to the long lead time and cost required to secure space for a physical clinic. More on the need for an interim location is discussed in Chapter 6.

In the multi-location scenario, patients are assumed to prefer one clinic (MGH or MGH-West), and they do not fluctuate between clinics, i.e. patients receive recurring treatment in the same location. Patient's home zip codes, and their distance to MGH-West, are used to provide a proxy estimate of preference. Based upon guidance from MGH, appointment data for those patients that received infusion on the main MGH Campus is filtered to find those patients that live within a 20 mile radius of MGH-West. The 20 mile radius was chosen after analyzing the zip codes of the patients that receive treatment at MGH-West and comparing how far they live from the main campus versus MGH-West. Next, within this radius, only those that live closer to MGH-West than the main MGH campus are retained. Finally, among those remaining patients, 75% are randomly selected and assumed to prefer MGH-West over MGH.

Scheduling Constraints

In addition to the core assumptions previously discussed, several constraints are placed upon the algorithmic scheduling process:

- The hours of operation, T_s to T_e are discretized into 15 minute increments[4].
- As described in equation 4.2, the maximum patient to nurse ratio is 2.5:1.
- Number of scheduled patient arrivals per time period t, $a(t)$, is limited to be one fewer than the number of nurses on staff in the clinic[5], that is:

$$a(t) \leq n(t) - 1 \tag{5.1}$$

- Appointments are scheduled in the same order in which they "arrived," that is, in the order of the "APPT SCHEDULED DATE TIME" field from the EPIC data[6]

[4]This level of discretization was chosen as the shortest appointment is 15 minutes.

[5]While nurses can serve more than one patient at a time, they can not start IVs on more than patient at a time. Thus the period of arrival is assumed to consume a nurse's time. One nurse is included as system slack, and is excluded from this calculation.

[6]APPT SCHEDULED DATE TIME is a data field representing the time and date an appointment was entered into the scheduling system. This is in contrast to the APPT DATE field, which records the date an appointment actually occurred.

In addition to the core assumptions, location choices, and scheduling constraints, each scenario includes additional levers that can be adjusted to influence the results. These are (i) the applied scheduling algorithm, (ii) the amount of interday flexibility, and (iii) the application of bin packing heuristics. All three items are discussed in detail in the Scheduling Algorithm Overview contained in Section 5.4.

Scenario Metrics

After the scheduling algorithm is applied in each scenario, the results are evaluated against a set of metrics. These include:

- Chair Requirement: The number of chairs required, C_{max} is the maximum of the patient census, P, over all time periods.

$$C_{max} = \max_{t \in T}\{P(t)\}, \tag{5.2}$$

  where T is the entire time period, July 2014 to July 2015.

- Nurse Requirement: The number of nurses required on staff, $n(t)$, is equal to the number of chairs required, C_{max}, divided by the patient to nurse ratio, n_{max}.

$$n(t) = C_max(t)/n_max \tag{5.3}$$

Useful statistics on these metrics (mean, median, percentiles, standard deviation, coefficient of variation) are calculated.

5.3.1 Modeling Scenarios

For labelling proposes, modelling scenarios are indicated with an alpha-numeric string. Each label first consists of P or an a M, to distinguish between physical and multi-location options. Next, the selected bin packing option is indicated, with an F (first available), L (last available), or an R (random). The next character is a number, indicating the amount of flexibility applied to the algorithim. Additionally, multi-location options are suffixed with an A or a B, to distinguish between the MGH Campus location (A) and MGH-West (B). For example, scenario P-F-0 represents a physical centralization scenario with first available packing and zero days of flexibility, while M-L-1-B indicates a multi-location scenario at MGH-West with last available packing, and one day of flexibility. Table 5.2 describes each of the physical centralization scenarios, and which settings for interday flexibility and bin packing they incorporate. Table 5.3 contains the same information for the multi-location scenarios. In total, the modeling includes 18 unique scenarios, exclusive of the reference model and sensitivity studies.

Table 5.2: Physical Centralization Modeling Scenarios

Scenario Label	**Flexibility**	**Bin-Packing**
P-F-0	0	First Available
P-L-0	0	Last Available
P-R-0	0	Random
P-F-1	1	First Available
P-L-1	1	Last Available
P-R-1	1	Random
P-F-2	2	First Available
P-L-2	2	Last Available
P-R-2	2	Random
P-F-3	3	First Available
P-L-3	3	Last Available
P-R-3	3	Random

Table 5.3: Multi-location Centralization Modeling Scenarios

Scenario Label	**Flexibility**	**Bin-Packing**
V-F-0-A	0	First Available
V-F-0-B	0	First Available
V-L-0-A	0	Last Available
V-L-0-B	0	Last Available
V-R-0-A	0	Random
V-R-0-B	0	Random
V-F-1-A	1	First Available
V-F-1-B	1	First Available
V-L-1-A	1	Last Available
V-L-1-B	1	Last Available
V-R-1-A	1	Random
V-R-1-B	1	Random

Reference Model

A high level view of the current state, referred to as the reference model, is also calculated, as it is useful to compare against any future state options that are developed. The reference model includes all of the infusion clinics, and assumes that all appointments are fixed in time, that is, they occur at the same date and time on which they were historically scheduled in the hospital data-set. The reference model provides a picture of how a centralized clinic would perform, if it were to make use of the current scheduling processes and procedures. The actual demand that the reference model portrays is, in reality, distributed amongst the clinics. This model, however, does give an indication of the resource requirements for a naively created centralized system, and it serves as an adequate baseline for comparison.

The evaluative metrics of the reference model are depicted in an aggregated box plot in figure 5-2. The number of chairs, C_{max}, equals 23, and the nurses required[7], n, is 10.

[7]Per equation 5.3, the number of nurses needed is 10. Note that in the actual current state system, the number of nurses is 15.

The coefficient of variation of day-to-day patient census is 0.74, indicating a high level of fluctuation in the number of patients present throughout the day, over the one year period of study.

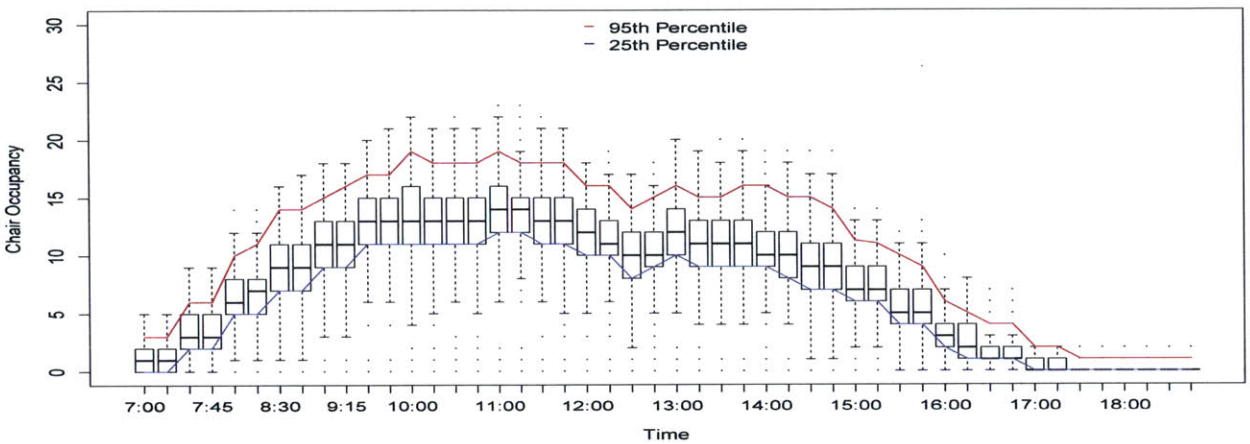

Figure 5-2: Reference Model: Infusion Chair Census, July 2014-July 2015

5.4 Scheduling Algorithm Overview

5.4.1 Max(Max-Min) Heuristic

For each of the above scenarios, one could, through the application of a scheduling process, attempt to optimize the schedule, reduce resource requirements, and smooth utilization. The models created for this project achieve that goal by using the *max(max-min) scheduling algorithm*, developed from insights gained after optimally scheduling a subset of MGH Cancer Center appointments using linear mixed-integer optimization. The algorithm seeks to reschedule new appointments while balancing infusion chair requirements. The algorithmic scheduling process is graphically described using a representative example in figures 5-3, 5-4, and 5-5. In the example, a sample 30 minute appointment is scheduled into a time period containing prior scheduled appointments. The process is described below.

1. The overall max patient census, U, is calculated for the existing state, and input into the algorithm. In Figure 5-3 the value of U is 6.

2. Next, feasible appointment times are determined. In Figure 5-3, the feasible times for the sample 30 minute appointment are any time that allows the appointment to be completed by the end of the prescribed time period[8]. There are, thus, 12 feasible appointment start times in the example.

3. For each feasible start time, the additional patient is added to the existing census across the corresponding appointment's interval. Then the minimum value of the census across the appointment's duration, if it were to be scheduled in a given time period, u, is calculated. Values for u are shown in Figure 5-3, and range from four to seven.

4. For each feasible start time, the value s, denotes the difference between U, from the existing state, and u from each potential appointment location is calculated. In Figure 5-3 values of s are shown for the feasible appointments, and range from negative one to two.

5. If any feasible appointment start times increase the overall value of the census[9], and there exist feasible options that do not increase the census, any options causing an increase are removed from consideration. In Figure 5-4, several options including 7:00am, 7:15am, 8:00am, and others are removed as they increased U from a value of 6 to a value of seven, and there were other options available where U was held constant.

6. The max(max-min) algorithm seeks to maximize[10] s. In Figure 5-5, options A through D have the maximum value of s, which is two. Option E, with $s = 1$, is thus removed from consideration.

7. Remaining appointment start time option(s) are then considered. Figure 5-5 depicts the remaining options after application of the heuristic. These four options (A,B,C,D) are considered "ties" as they have equal impact on the metrics of import to the

[8] 10:00 am in this example. Note that in a real scenario the end of the available time period is equivalent to the closing time for the clinic, 7:00 PM

[9] Potential slots are also excluded if they exceed the maximum allowable appointment starts per time period

[10] Although the process of scheduling is explained in this manner, in practice, maximizing s is same as minimizing u, because U is fixed.

max(max-min) heuristic. The scheduling process can choose between the four options (randomly or other methods).

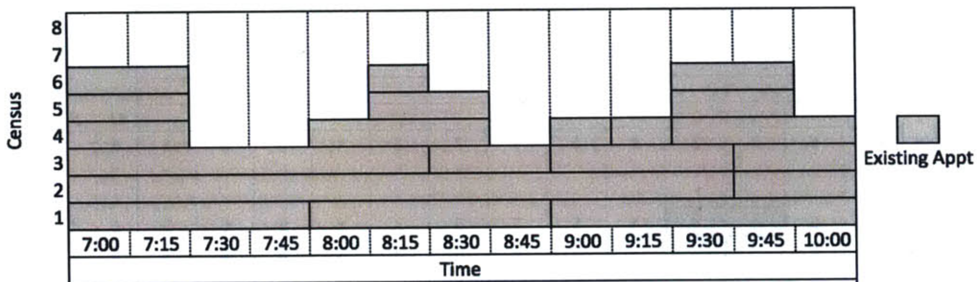

Census	7:00	7:15	7:30	7:45	8:00	8:15	8:30	8:45	9:00	9:15	9:30	9:45	10:00	*U*	*u*	*s*
Current	6	6	3	3	4	6	5	3	3	4	6	6	5	6	–	–
Start = 7:00	7	7	3	3	4	6	5	3	3	4	6	6	5	7	7	-1
Start = 7:15	6	7	4	3	4	6	5	3	3	4	6	6	5	7	4	2
Start = 7:30	6	6	4	4	4	6	5	3	3	4	6	6	5	6	4	2
Start = 7:45	6	6	3	4	5	6	5	3	3	4	6	6	5	6	4	2
Start = 8:00	6	6	3	3	5	7	5	3	3	4	6	6	5	7	5	1
Start = 8:15	6	6	3	3	4	7	6	3	3	4	6	6	5	7	6	0
Start = 8:30	6	6	3	3	4	6	6	4	3	4	6	6	5	6	4	2
Start = 8:45	6	6	3	3	4	6	5	4	4	4	6	6	5	6	4	2
Start = 9:00	6	6	3	3	4	6	5	3	5	5	6	6	5	6	5	1
Start = 9:15	6	6	3	3	4	6	5	3	3	5	7	6	5	7	5	1
Start = 9:30	6	6	3	3	4	6	5	3	3	4	7	7	5	7	7	-1
Start = 9:45	6	6	3	3	4	6	5	3	3	4	6	7	6	7	7	-1

Figure 5-3: Sample Application of Scheduling Algorithm, Part I

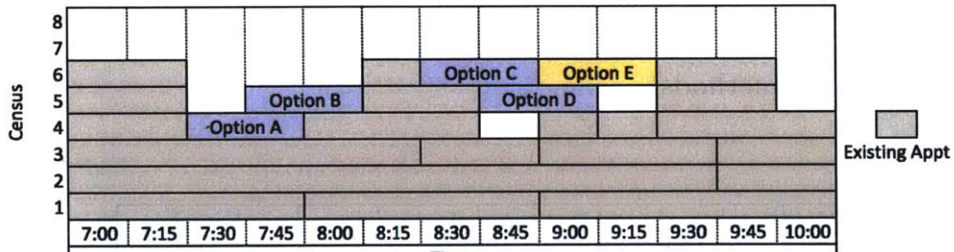

	7:00	7:15	7:30	7:45	8:00	8:15	8:30	8:45	9:00	9:15	9:30	9:45	10:00	*U*	*u*	*s*
Current	**6**	**6**	**3**	**3**	**4**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**–**	**–**
Start = 7:00	7	7	3	3	4	6	5	3	3	4	6	6	5	7	7	-1
Start = 7:15	6	7	4	3	4	6	5	3	3	4	6	6	5	7	4	2
Start = 7:30	**6**	**6**	**4**	**4**	**4**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 7:45	**6**	**6**	**3**	**4**	**5**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 8:00	6	6	3	3	5	7	5	3	3	4	6	6	5	7	5	1
Start = 8:15	5	6	3	3	4	7	6	3	3	4	6	6	5	7	6	0
Start = 8:30	**6**	**6**	**3**	**3**	**4**	**6**	**6**	**4**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 8:45	**6**	**6**	**3**	**3**	**4**	**6**	**5**	**4**	**5**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 9:00	**6**	**6**	**3**	**3**	**4**	**6**	**5**	**3**	**5**	**5**	**6**	**6**	**5**	**6**	**5**	**1**
Start = 9:15	6	6	3	3	4	6	5	3	3	5	7	6	5	7	5	1
Start = 9:30	6	6	3	3	4	6	5	3	3	4	7	7	5	7	7	-1
Start = 9:45	6	6	3	3	4	6	5	3	3	4	6	7	6	7	7	-1

Figure 5-4: Sample Application of Scheduling Algorithm, Part II

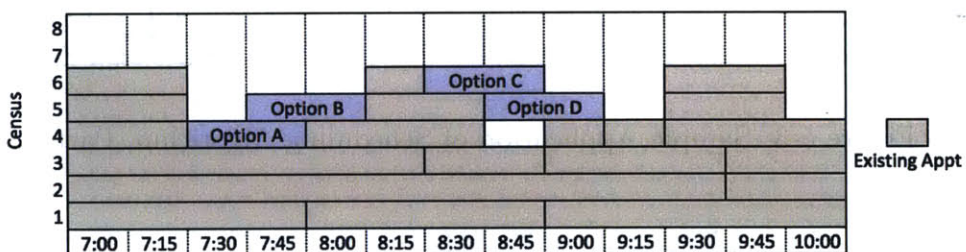

	7:00	7:15	7:30	7:45	8:00	8:15	8:30	8:45	9:00	9:15	9:30	9:45	10:00	*U*	*u*	*s*
Current	**6**	**6**	**3**	**3**	**4**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**–**	**–**
Start = 7:00	7	7	3	3	4	6	5	3	3	4	6	6	5	7	7	-1
Start = 7:15	6	7	4	3	4	6	5	3	3	4	6	6	5	7	4	2
Start = 7:30	**6**	**6**	**4**	**4**	**4**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 7:45	**6**	**6**	**3**	**4**	**5**	**6**	**5**	**3**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 8:00	6	6	3	3	5	7	5	3	3	4	6	6	5	7	5	1
Start = 8:15	6	6	3	3	4	7	6	3	3	4	6	6	5	7	6	0
Start = 8:30	**6**	**6**	**3**	**3**	**4**	**6**	**6**	**4**	**3**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 8:45	**6**	**6**	**3**	**3**	**4**	**6**	**5**	**4**	**5**	**4**	**6**	**6**	**5**	**6**	**4**	**2**
Start = 9:00	6	6	3	3	4	6	5	3	5	5	6	6	5	6	5	1
Start = 9:15	6	6	3	3	4	6	5	3	3	5	7	6	5	7	5	1
Start = 9:30	6	6	3	3	4	6	5	3	3	4	7	7	5	7	7	-1
Start = 9:45	6	6	3	3	4	6	5	3	3	4	6	7	6	7	7	-1

Figure 5-5: Sample Application of Scheduling Algorithm, Part III

5.4.2 Intraday and Interday Algorithm

The max(max-min) heuristic can be applied over various time windows. In the first option, intraday, the process is only used to reschedule appointments within the same day on which they occurred in the past. The second, approach, interday, allows appointments to float within a prescribed interday window of time. The data on which the appointment historically occurred, is referred to as the "target date".

The value of the the target date is determined directly from the appointment dataset. For interday scheduling, however, flexibility, f, must be prescribed for each appointment. The modeling approach allows flexibility to be defined on both the left, f^-, the right, f^+, of an original appointment date[11]. When both f^- and f^+ are set equal to zero, the appointment has no interday flexibility, and is scheduled the intraday process. The complete algorithm, as it was implemented programatically in code (using the R programming language), is described in Algorithm 1.

[11]When rescheduling an appointment, the algorithm only uses weekdays. If a an appointment occurred on a Friday, and one day of flexibility is added, the window of potential appointments would span from the prior Thursday to the next Monday.

Algorithm 1 Scheduling Algorithm

Require: $\mathcal{A} : (p, d, h, f^+, f^-, fixed) \leftarrow$ set of appointments ordered by date of arrival (includes appointments p, date d, original appt time h, flexibility f^+ and f^- and fixed $fixed$); $T_s, T_e \leftarrow$ start and end time of infusion center (7:00am and 7:00pm, respectively); $\delta \leftarrow$ time interval default at 15 minutes;

$\mathcal{U} \leftarrow \emptyset$ (set of scheduled appointments)
$\mathcal{S} \leftarrow \emptyset$ (set of number of appointment starts)
for $d \in$ dates **do**
  for $p \in$ appointments in day d **do**
    if p is $fixed$ **then**
      $\mathcal{U}(p) \leftarrow$ original (fixed) time h
      $\mathcal{S}(h) \leftarrow \mathcal{S}(h) + 1$
    else
      dayRange $\leftarrow [\max(1, d - f^-), \min(d + f^+, size(\mathcal{U}))]$
      bestTimeAll $\leftarrow$ default (when p is scheduled at the first available slot)
      bestDayAll $\leftarrow$ default (when p is scheduled at the first available day)
      bestPeakAll $\leftarrow$ default (based on bestDayAll and bestTimeAll)
      for $j \in$ dayRange **do**
        timeSlots $\leftarrow$ all valid time slots in the preferred order
        timeIntra $\leftarrow$ timeSlots(1)
        peakIntra $\leftarrow \max(\mathcal{U}(p))$
        for $t \in$ timeSlots **do**
          peakCur $\leftarrow \max(\mathcal{U}|$appointment p is scheduled at day j at time $t)$
          startCur $\leftarrow \mathcal{S}(t, j)$
          if $peakCur < peakIntra$ OR $peakCur$ is equal but more attractive than $peakIntra$ **then**
            peakIntra $\leftarrow$ peakCur
            timeIntra $\leftarrow t$
          end if
        end for
        if $peakIntra < bestPeakAll$ OR $peakIntra$ is equal but more attractive than $bestPeakAll$ **then**
          bestPeakAll $\leftarrow$ peakIntra
          bestTimeAll $\leftarrow$ timeIntra
          bestDayAll $\leftarrow j$
        end if
      end for
      $\mathcal{U}(p) \leftarrow$ on $bestDayAll$ at $bestTimeAll$
      $\mathcal{S}(bestTimeAll, bestDayAll) \leftarrow \mathcal{S}(bestTimeAll, bestDayAll) + 1$
    end if
  end for
end for

5.4.3 Application of Bin Packing Heuristics

The selection of appointment times in the existence of ties quickly becomes important, given that there are often many feasible start times, especially when employing interday scheduling. In this modelling effort, modified approaches to the bin packing problem, as discussed in Chapter 2, are used to select between feasible appointment start times. As mentioned, bin backing approaches are often used in scheduling problems, most often when scheduling the use of specific resources (chairs, machines, rooms). While the algorithm used in this thesis does not schedule appointments into specific chairs, bin packing heuristics provide a logical and repeatable means of scheduling appointments[12]. Their use also provides an extensible framework for enhancement of the scheduling process, should the need to schedule to specific chairs become necessary in the future. The selected approaches used in this study include:

- First available (first fit): select the appointment start time that has the earliest start time. This is a version of the classic "first fit" bin packing heuristic.
- Last available: select the appointment start time that has the latest start time. Also an application of the "first fit" algorithm, but instead of searching left to right (in time series order) the search is made right to left.
- Random (random fit): select any feasible start time, randomly.

The graphical sample example from Figure 5-5, with four options of equal impact, is amended in Figure 5-6 to differentiate between these bin packing options. In Figure 5-6, the first available appointment, 7:30am, is selected if the first available heuristic were applied. 8:45am, the latest starting feasible appointment, is chosen in a last available application. Any feasible start time, 7:30am, 7:45am, 8:30am, or 8:45am, can be chosen when using random selection.

First fit and last fit are very similar to one-another, as they only differ in the "end" of the time-line (start of the day or the end of the day) to which appointments are driven. With so many options (ties), random selection yields results different from either first fit or last

[12] It is of note that the scheduling algorithm assumes that patients agree with appointment selected by the algorithm. Potential challenges to this assumption are addressed in the Sensitivity Study section, below.

fit. Discussion on the differences in performance of each of the heuristics, as applied to this problem, is conveyed in the Results section, below.

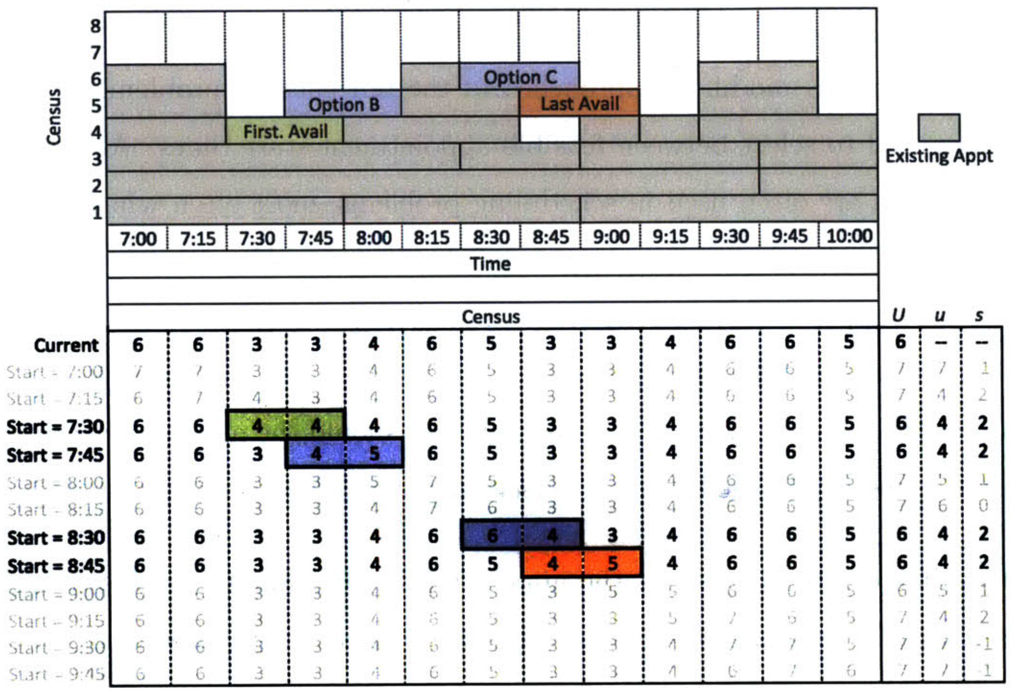

Figure 5-6: Sample Application of Scheduling Algorithm, Part IV

5.5 Results

Table 5.4 contains a summary of the statistics derived after the modeling of each scenario. The reference model is also included for comparison. Details and and data from selected physical and multi-location models are included below. Appendix C includes box plots of chair census for all modeled scenarios. The following definitions apply to the data presented in Table 5.4:

- Scenario: The scenario name, as described in Section 5.3.1.
- Max C: The maximum of the patient census, also equal to the maximum number of chairs required, across all days modeled in the scenario.
- Median C: The median number of the census across all days modeled in the scenario.

- 95% C: The 95th percentile of the census, across all days modeled in the scenario.
- Std. Dev. C: The standard deviation of the the patient census.
- Nurse Requirement: The number of nurses required, calculated by equation 5.3.
- CV Census: The coefficient of variation of the patient census (standard deviation divided by mean).

Table 5.4: Scenario Data Summary

Scenario	Max C	Median C	95% C	Std. Dev C	Nurse Reqt	CV Census
Reference	23	8	19	5.48	10	0.74
P-F-0	12	8	11	2.37	5	0.33
P-L-0	12	7	10	1.85	5	0.26
P-R-0	12	7	11	1.95	6	0.27
P-F-1	10	7	9	1.54	4	0.21
P-L-1	11	7	9	1.35	5	0.19
P-R-1	11	7	10	1.74	5	0.24
P-F-2	10	7	9	1.39	4	0.19
P-L-2	11	7	9	1.33	5	0.18
P-R-2	11	7	10	1.40	5	0.19
P-F-3	11	7	9	1.36	5	0.19
P-L-3	11	7	9	1.35	5	0.19
P-R-3	12	7	9	1.43	5	0.20
M-F-0-A	7	4	7	1.50	3	0.36
M-F-0-B	8	3	6	1.37	4	0.44
M-L-0-A	7	4	6	1.20	3	0.29
M-L-0-B	8	3	5	1.18	4	0.38
M-R-0-A	8	4	7	1.30	4	0.31
M-R-0-B	8	3	6	1.27	4	0.41
M-F-1-A	7	4	5	0.89	3	0.21
M-F-1-B	5	3	4	0.81	2	0.26
M-L-1-A	7	4	6	0.87	3	0.21
M-L-1-B	5	3	5	0.88	2	0.28
M-R-1-A	8	4	6	1.08	3	0.26
M-R-1-B	5	3	5	0.99	2	0.31

Discussion: Physical Scenarios

Application of intraday scheduling, observed in scenarios P-F-0 through P-R-0 yields a significant reduction in the required chairs (23 reduced to 12). Subsequently, nurse staffing is also reduced (from ten to five). Also of import, the day-to-day variation in patient census is drastically diminished, dropping from 0.74 in the reference model to a value of 0.26 in scenario P-L-0. This indicates that by making use of the advanced scheduling algorithm the nurses should expect a more stable number of patients per day.

In the intraday options, slight differences between the performance of the various bin packing heuristics can also be observed. First available and last available perform similarly, when evaluated on the number of chairs required (12). Last available scheduling, does, however, yield a slightly better coefficient of variation than does first available (0.26 vs 0.33). Although the coefficient of variation for intraday random available (scenario P-R-0) is competitive with last available, the number of chairs required is increased over the other two options.

The reason for this increase is "adversarial" scheduling. As patients are scheduled in order, and knowledge of the next patient to be scheduled does not exist, a decision to place an appointment can have significant impact on the possible locations of fit for the next appointment. More examples of adversarial scheduling will arise in the discussion of interday scheduling scenarios.

Figure 5-8 depicts the aggregated chair census throughout the period of study, after the algorithm has been applied. It is observable, as compared to the reference model in shown in figure 5-2, that the profile is much flatter, and lacks a significant mid-day peak. The impact of the maximum arrivals per time period, $a(t)$, can easily be observed in the ramp effect from 7:00am until the census stabilized. Also, as this figure depicts an implementation of the first available heuristic, appointments are "pushed" to the left, the earlier part of the day. This pushing results in a steep downward slope in chair occupancy towards the end of the day.

Interday results for scenarios P-F-1, P-L-1, and P-R-1, each of which include one day of flexibility on either side of the original appointment date, are depicted in Figures 5-8 through 5-10. Each of the three figures is included here, as the differences and similarities are of note. First available performs the best when measured by chair requirements (10), while last available has the best coefficient of variation (0.19). Both of these numbers are improvements over the intraday result.

Graphical analysis of the three interday figures reveals the same appointment start ramp observed in intraday. The first available application, in scenario P-F-1, also shows the same end of day decline. The last available box plot, developed from scenario P-L-1, does not show the same end of day drop off. In this scenario the clinic remains busy until closing. The end of day business is also due to the heuristic "pushing" appointments, although this time to the right. Random selection is shown to require more chairs and have a higher amount of day to day variation. This is likely due to an increased amount of adversarial scheduling. Inspection of 5-10 reveals small intraday peaks and valleys, a phenomena not observed when using the other two bin packing heuristics.

Further analysis of intraday vs interday shows the substantial impact that a small amount of flexibility has on the performance of the system. The addition of two total days of flexibility reduces chair requirements by 17% (from 12 to 10). Variation is reduced as the appointments have the opportunity to become more efficiently packed. The diminishing value of interday scheduling is also of import. Scenarios P-F-2 through P-R-3 add more flexibility, to a maximum of 6 additional days (three on either side of the original appointment date). The number of chairs required actually increases, again due to adversarial scheduling. Variation stabilizes, but never outperforms the value observed in scenario P-L-1, which was obtained with only one day of flexibility.

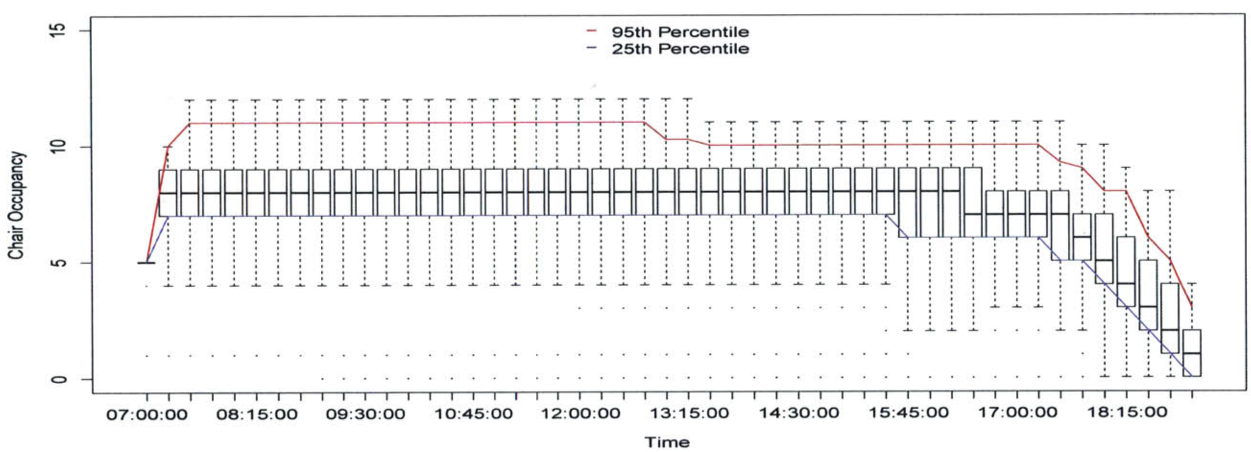

Figure 5-7: Scenario P-F-0 Modeled Chair Requirements. Physical center with Intraday Scheduling

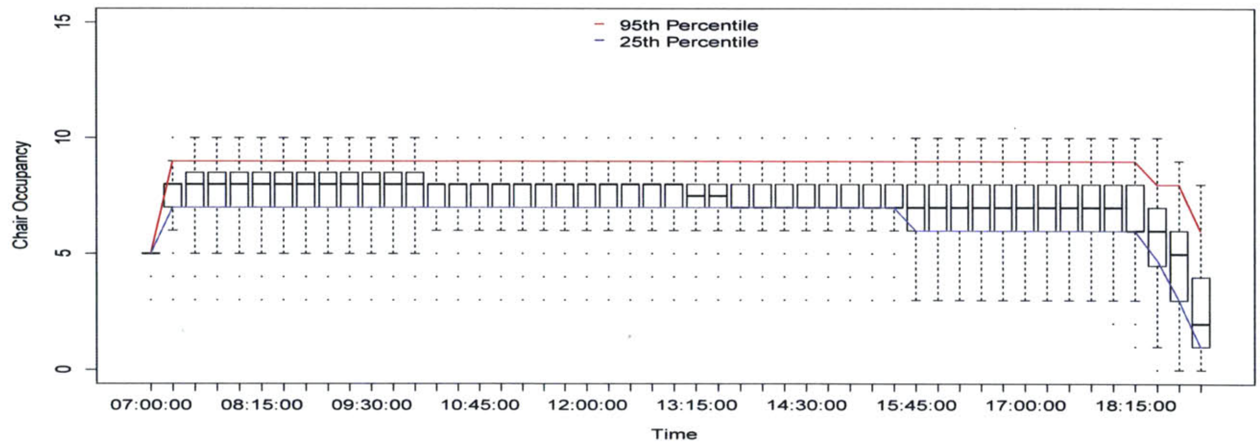

Figure 5-8: Scenario P-F-1 Modeled Chair Requirements. Physical Center with One Day

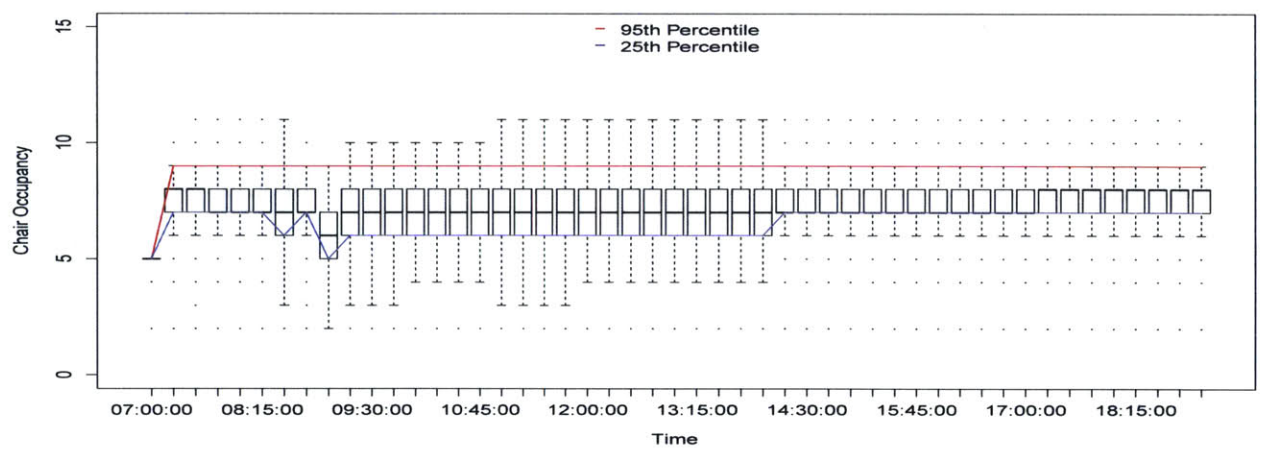

Figure 5-9: Scenario P-L-1 Modeled Chair Requirements

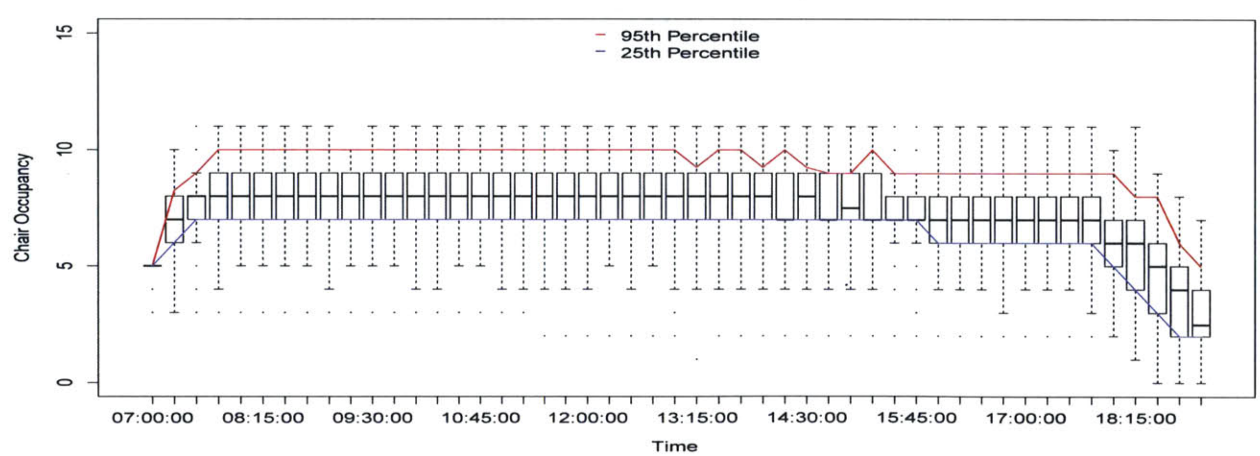

Figure 5-10: Scenario P-R-1 Modeled Chair Requirements

In addition to chair occupancy (census) of each modeled scenario, several additional metrics of interest are calculated. Two of these metrics, appointment starts per time period and nurse workload, are shown in figures 5-11 and 5-12. Both figures included here are from scenario P-L-1 (interday scheduling, one day of flexibility, last available bin packing).

Figure 5-11 is an aggregated box plot of the arrivals, $a(t)$ per fifteen minute time period. The number of nurses is assumed to be six, and thus per equation 5.1 the number of arrivals per time period is held to five, at maximum. Thus when the clinic opens, at 7:00am, there are always five patients scheduled, and this is observable in the figure. Throughout the day, arrivals fluctuate, but never surpass the imposed limit.

Nurse workload in scenario P-L-1 is shown in Figure 5-12. Workload is calculated using equation 4.2. As previously discussed in Chapter 3, some clinics experience workloads greater than one. In this scenario, the workload average is near 0.5, and never exceeds one. This outcome is partly due to the assumption of a nurse to provide system slack. While the model projects that only five nurses are required, at minimum, the figure shows workload calculated with six nurses. This additional nurse, bringing the staffing total to six, is recommended to ensure that breaks, shift changes, disruption, and other necessary tasks are addressed without overloading the staff.

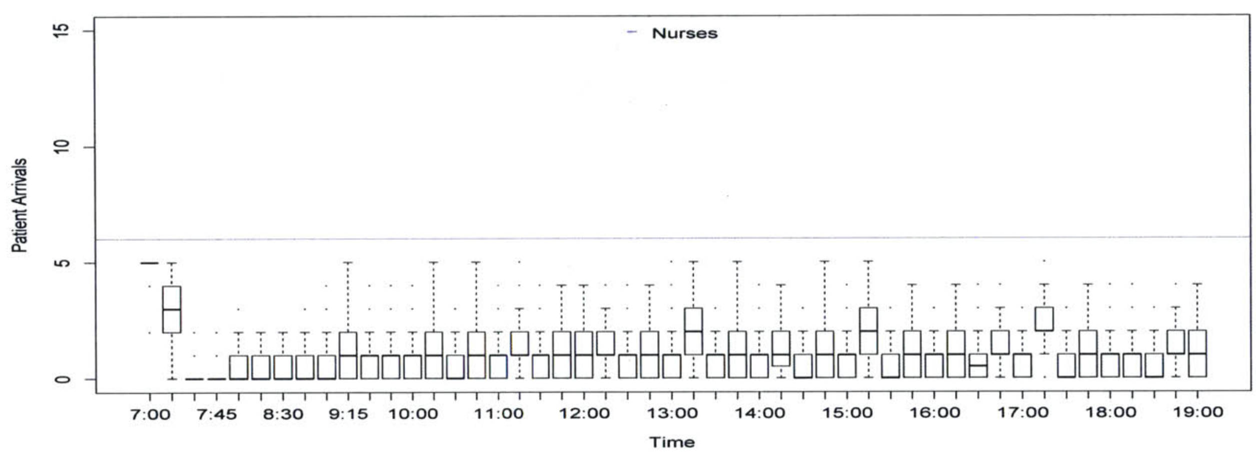

Figure 5-11: Scenario P-L-1 Modeled Appointment Starts

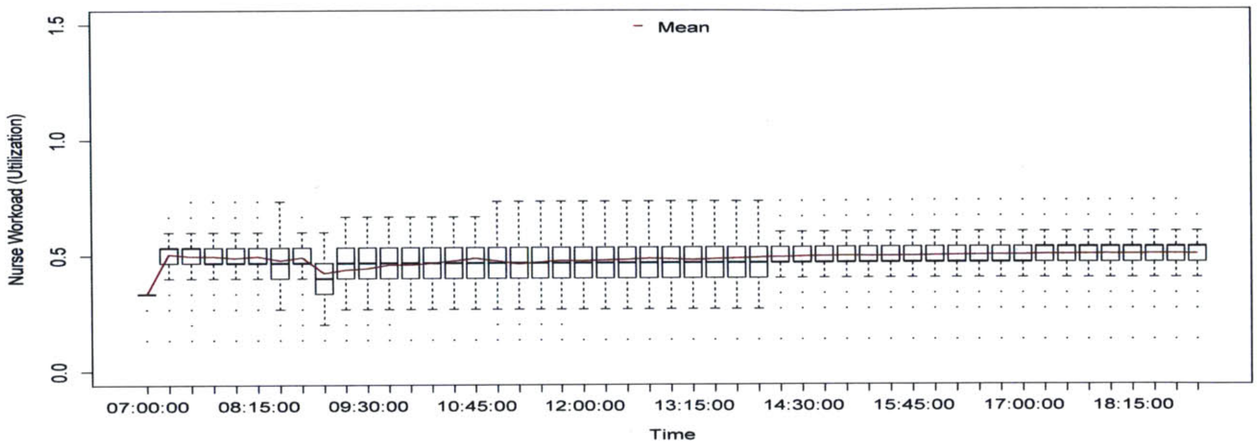

Figure 5-12: Scenario P-L-1 Modeled Nurse Workload

Discussion: Multi-location Scenarios

Multi-location scenarios M-F-0-A through M-R-0-B, each comprised of two results, one for MGH main campus and one for MGH-West, have similar results to the physical scenarios, except that there is an excess of resources due to the split capacity. The impact of interday and intraday scheduling is also similar, as is the performance of the tested bin packing heuristics.

Figure 5-13 shows the modeled census for the MGH wing of the solution, while Figure 5-14 shows the same information for the MGH-West location. The same features, for example the ramp at the start and end of the day, can be observed. It is of note that employing this intraday multi-location option requires three additional chairs (15 total) and two additional nurses (seven total), exclusive of any resources designated to provide slack capacity. It is also of note that the day to day variation observed within each of these clinics, is, as expected, higher than that found in the physical option.

Interday solutions for the multi-location scenarios are depicted in Figures 5-15 and 5-16. The comparisons between these results and those of the physical interday are also similar. The total number of required chairs is two more than the physical option (12 versus 10); the nursing requirement between the two options, however, is equivalent at five. This result once again shows the significant impact and improvement offered by the addition of a small amount of flexibility.

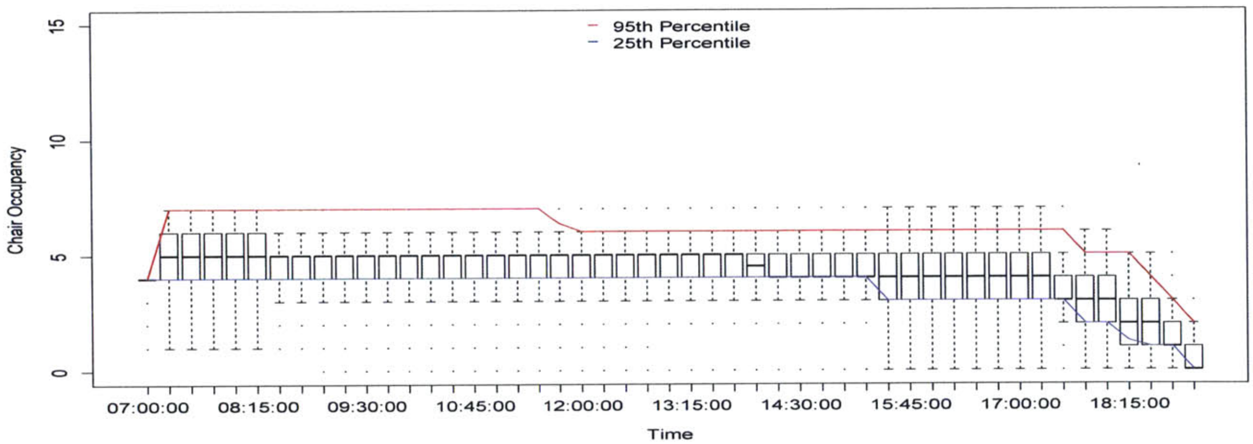

Figure 5-13: Scenario M-F-0-A Modeled Chair Requirement

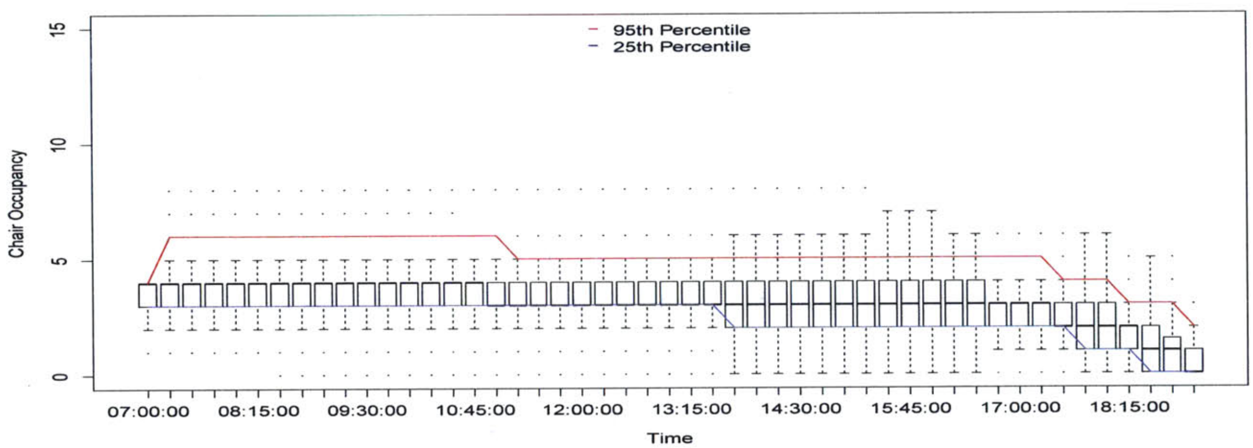

Figure 5-14: Scenario M-F-0-B Modeled Chair Requirements

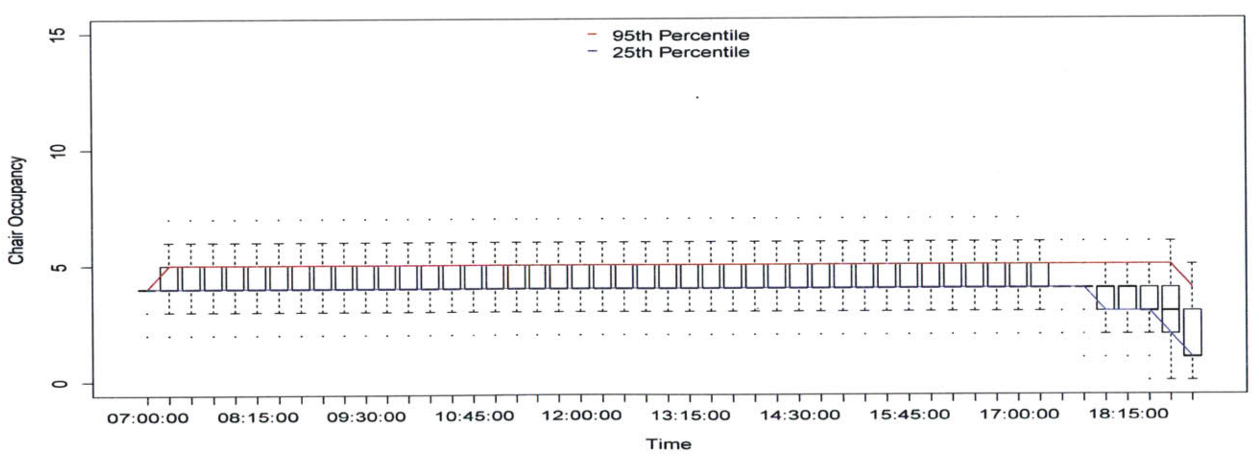

Figure 5-15: Scenario M-F-1-A Modeled Chair Requirement

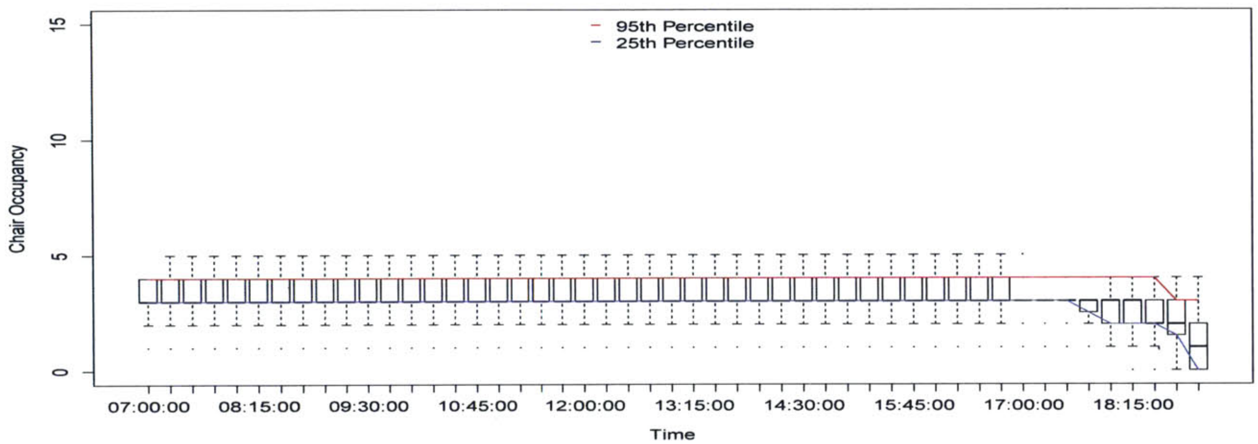

Figure 5-16: Scenario M-F-1-B Modeled Chair Requirements

5.6 Discussion and Sensitivity Analysis

Analysis of the multiple scenarios in Section 5.5 reveals that (i) centralization and improved scheduling yield significant reduction in resource requirements, (ii) interday scheduling further improves the number of required resources and (iii) only a small amount of flexibility is required in order to provide improvement, and additional flexibility has diminishing returns, or may be detrimental. These results are further tested in several sensitivity studies, below, and are coalesced into tangible recommendations in Chapter 6.

5.6.1 Sensitivity Studies

Variations of the models were created and analyzed in order to test that the process is robust. Two sensitivity study categories are examined, the first introduces inflexible, or fixed appointments, and the second includes patient preference for appointment time periods (AM or PM). Both sensitivity studies reveal that a small amount of flexibility in scheduling can alleviate the challenges introduced by adhering to patient preference over system performance. Furthermore, the sensitivity studies show that the selected scheduling algorithm is robust, and performs well under the constraint of individual patient preference.

Fixed Appointments

The previously described interday scenarios permit all appointments to move within the prescribed scheduling window, as constrained by the values of f^+ and f^-. It is possible, however, that a patient or physician may require an appointment to occur at a fixed time. This requirement could be driven by patient preference or constraint, or perhaps medical necessity. For example, a patient may only be able to obtain transportation to the clinic for an 8:00am appointment, or a physician may require an appointment to be at 7:00am for medical reasons. Although interviews and shadowing indicate that occurrences of these constrained appointments are low, it is important to understand the impact they may have on the clinic's resource requirements.

Table 5.5: Fixed Appointment Sensitivity Scenarios

Scenario Label	Fixed Appt. Percentage	Flexibility	Bin-Packing
PF0-1S	1%	0	First Available
PF0-3S	3%	0	First Available
PF0-5S	5%	0	First Available
PF1-1S	1%	1	First Available
PF1-3S	3%	1	First Available
PF1-5S	5%	1	First Available
PF2-1S	1%	2	First Available
PF2-3S	3%	2	First Available
PF2-5S	5%	2	First Available

The fixed appointment studies consider only the physical centralization option, and include the application of the interday scheduling algorithm with parameter f^+ and f^- set equal to 0, 1, and 2 days, respectively. The level of fixed appointments are set to 1%, 3%, and 5%. First available bin packing is applied in all cases, resulting in a total of nine additional models. The parameters for these scenarios are summarized in Table 5.5. For reference, box plots of chair requirements for all fixed appointment sensitivity scenarios are included in Appendix C.

The results, compiled in Table 5.6, show how a small number of inflexible appointments can impact resource requirements. The baseline data, provided in the table for reference, is from scenario P-F-0. The addition of only 3% of inflexible appointments, about two

Table 5.6: Fixed Appointment Sensitivity Study Results

Scenario	Max C	Median C	95% C	Std. Dev C	Nurse Reqt	CV Census
Baseline	12	8	11	2.37	5	0.33
PF0-1S	12	8	11	2.41	5	0.33
PF0-3S	13	8	11	2.49	6	0.35
PF0-5S	13	8	11	2.57	6	0.36
PF1-1S	12	7	10	1.70	5	0.23
PF1-3S	12	7	10	1.88	5	0.26
PF1-5S	13	7	11	2.00	6	0.28
PF2-1S	11	7	09	1.47	5	0.20
PF2-3S	11	7	10	1.61	5	0.22
PF2-5S	13	7	10	1.74	6	0.24

per day, requires the clinic to have an additional chair. Some of these increased resource requirements, however, can be overcome through the addition of scheduling flexibility. In sensitivity scenario PF1-3S, chair requirements return to previous, pre-fixed appointment levels (12) with the addition of only one day of flexibility. It is also of note that increasing levels of inflexible appointments cannot be overcome and do require additional resources. For example, scenario FP2-5S, which has 5% inflexible appointments and two days of flexibility, still requires one more chair than the unmodified scenario P-F-0.

AM-PM Patient Preference

To further explore the impact of patient preference, appointments in the data set were classified as AM (occurring prior to 13:00, exclusive), and PM, occurring after 13:00, inclusive. The assumption in this study is that in the current system, patients are, more often than not, scheduled at a time of day (AM or PM) that is, in their view, convenient and acceptable. In this study, when the scheduling algorithm is choosing between acceptable appointment intervals, as depicted in Algorithm 1, ties are broken based upon when the appointment occurred in the past. Any patient with a previously defined AM preference would be placed, when possible, in an AM appointed. The same would occur for a previously scheduled PM patient. The AM-PM appointment studies consider only the physical centralization option, and include the application of the interday scheduling algorithm with parameter f^+ and f^- set equal to 0, 1, and 2 days, respectively

Table 5.7: AM-PM Patient Preference Sensitivity Study Results

Scenario	Max C	Median C	95% C	Std. Dev C	Nurse Reqt	CV Census
Baseline	12	8	11	2.37	5	0.33
PF0-AMPM	13	8	11	2.69	5	0.37
PF1-AMPM	15	7	12	2.26	6	0.31
PF2-AMPM	14	7	11	1.91	6	0.26

The results of this study show that patient preference , as expected, impacts the centralized clinic's resource requirements (chair and nurses). [13] Peak chair requirements increase from the baseline of 12 to a new value of 13, with intraday scheduling (scenario PF0-AMPM). Adversarial scheduling is apparent in interday options, with an increase in resources observed when flexibility is set to either one or two days. Although these scenarios increase the resource requirements, they provide further evidence of the robust performance of the scheduling algorithm. Even in this somewhat unrealistic scheduling scenario, the resource requirements are only slightly increased over the baseline. Box plots for each of the three AM-PM scenarios are available in Appendix D.

[13] Note that the appointment start stagger rule is deactivated for this study, due to the tight packing observed in the AM time period.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Recommendation and Conclusions

6.1 Operational Recommendations

Based upon the modeling results and in-depth analysis of the MGH infusion system, the creation of a new, main-campus multi-specialty infusion clinic is recommended. The new clinic will work in conjunction with a satellite location at MGH-West, and both will be served by a single administrative and staffing unit, described below, the Administrative Recommendations Section. The proposed main-campus clinic shall include a minimum of 12 chairs (based upon the intraday scheduling results, for conservatism), a reduction of 57% over the current state. Four additional chairs shall be included to account for growth, emergent appointments, and intraday disruption. Hours of operation are to be as modeled, 7:00am to 7:00pm, Monday through Friday. Weekend appointments, if necessary for special cases, can be made via appointment, but the clinic is not proposed to maintain regular weekend operating hours. Staffing of the clinic is derived from the number of chairs and max patient to nurse ratio of 2.5:1. Based upon the minimum number of chairs (12), the clinic requires a staff of six nurses[1], a reduction of 60% compared to the current state. Five of these nurses are required, and an additional nurse is recommended to provide for system slack and allow for breaks and disruption. The number of nurses on staff, however, is higher than six, due to the operating hours of the clinic. As the clinic will be open 12 hours per day, five days

[1] As described in Chapters 3 and 4, the nurse requirement only relates to treating patients (providing infusion). If additional tasks are required of nurses, such as scheduling or other administrative tasks, this number will need to be increased.

per week, a total of nine nursing FTE (full time equivalent) staff is recommended[2].

The clinic must also make use of the scheduling algorithm used in the modeling effort described in Chapter 5. Information technology systems shall be modified to incorporate the algorithm and the concept of appointment flexibility. Based upon the evaluation of the three bin packing heuristics, the implemented algorithm should make use of the first available bin packing heuristic. While the last available heuristic may perform slightly better in some scenarios, the first available option ensures a declining patient census at the end of the day, which ensures a better working environment for the nursing staff. It is critical to note that creation of a centralized clinic without improved scheduling techniques will lead to many of the same operational issues observed today.

As of May 2016, MGH has started the process to construct this new, multi-specialty main-campus infusion clinic. As previously discussed, due to the long lead time necessary to secure space, a temporary multi-location solution is required. For the temporary two-location "virtual" clinic, a split of capacity between the current MGH Neurology clinic and MGH-West is recommended. The Neurology clinic has room for two additional chairs, bringing the total to ten. MGH-West has eight chairs[3]. Staffing requirements are five nurses for the MGH location, and four for MGH west. This results in a total of nine nurses (13.5 FTE). During the temporary period, both clinics are to be served by a single staffing unit, including nurses and administrative personnel.

It is not surprising that the multi-location solutions require more resources to serve the same number of patients as a single location, as the resources are not centralized and some inefficiencies are realized. The two temporary two location clinic will, however, provide a near term solution that will alleviate the access issues currently observed in the clinic's today. Administratively, however, the any two location clinic will be more complex to manage, and must be supported by the single scheduling and PA unit, discussed below. Both the physical and virtual clinics will address the issues currently existing within the non-oncology infusion clinics. Patient access will improve through increased capacity and streamlined administration (below). Same day and emergent treatment will be available,

[2] A normal working week is 40 hours. Thus the clinics requires 1.5 times FTEs as the number of personnel: $1.5 \times 6 = 9$.

[3] The modeling effort was not used to develop this chair requirement, the capacity was already in place

which will reduce or eliminate avoidable admissions for infusion. Finally, the pooling of resources provides for more efficient use of clinical, physical, and administrative resources.

6.2 Administrative Recommendations

In chapter four, the two impediments to patient access, administrative and scheduling, were discussed. While the pooling of clinical resources and the addition of advanced scheduling techniques increase capacity, neither addresses the administrative delay associated with prior authorization. In order to ensure the timely processing of authorizations, the creation of a centralized non-oncology infusion prior authorization team is recommended.

The new team, to be composed of three or four members, shall be modeled after the similar effort underway in the MGH Cancer Center. In that group, a small staff of experts manages all requests for authorization. The centralization of this staff allows for specialization and the development of expertise in both payer and drug requirements. Using the Cancer Center team's performance as a guideline, the use of this specialized team should result in typical authorization processing of three days or less, with a maximum of seven days expected on extremely complex cases.

Also, the specialization and expertise developed on this team will allow for significant reduction in write off due to authorization denial. The PA team will also be able to build knowledge in the patient assistance programs offered by drug manufactures, and offer them to patients, a further means of increasing reimbursement. Although hospital financial data cannot be specifically included in this document, a 10% reduction in write offs for the major non-oncology infusion drugs results in more than a million dollars in savings. Additionally, the 15-20 part time staff that currently perform authorizations, some of which are nurses, will be unburdened and made available to work on other tasks.

Most importantly, however, patient access will be significantly improved. The expected wait time for an authorization will decrease from weeks to days. As the scheduling system and physical capacity will be available to accommodate new patients in a timely manner, wait time for a new appointment is expected to be no more than 14 days, a 72% reduction from the 2014 average. The updated patient flow, previously presented in Chapter 3, is depicted below, in figure 6-1.

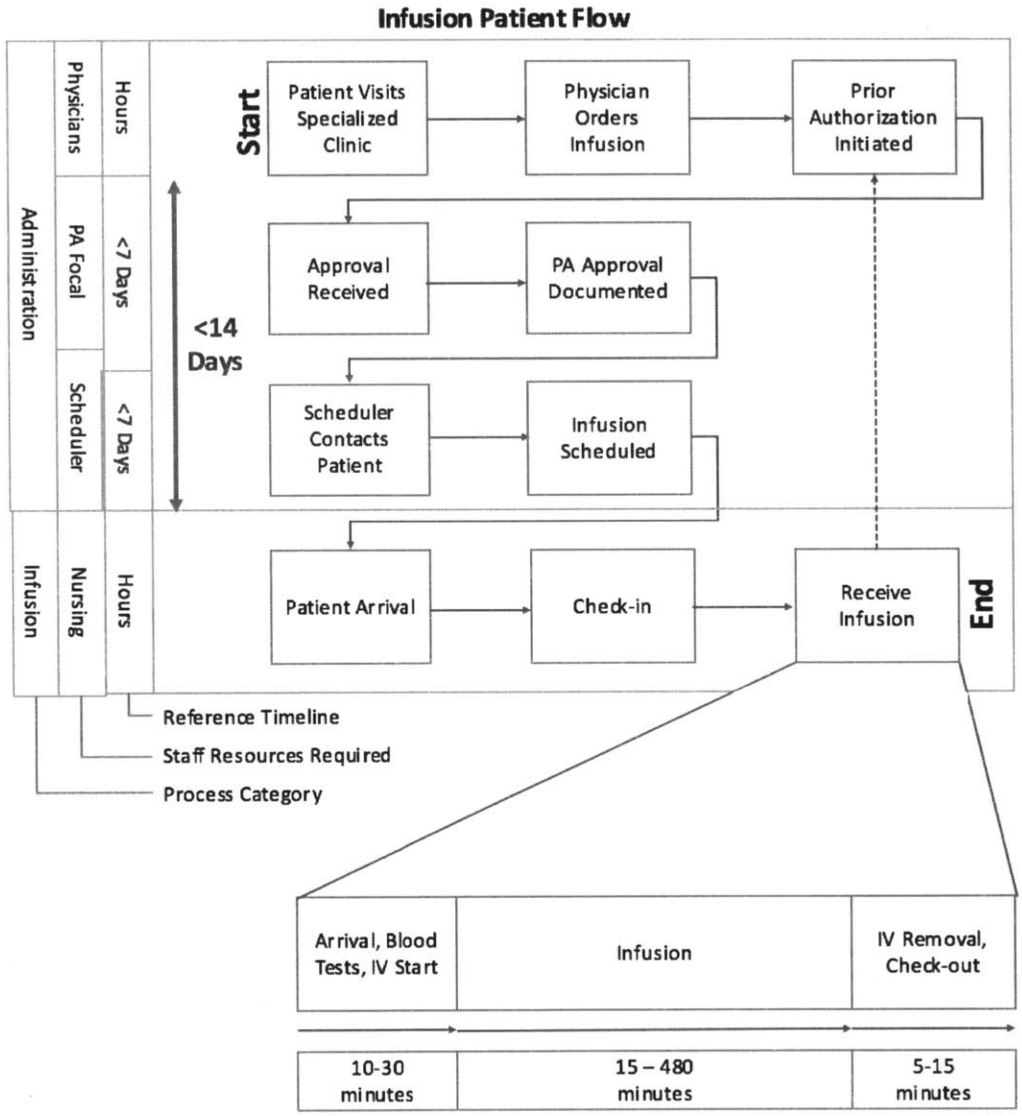

Figure 6-1: Basic overview of infusion patient flow, from physician appointment to treatment. Required resources and reference time line are indicated.

6.3 Potential for Future Research

Throughout the project, several ideas for further exploration were developed. Some of these areas are academic, and are focused on improving the performance of the scheduling process. Others, however, are potential operational areas that could be studied and perhaps improved in the new infusion clinic.

- **Offline versus online:** The recommended scheduling algorithm is on-line, that is it schedules patients in real-time as they arrive into the scheduling system. Evaluation of off-line scheduling, which would show a theoretical optimum, versus the proposed on-line scheduling, would show the relative performance of the selected heuristic, a comparison called the competitive ratio. By researching this benchmark, additional heuristics could be developed and evaluated.

- **Batched scheduling:** While implementing true offline scheduling is impossible, given that the appointment population is unknown, partial offline scheduling is potentially a realistic way to further improve scheduling performance. Rather than schedule each patient as they arrive, batches of patients could be scheduled together, in an offline manner, perhaps every day. Implementing this process would present a significant operational and cultural shift, but it may be a way to gain even more efficiency without adding physical capacity.

- **Integrated PICC Scheduling:** Many infusion patients, in particular those receiving recurring antibiotic treatment, require a peripherally inserted central catheter (PICC). PICCs are placed by a specialized team that also experiences severe operational challenges resulting delays of several days for PICC insertion. The PICC service does not allow for patients to be scheduled; they are served on a first come, first served basis. The creation of a scheduling mechanism between the infusion center and the PICC service will further decrease patient delay and help to avoid hospital admission.

- **Pharmacy:** The existing infusion clinics are served by multiple pharmacies. Centralization of infusion resources will impact the pharmacy. As of the time of writing, a new centralized specialty pharmacy is planned for MGH. It is possible to incorporate pharmacy constraints, such as staffing and demand for a single drug, into the scheduling process. Doing so will require research into these constraints, algorithm revision and testing, and updating the clinic's scheduling software.

- **Prior authorization analysis:** One of the challenges encountered on this project was the lack of data on administrative processes, in particular prior authorization. The creation of a centralized authorization team provides the opportunity to track and monitor authorization progression, approval, and rejection. These data, which could be analyzed by payer, drug, diagnoses, and many other categories, would be valuable to MGH as it works to continuously improve administrative processes and lower costs.

- **Operational studies:** Currently, none of the clinics at MGH track actual chair time. The centralized clinic provides the opportunity to implement more robust operational metrics and data collection. Tracking of chair time, as an example, allows for the identification of treatments that over utilize or under utilize the scheduled time. Highly disruptive treatments could be targeted for process improvement, or, the scheduling system could be updated to reflect the actual time required.

- **Avoidable admission:** Whether or not a hospital admission is avoidable is highly subjective. Additionally, data indicating if a clinician determines an admission to be avoidable, are not kept in any hospital system. Although attempts were made during this project to quantify avoidable admission for infusion treatment, additional work could be done to better quantify avoidable admission. Additionally, studies could be performed to expand the work beyond infusion admission cases.

6.4 Conclusions

This project was charted with a short but complex goal: improve patient access to non-oncology infusion. Through the pooling of resources and the implementation of advanced scheduling techniques, resources are used more efficiently, and that initial goal is achieved. Additionally, through the centralization and standardization of prior authorization and other administrative processes, patient access is further improved and significant financial benefits are realized through reduced write-offs. Finally, the proposed centralized infusion clinic will provide the emergent capacity necessary to greatly reduce or eliminate the avoidable admission of patients to the hospital to receive infusion.

Appendix A

Infusion Treatments and Appointment Codes

The tables below, A.1 and A.2, contain the EPIC infusion appointment codes used in this study, organized by clinic. Treatment codes correspond to the process name field in the EPIC data extraction, for the period of study (July 2014 to July 2015). While most clinics code appointments by the drug or treatment being administered, some appointments are coded as "INFUSION" or "OTHER." Note that injection appointments are rarely coded by the drug being administered, and are instead labeled as "INJECTION." Some of the infusion codes are also coupled with department or organizational labels, including "MGH NEURO","MGPNCPCZ8",and "MGH GI." Also of note, some clinic indicate referral treatments in the appointment code, "TX" and "NON-TX" for transplant and non-transplant, respectively.

For the studies described in Chapters 4 and 5, appointments classified in EPIC as "Completed" and "No-Show" were used. Appointments classified as "Cancelled" were excluded as they were assumed to be rescheduled within the interval of study.

Table A.1: Infusion Treatment Codes, Part I

Clinic	Infusion Visit Type
Allergy	XOLAIR
Blood Transfusion Service	25 % ALBUMIN INFUSION
	AUTO PLATELET
	AUTO WB
	BLOOD SPECIMEN
	CATHETER DRESSING FLUSH
	FERAHEME INJECTION
	FRESH FROZEN PLASMA INFUSION
	IMMUNOGLOBULIN INFUSION
	IMMUNOGLOBULIN INFUSION- LONG
	IRON INFUSION
	LEUKODEPLETION
	OTHER
	PLASMA EXCHANGE
	PLATELET INFUSION
	PROLASTIN INFUSION
	RED BLOOD CELL TRANSFUSION
	RED CELL EXCHANGE
	RESEARCH WB
	VEIN CHECK
Dermatology	INFUSION
	INJECTION
Endocrine	AREDIA PAMIDRONATE
	CORTROSYN STIM TEST
	INFUSION
	INJECTION
	ZOLENDRONATE
GI	INFUSION - OTHER
	INFUSION - REMI
	INFUSION - VEDO
	INFUSION- MGH GI
	INJECTION- MGH GI
NeuroEndocrine	CORTROSYN STIM TEST
	CORTROSYN STIM TEST MGH NEURO
	INJECTION
	ZOLENDRONATE

Table A.2: Infusion Treatment Codes, Part II

Clinic	Infusion Visit Type
Neurology	ENZYME REPLACEMENT
	ENZYME REPLACEMENT MGPNCPCZ8
	HYDRATION MGPNCPCZ8
	INFUSION CHEMO MGPNCPCZ8
	INFUSION IVIG MGPNCPCZ8
	INFUSION RITUXIMAB MGPNCPCZ8
	INFUSION STEROID
	INFUSION STEROID MGPNCPCZ8
	INFUSION TYSABRI
	INFUSION TYSABRI MGPNCPCZ8
	INFUSION VENOFER MGPNCPCZ8
	INJECTION BOTOX
Rheumatology	ACTEMRA
	ANTIBIOTIC
	AREDIA PAMIDRONATE
	BENLYSTA
	CORTROSYN STIM TEST
	CYCLOPHOSPHAMIDE
	INFUSION SOLUMEDROL MGHRHEUYAW
	MAGNESIUM
	ORENCIA
	PEGLOTICASE
	RECLAST
	REMICADE
	REMICADE WITH MEDICATION
	RITUXAN 1
	RITUXAN 2
	TEACHING
Transplant	INFUSION
	INFUSION NON-TX
	INFUSION TX
	INJECTION
	INJECTION ESRD
	INJECTION NON-TX
	INJECTION TX

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

Avoidable Admit Data Summary

This appendix describes the estimates for avoidable infusion admissions and the steps taken to validate those estimates using hospital records and data. The results of these estimates are summarized in Table 4.2.

B.1 Overview

Avoidable Admit Definition: an admission to the hospital, where the patient is healthy enough to receive treatment in an outpatient setting, but is admitted because no outpatient resources are available. In addition to the infusion or transfusion, avoidable admits may require services such as PICC (peripherally inserted central catheter) placement and home infusion coordination.

Savings Calculations:

Potential bed days saved are calculated by:

$$Est_{low} = Est_{admit} \times P_{avoidable-low} \times LOS_{avg}, \tag{B.1}$$

where Est_{admit} is the estimate of admissions, $P_{avoidable-low}$ is the lower end of the percentage of admissions that are avoidable, and LOS_{avg} is the average length of stay for patients located in the hospital data. The same approach can be used to calculate the high end of the estimate, by replacing $P_{avoidable-low}$ with $P_{avoidable-high}$, the upper end of the percentage of avoidable admissions.

B.2 Patient Categories

B.2.1 Antibiotics - OPAT Start

Description: Outpatient Antibiotic Therapy (OPAT). Broad category of admissions describing patients that need to be seen in a medical setting to receive a first dose of antibiotics. They also often need to have home infusion service arranged. Some may need a PICC line placed.

Estimate: One patient per week, 50 per year. All of these admissions are considered avoidable.

Validation: The OPAT program keeps record of all of these patients. The ID doctors used that data to make the avoid-ability estimate; a search of admission data was performed looking for patients that received a PICC line, stayed less than 3 days, and had seen an ID doctor. This is a conservative method of search recommended by the Infectious Disease clinic. Using this search 26 patients were identified, average LOS was 1.59 days. LOS calculated by (discharge date – admission date).

Given a sample patients for osteomyelitis, we tried to find a general rule to identify avoidable admissions. Many test searches were run using different J-codes and different restrictions with length of stay (LOS), department, PICC line, IV starters, etc. They all failed. ID Subject matter experts agreed that the population is too diverse to be able to catch it retrospectively in data. These numbers have been verified by many doctors within ID clinic.

B.2.2 Antibiotics - Cellulitis

Description: Broad category of admissions describing patients with cellulitis/abscess that need antibiotics. Treatment course often includes two doses of antibiotics given over two days (one hour per treatment).

Estimate: 250-750 patients admitted per year; 25%-75% are considered avoidable by hospitalists and ID clinic subject mater experts.

Validation: A search of EPSi (finance) data was performed, looking for patients with a primary diagnosis of cellulitis (ICD9 Code 682.X), and a LOS of less than 4 days. In 2014, 646 admissions met these criteria. The median LOS was 1.34 days, average 1.58 days, range 0.38-3.82 days. Note: The 4 day maximum is based upon confirmed avoidable admissions submitted to the avoidable admit inbox. LOS is calculated by (discharge time – admission time).

B.2.3 Antibiotics - Adult Cystic Fibrosis

Description: Patients within the Adult Cystic Fibrosis program that need to be admitted to the hospital for a course of antibiotics.

Estimate: Per year, 110-120 patients in the program need to be admitted. A single doctor from the Cystic Fibrosis Clinic oversees all of the patients and estimates that 15-20% of the admissions could be avoided. The length of stay is estimated to be four days (minimum).

Validation: A search for this population was not performed due to the small patient population; the Cystic Fibrosis clinic has data necessary to validate.

B.2.4 Emergent Infusion - Rheumatology

Description: Rheumatology infusion patients requiring emergent treatment (within 1-3 days).

Estimate: 10 patients per year, all of which are assumed to be avoidable.

Validation: A search of inpatient data was performed; LOS was limited to 3 days. Results included patients that had visited Rheumatology and received an inpatient infusion (identified with drug J-codes). Seven patients were identified. The average LOS was 2.43 days.

B.2.5 Emergent Blood Transfusion - BTS

Description: Patients referred to Blood Transfusion Service (BTS) needing an emergent appointment for blood transfusion (within 1-3 days).

Estimate: The BTS estimates that they turn away 1-2 patients per week; it is assumed that the patients are then admitted to receive treatment. All of these admissions are assumed to be avoidable.

Validation: As of yet, no search validation of this data has been performed as a methodology to identify them in the hospital records has not been identified. Note: over a three-week period the avoidable admit inbox received indication of six patients turned away by the BTS. The lower end of the estimate (50/year) is used for conservatism.

B.2.6 Allergy - Chemo Desensitization

Description: Oncology patients referred to allergy for chemo desensitization. Desensitization is an intensive process (8-10 hours minimum) where chemo drugs are slowly infused and the patient is closely monitored.

Estimate: Chemo desens patients are closely tracked, in 2014 there were 180 admitted. Allegy subject matter experts estimate that between 50% and 75% of the admissions are avoidable. LOS is estimated to be 1.5 days.

Validation: The Allergy clinic maintains a detailed log of all chemo desensitization. In FY 2014 there were 362 admissions.

Appendix C

Detailed Modeling Results

This appendix contains detailed modeling results for each of the scenarios and sensitivity studies discussed in Chapter 5. Patient census box plots for the period of study (July 15, 2014 - July 15, 2015) are presented for each scenario.

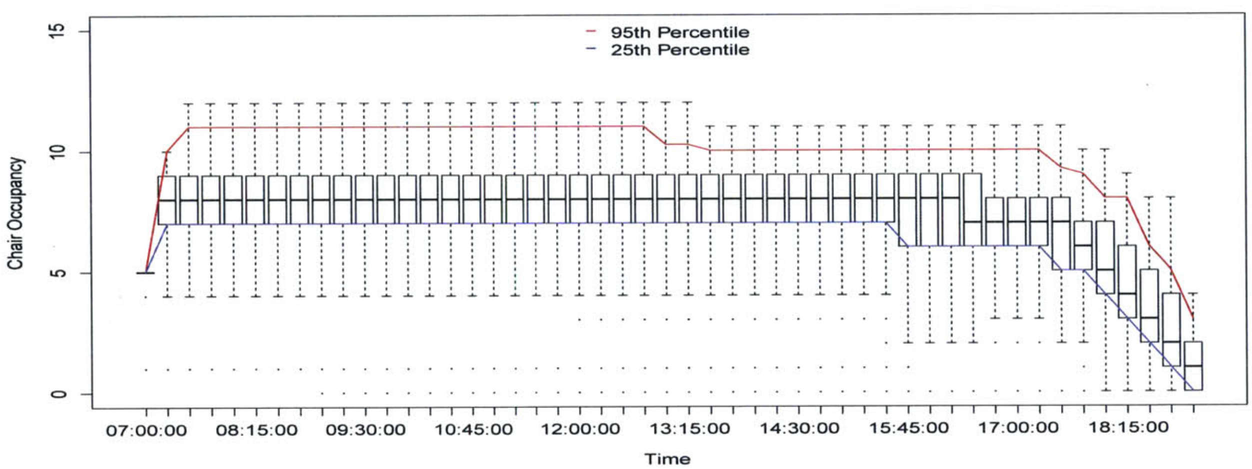

Figure C-1: Scenario P-F-0 Modeled Chair Requirements

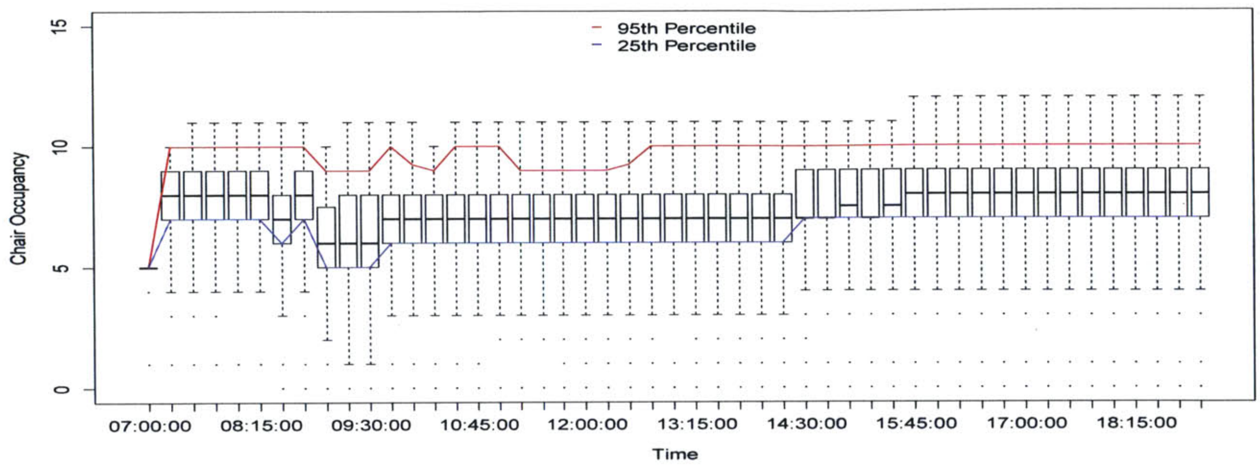

Figure C-2: Scenario P-L-0 Modeled Chair Requirements

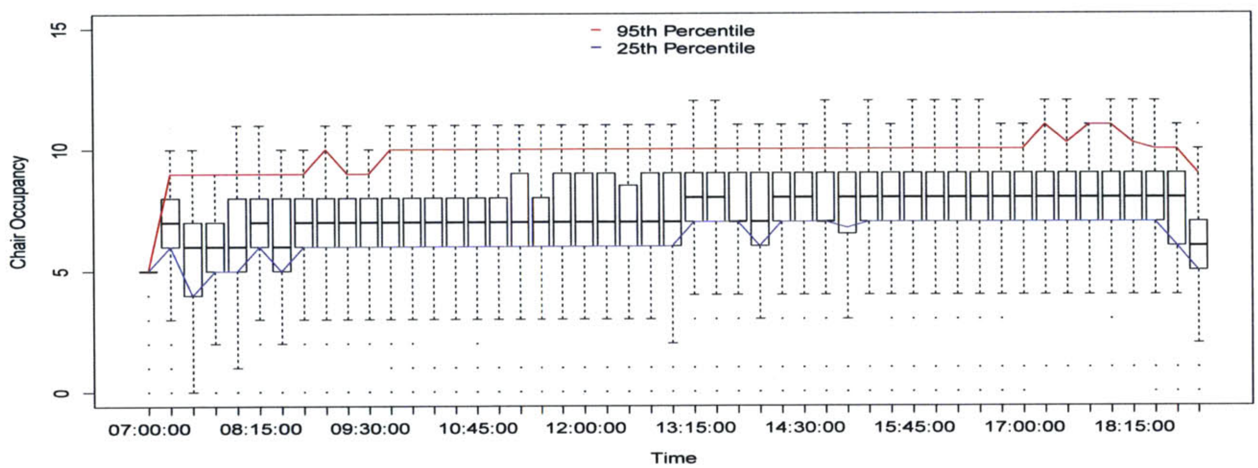

Figure C-3: Scenario P-R-0 Modeled Chair Requirements

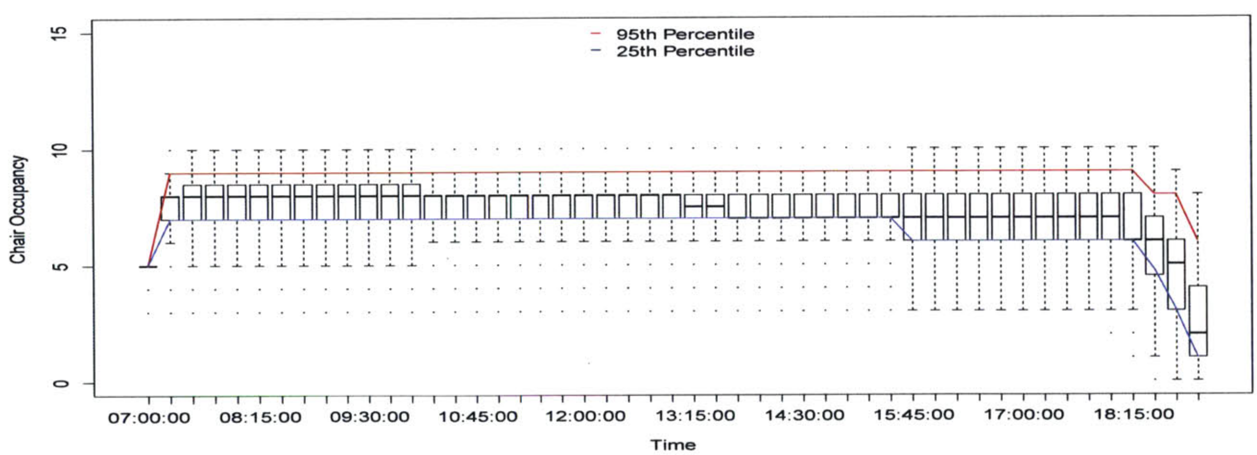

Figure C-4: Scenario P-F-1 Modeled Chair Requirements

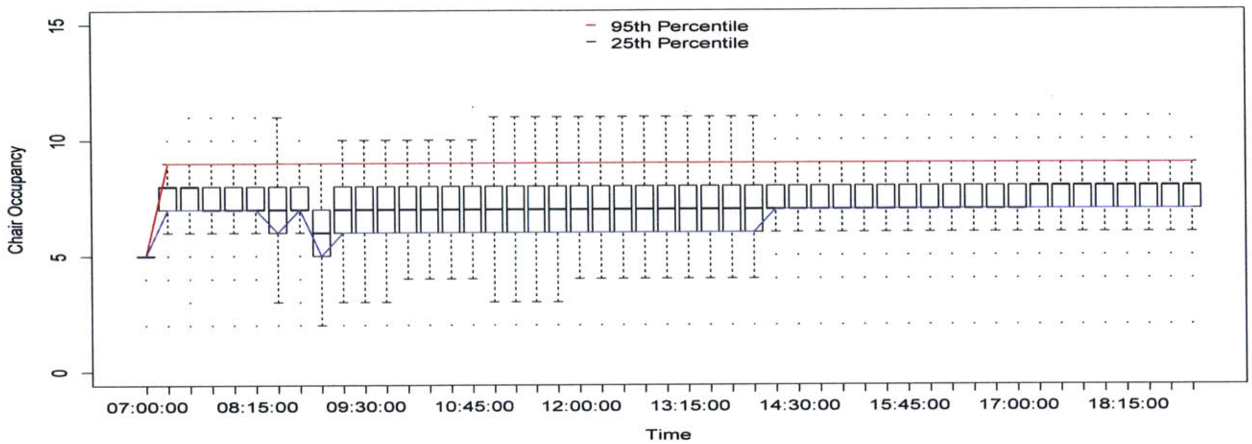

Figure C-5: Scenario P-L-1 Modeled Chair Requirements

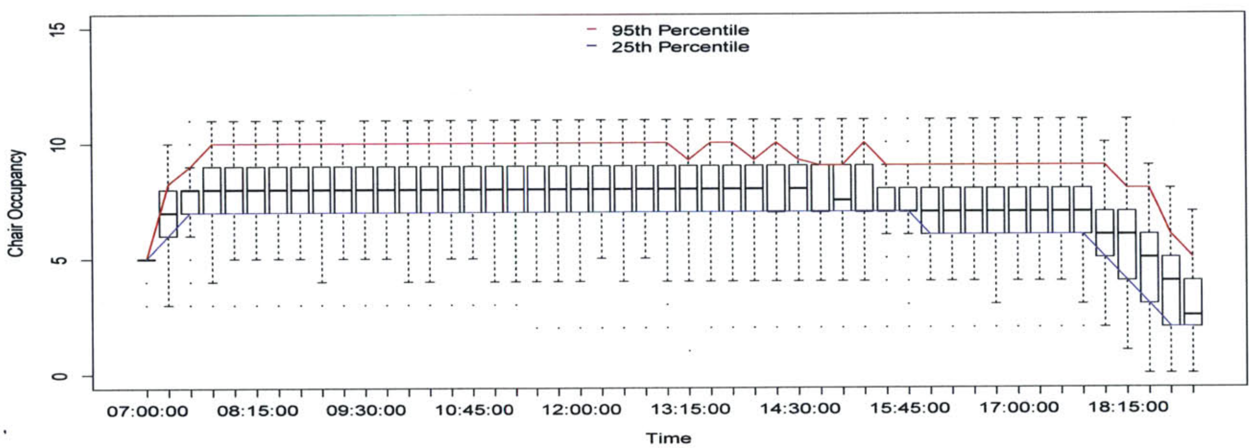

Figure C-6: Scenario P-R-1 Modeled Chair Requirements

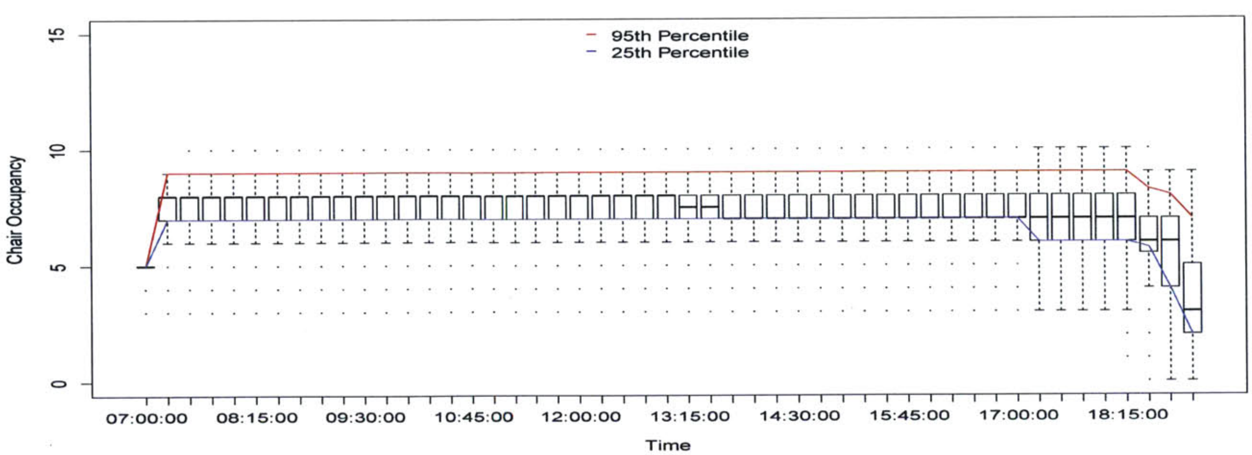

Figure C-7: Scenario P-F-2 Modeled Chair Requirements

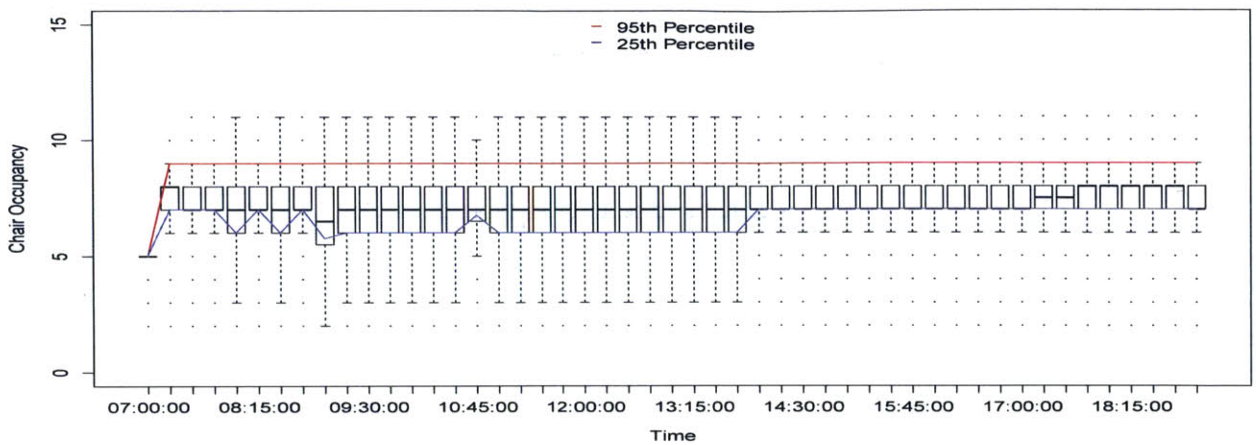

Figure C-8: Scenario P-L-2 Modeled Chair Requirements

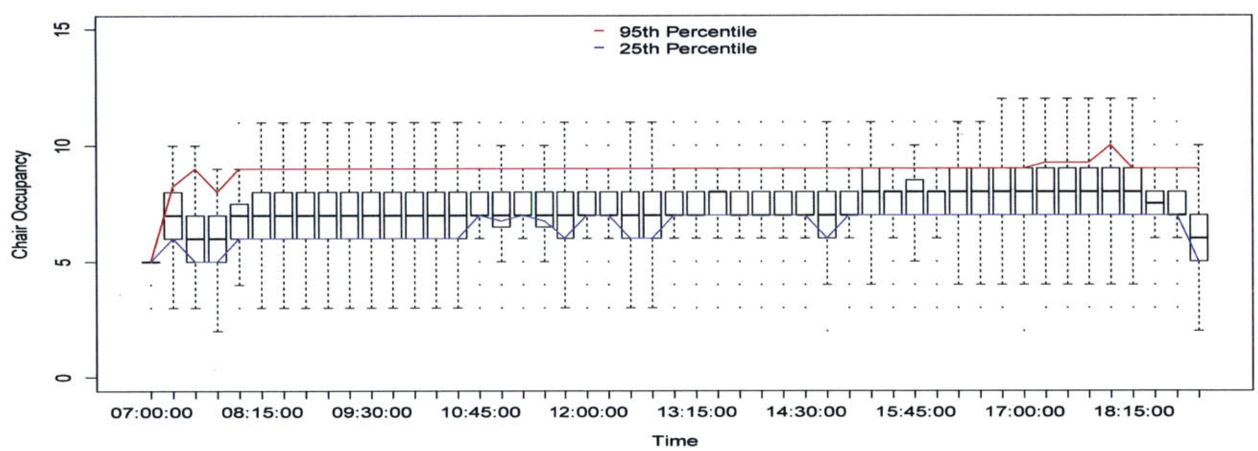

Figure C-9: Scenario P-R-2 Modeled Chair Requirements

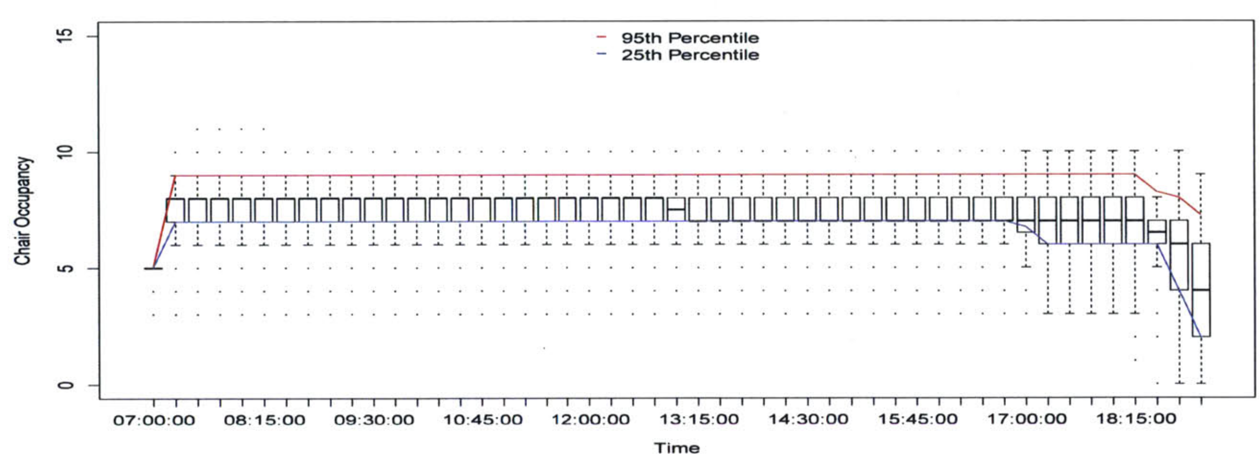

Figure C-10: Scenario P-F-3 Modeled Chair Requirements

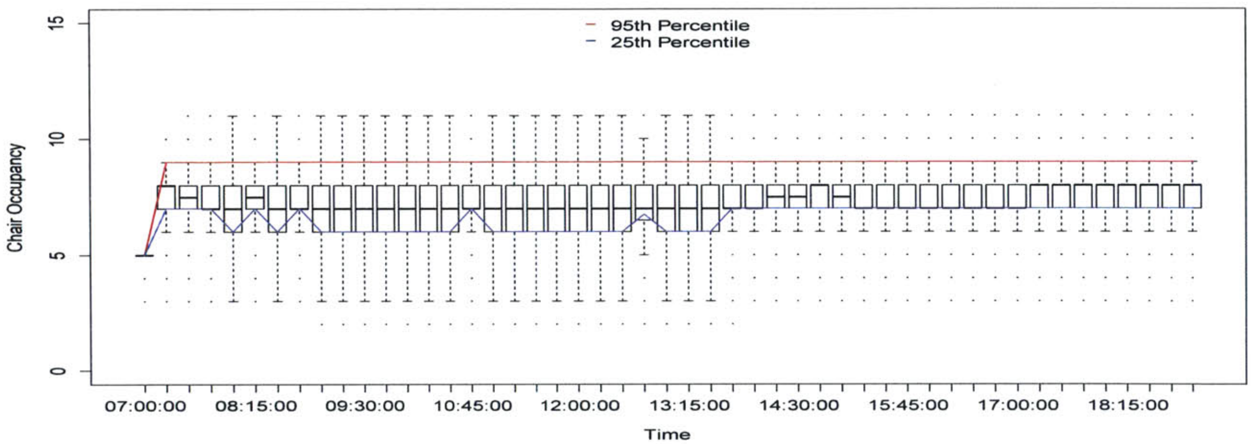

Figure C-11: Scenario P-L-3 Modeled Chair Requirements

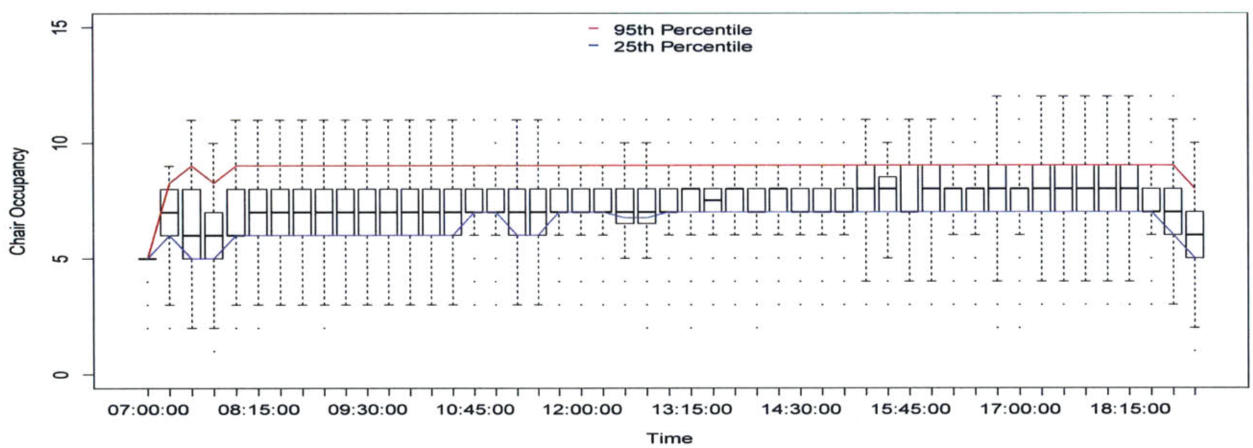

Figure C-12: Scenario P-R-3 Modeled Chair Requirements

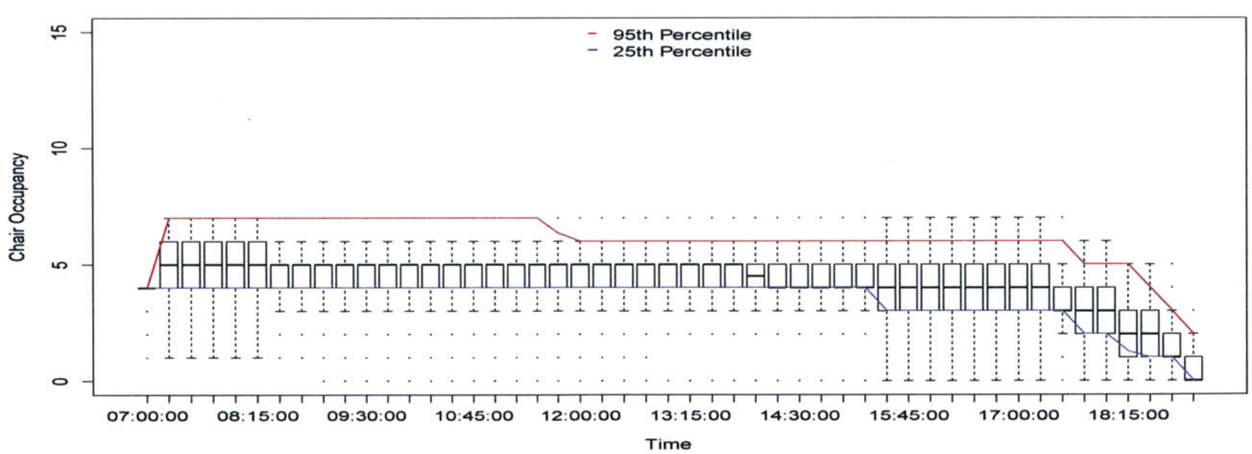

Figure C-13: Scenario V-F-0A Modeled Chair Requirements

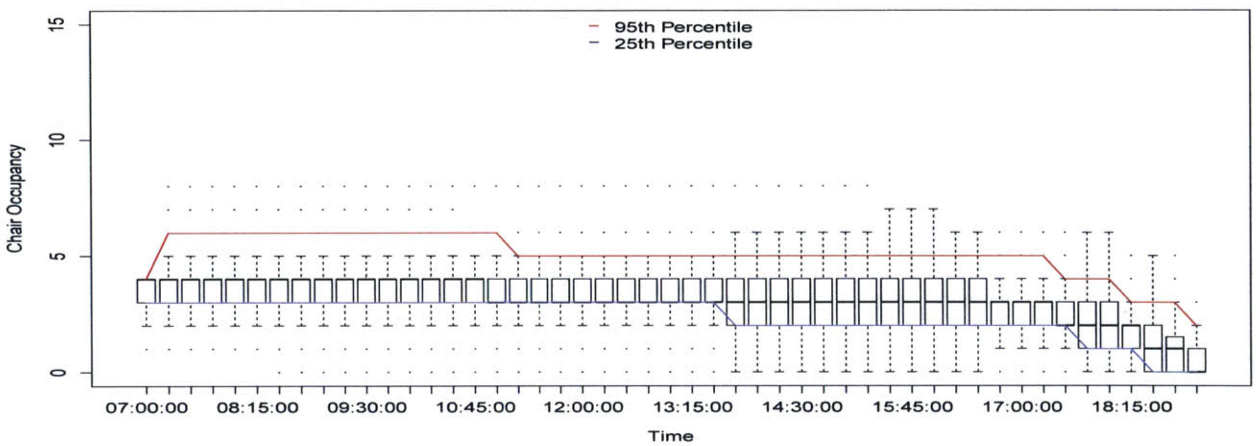

Figure C-14: Scenario V-F-0B Modeled Chair Requirements

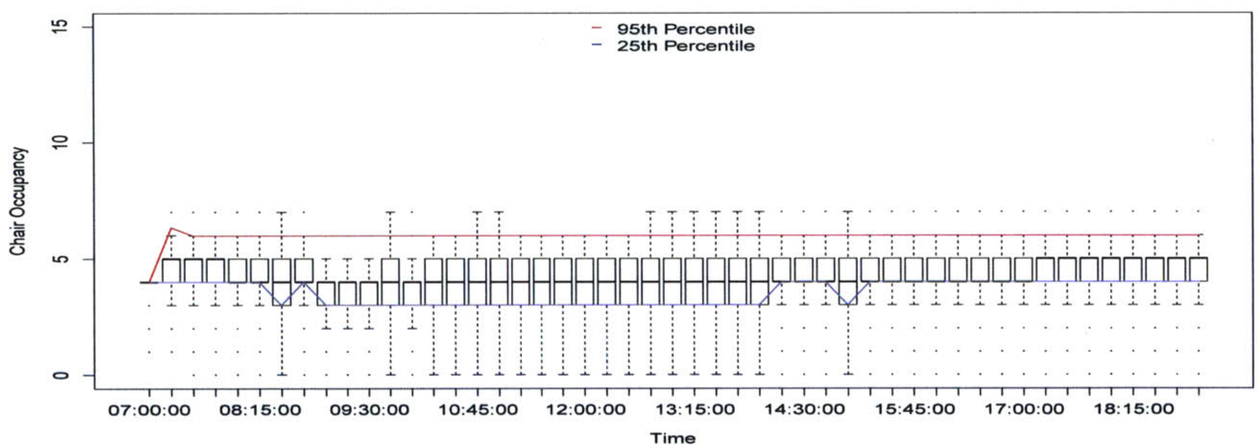

Figure C-15: Scenario V-L-0A Modeled Chair Requirements

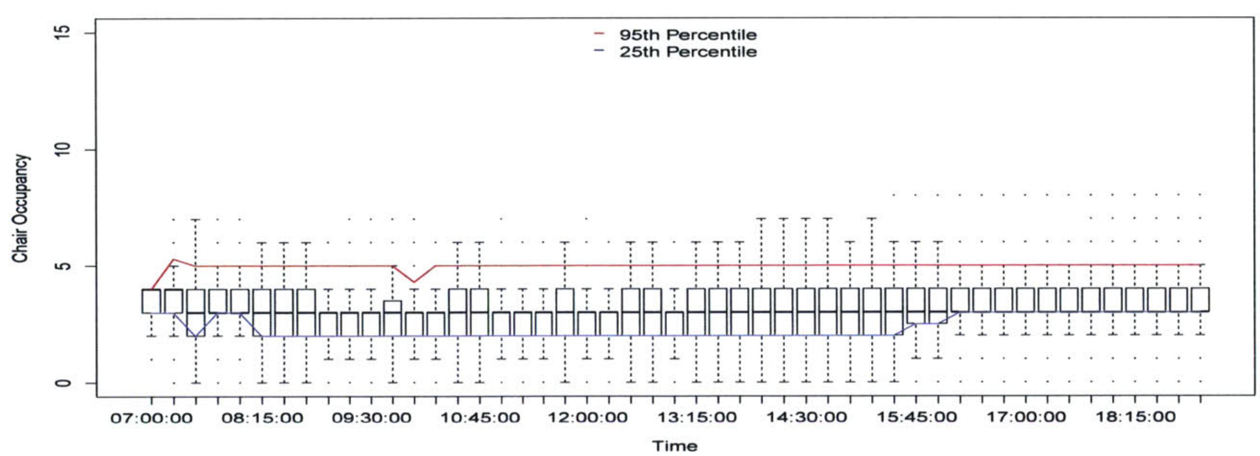

Figure C-16: Scenario V-L-0B Modeled Chair Requirements

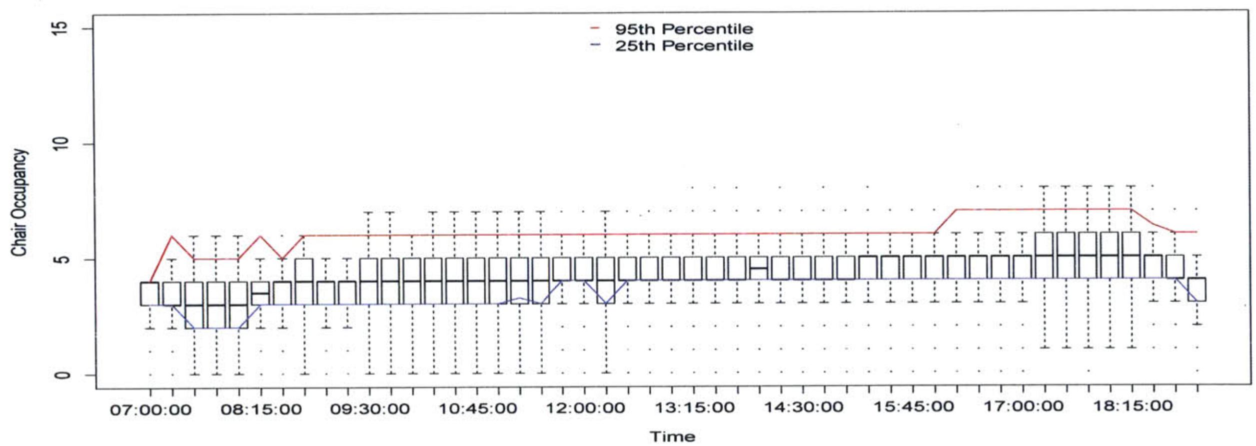

Figure C-17: Scenario V-R-0A Modeled Chair Requirements

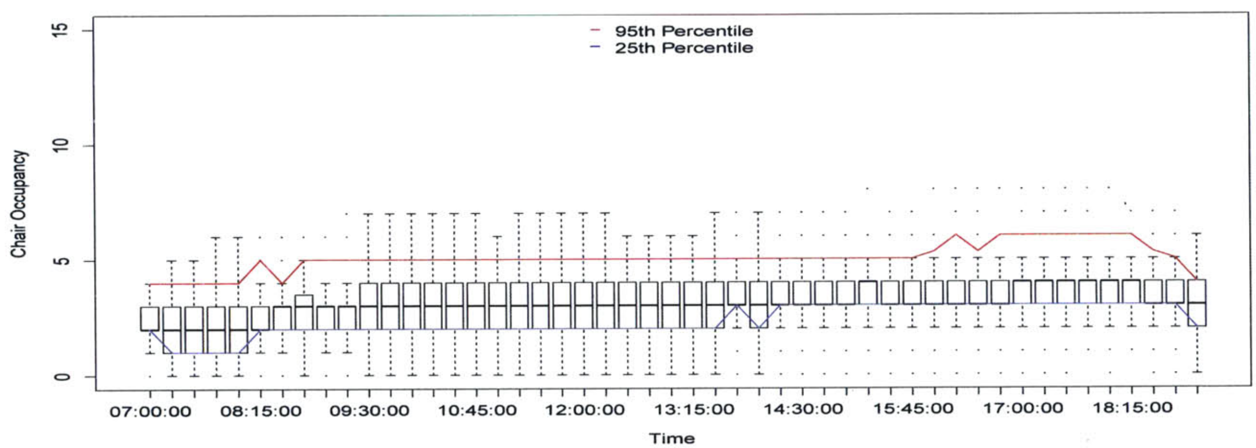

Figure C-18: Scenario V-R-0B Modeled Chair Requirements

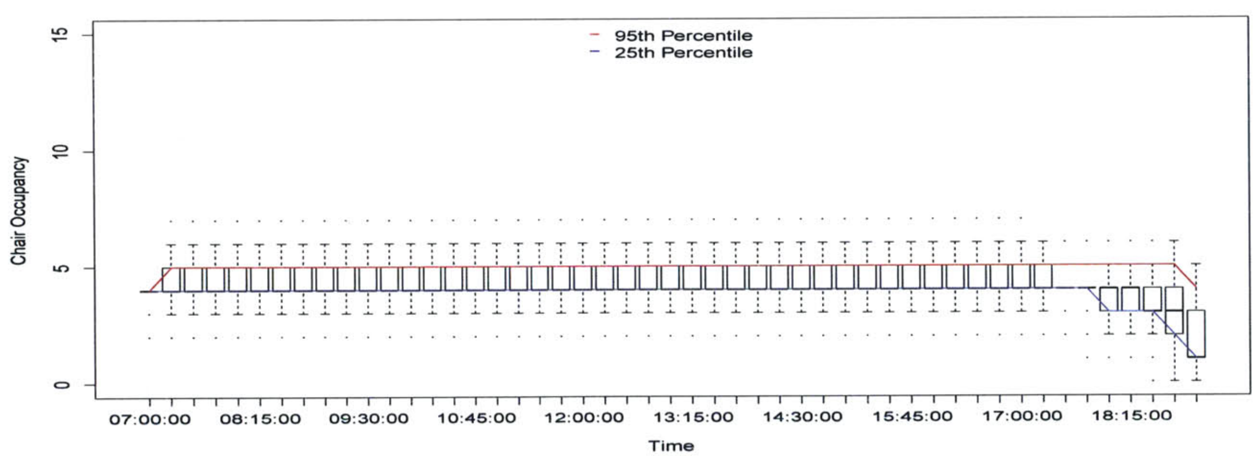

Figure C-19: Scenario V-F-1A Modeled Chair Requirements

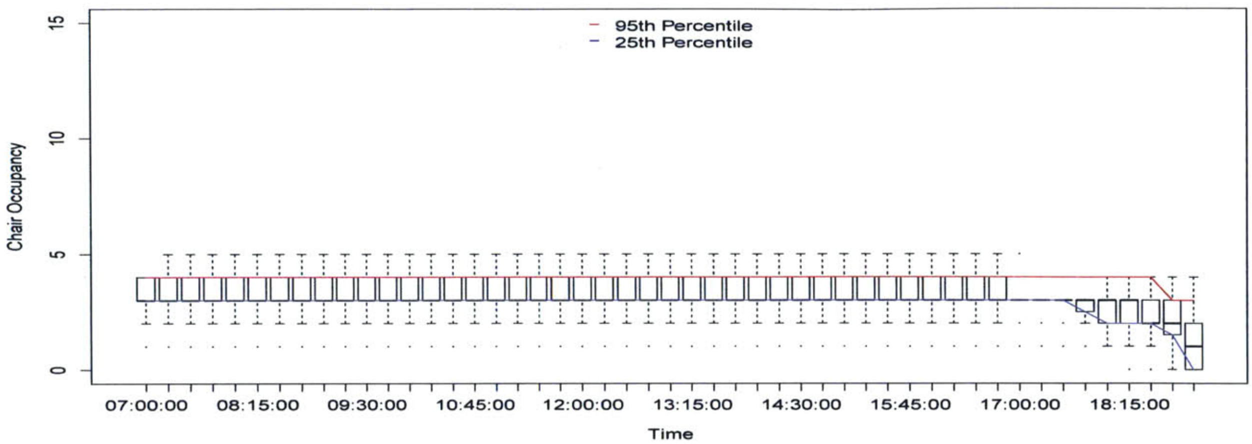

Figure C-20: Scenario V-F-1B Modeled Chair Requirements

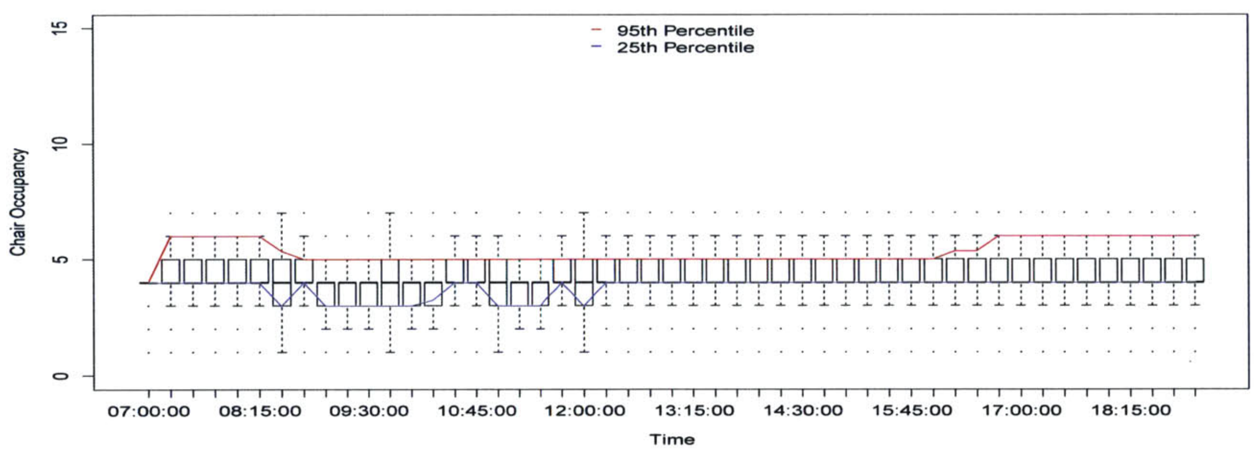

Figure C-21: Scenario V-L-1A Modeled Chair Requirements

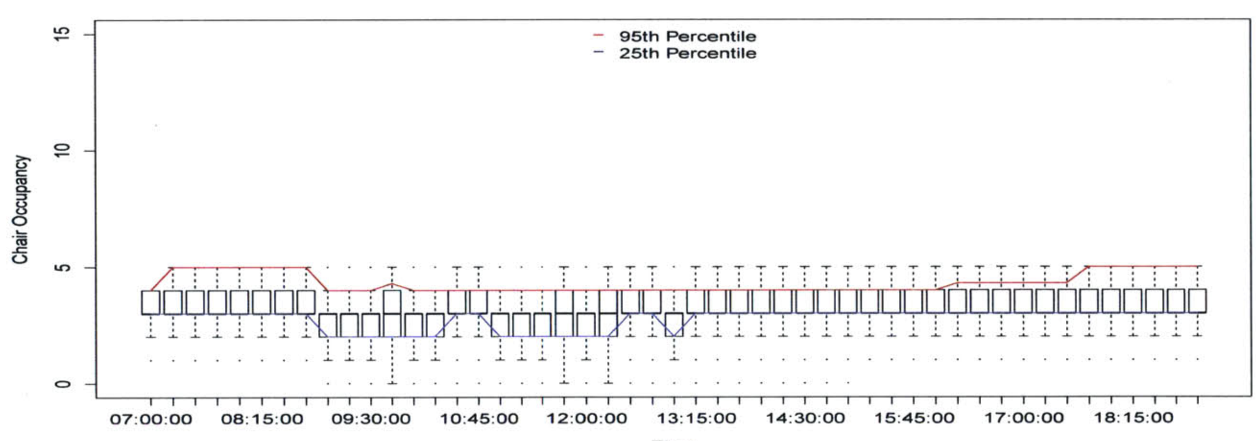

Figure C-22: Scenario V-L-1B Modeled Chair Requirements

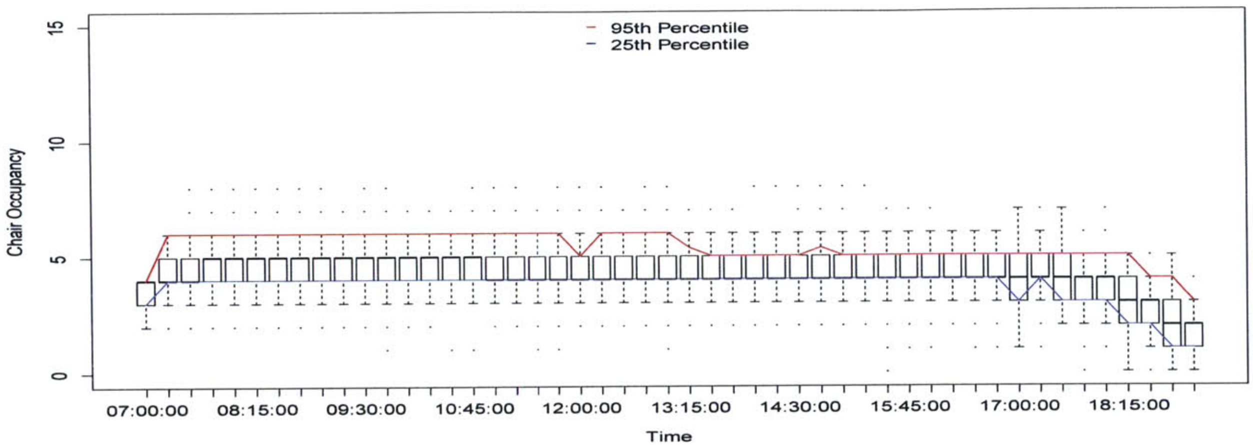

Figure C-23: Scenario V-R-1A Modeled Chair Requirements

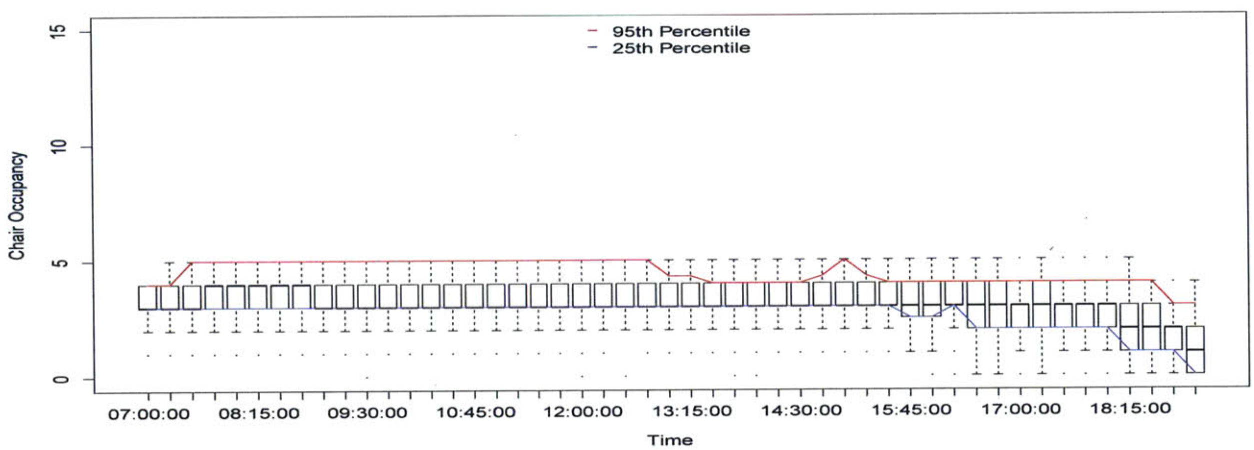

Figure C-24: Scenario V-R-1B Modeled Chair Requirements

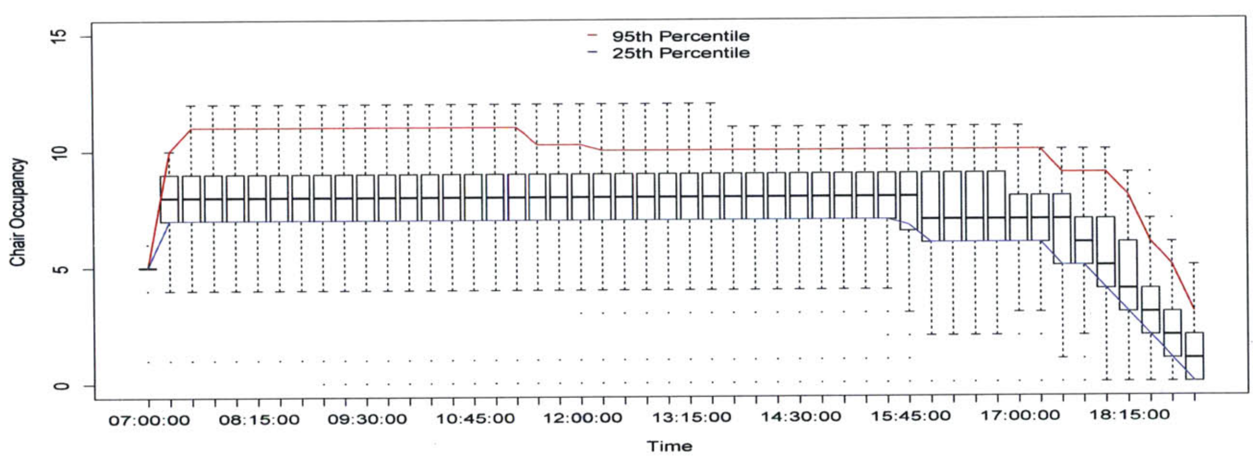

Figure C-25: Sensitivity Scenario PF0-1S Modeled Chair Requirements

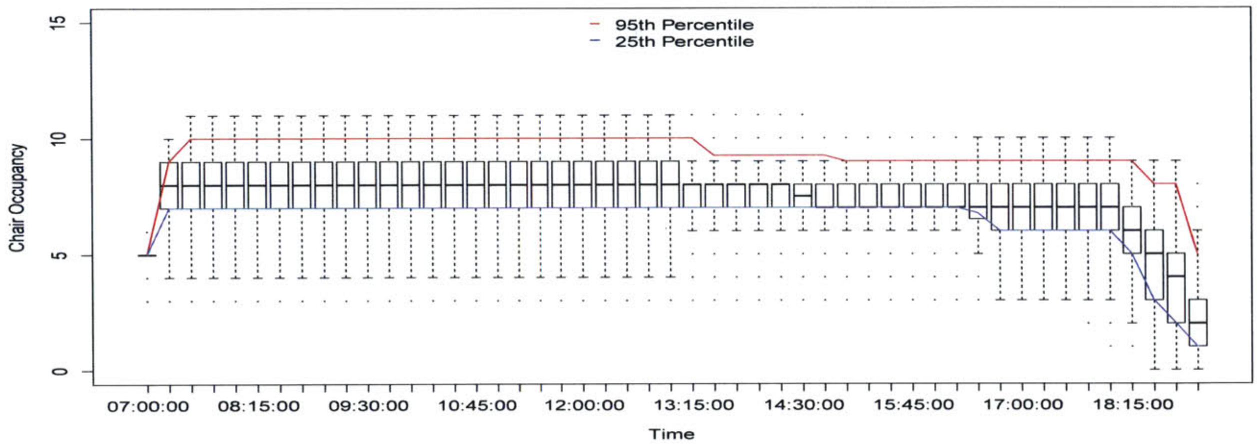

Figure C-26: Sensitivity Scenario PF1-1S Modeled Chair Requirements

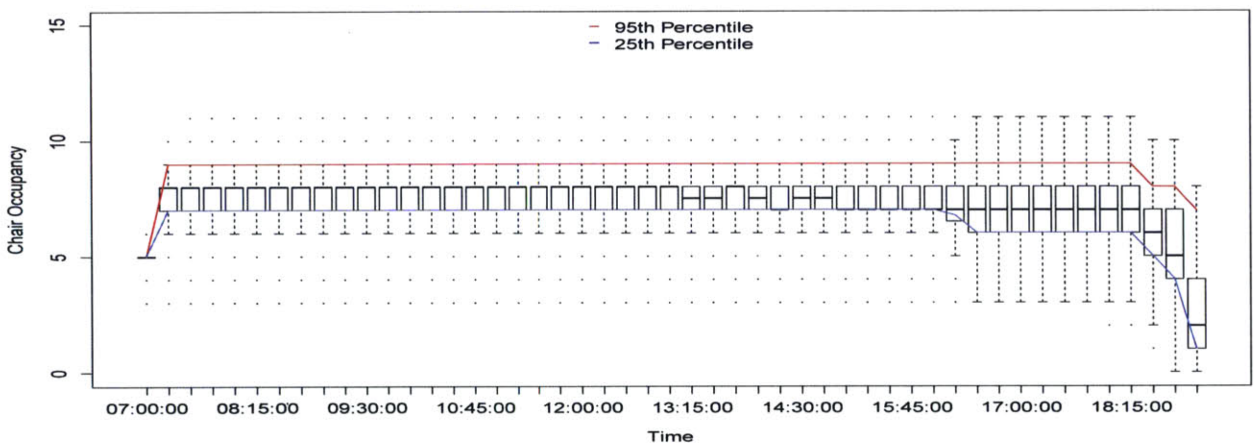

Figure C-27: Sensitivity Scenario PF2-1S Modeled Chair Requirements

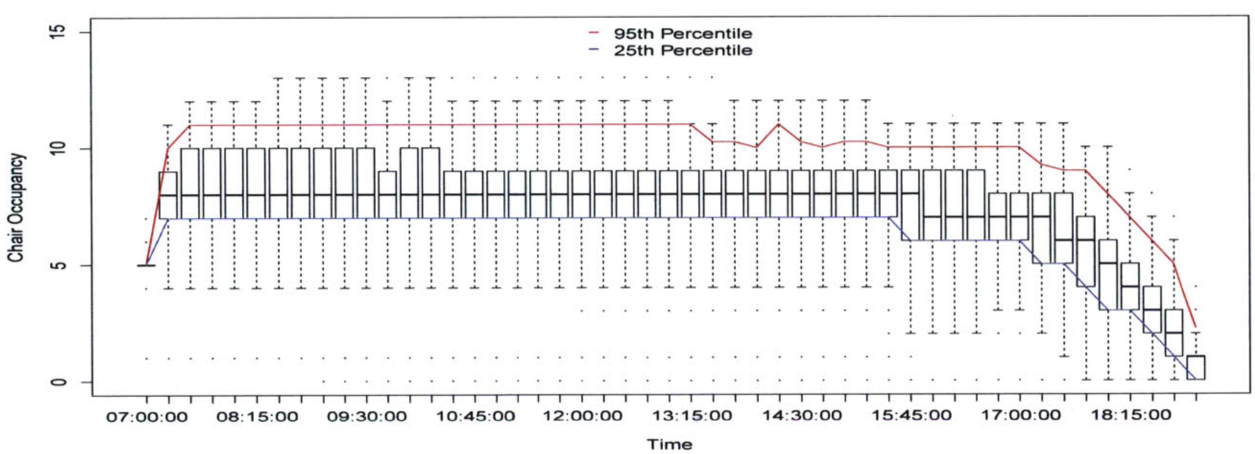

Figure C-28: Sensitivity Scenario PF0-3S Modeled Chair Requirements

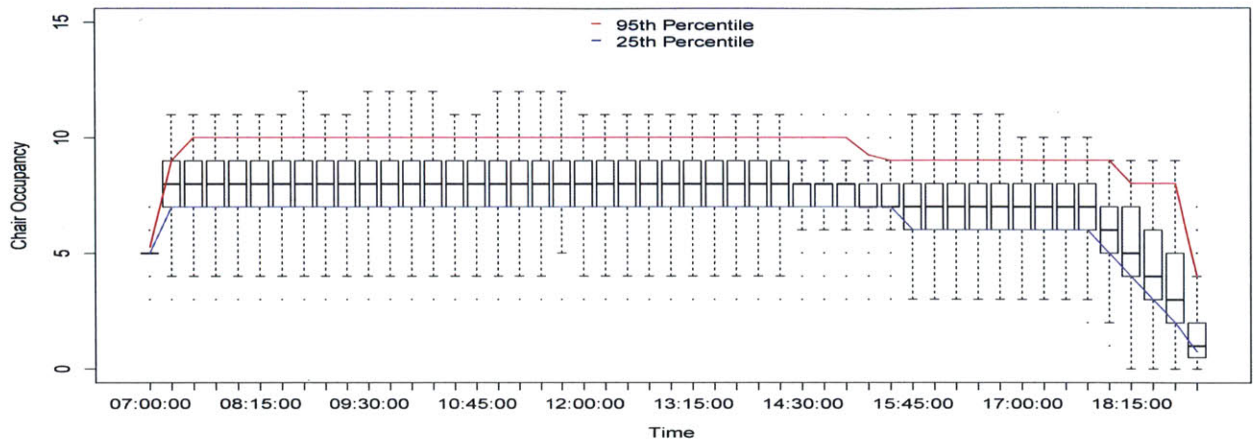

Figure C-29: Sensitivity Scenario PF1-3S Modeled Chair Requirements

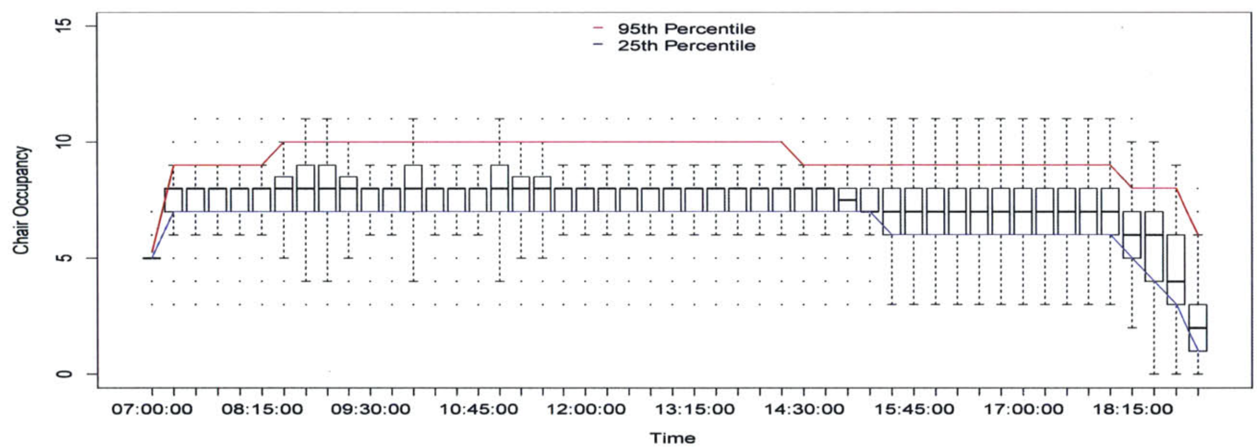

Figure C-30: Sensitivity Scenario PF2-3S Modeled Chair Requirements

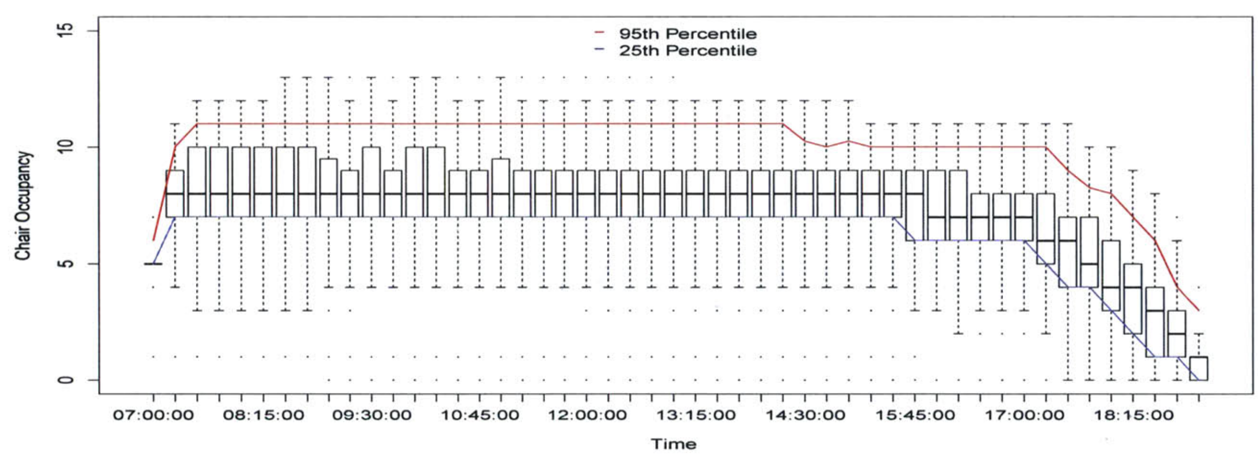

Figure C-31: Sensitivity Scenario PF0-5S Modeled Chair Requirements

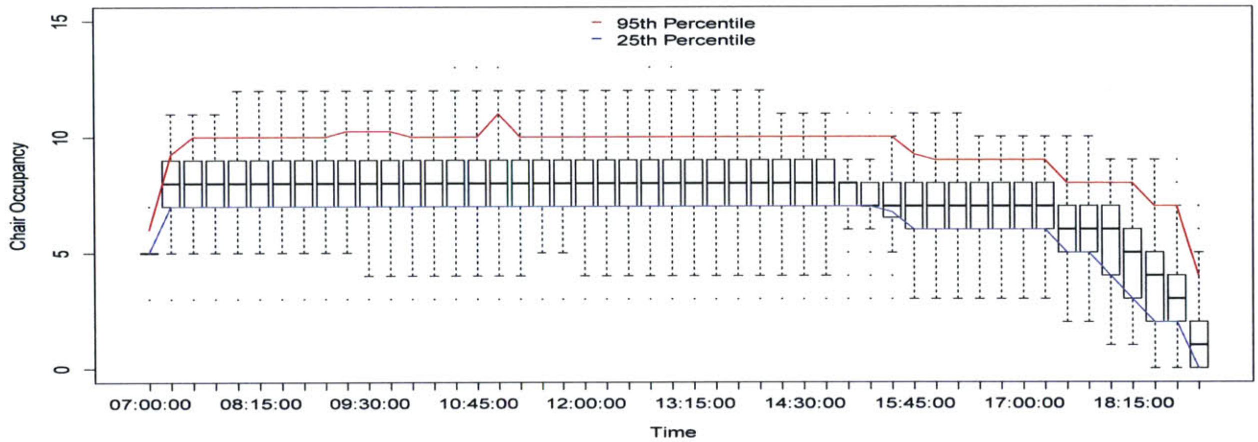

Figure C-32: Sensitivity Scenario PF1-5S Modeled Chair Requirements

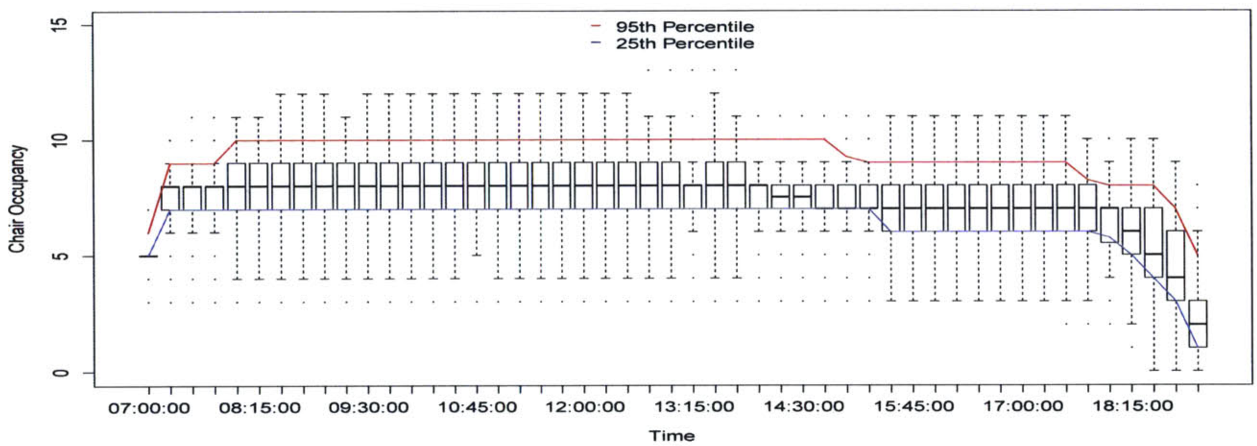

Figure C-33: Sensitivity Scenario PF2-5S Modeled Chair Requirements

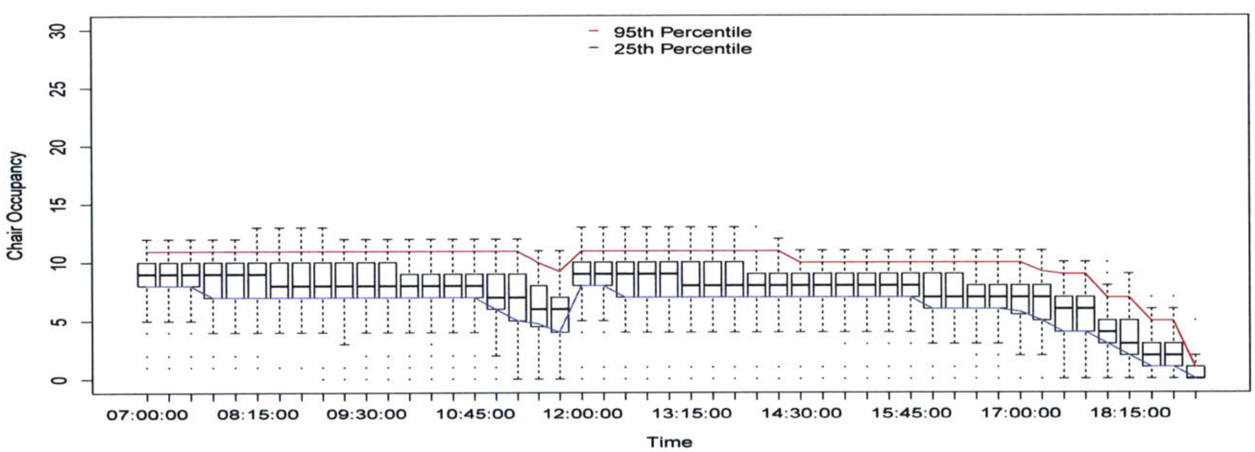

Figure C-34: Sensitivity Scenario PF0-AMPM Modeled Chair Requirements

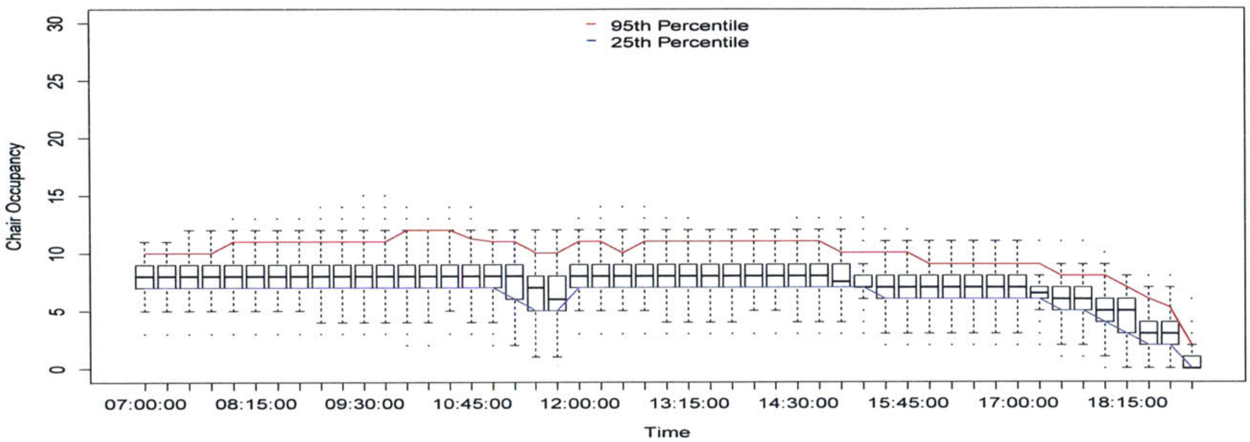

Figure C-35: Sensitivity Scenario PF1-AMPM Modeled Chair Requirements

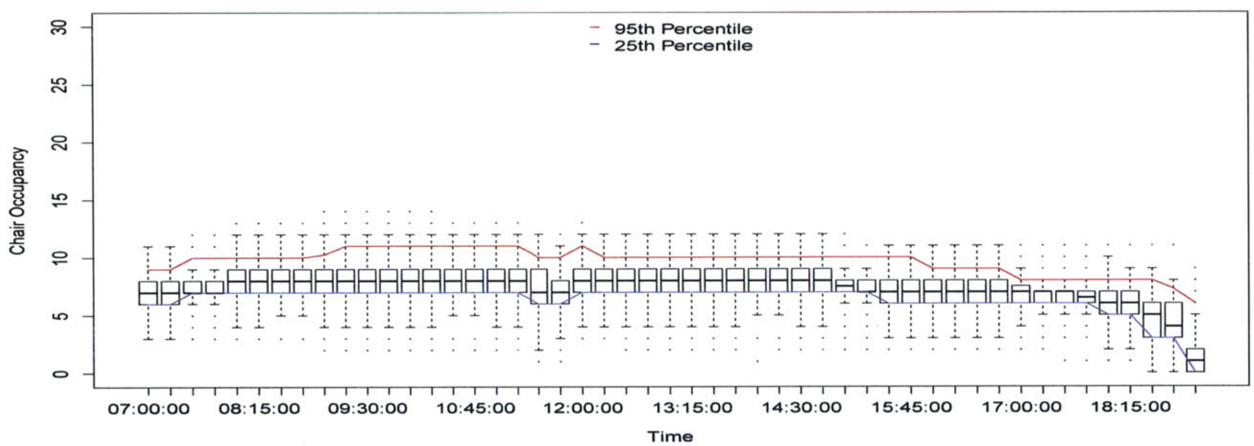

Figure C-36: Sensitivity Scenario PF2-AMPM Modeled Chair Requirements

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix D

Impact of Start Stagger Rule

As described in Chapter 5, when scheduling appointments the algorithm constrains the number of scheduled patient arrivals per time period t, $a(t)$, to be one fewer than the number of nurses on staff in the clinic, or:

$$a(t) \leq n(t) - 1 \tag{D.1}$$

The impact of this rule, called "start stagger", on resource requirements is detailed below. The results of scenario P-F-1 (physical, first available, 1 day of flexibility) are reproduced in Figure D-1. The number of chairs required in this scenario, which has the start stagger implemented, is ten, the standard deviation of the census is 1.54, and the coefficient of variation of census is 0.21.

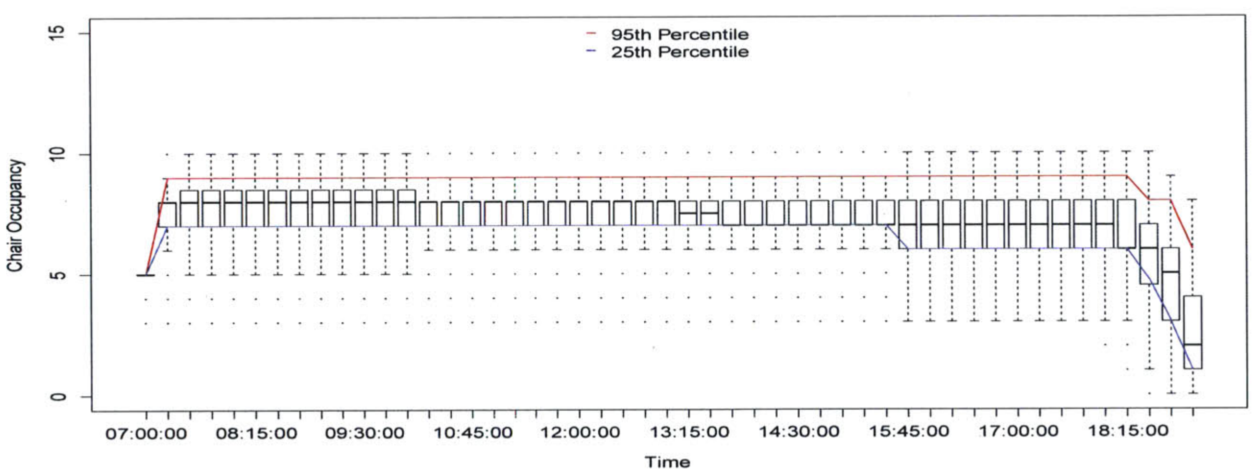

Figure D-1: Stagger Rule Impact, Baseline PF1 Scenario

Figure D-2 shows the same scenario, but with the start stagger rule completely removed. The chair requirement in this case is still ten, and the other metrics fluctuate only slightly (standard deviation decreases to 1.52, coefficient of variation is equivalent). It is apparent that the start of day "ramp" observed in Figure D-1 does not exist, due to the removal of the stagger rule. Inspection of Figure D-3, however, reveals the issues with removal of start stagger. The relaxation of this constraint leads to a maximum of ten patients arriving during one time interval, 7:00am. While the proposed clinic has enough physical space to accommodate these appointments, an additional four nurses would be required at 7:00am if the rule is removed. Thus, the rule creates a more realistic staffing scenario, and has essentially no impact on the chair requirements.

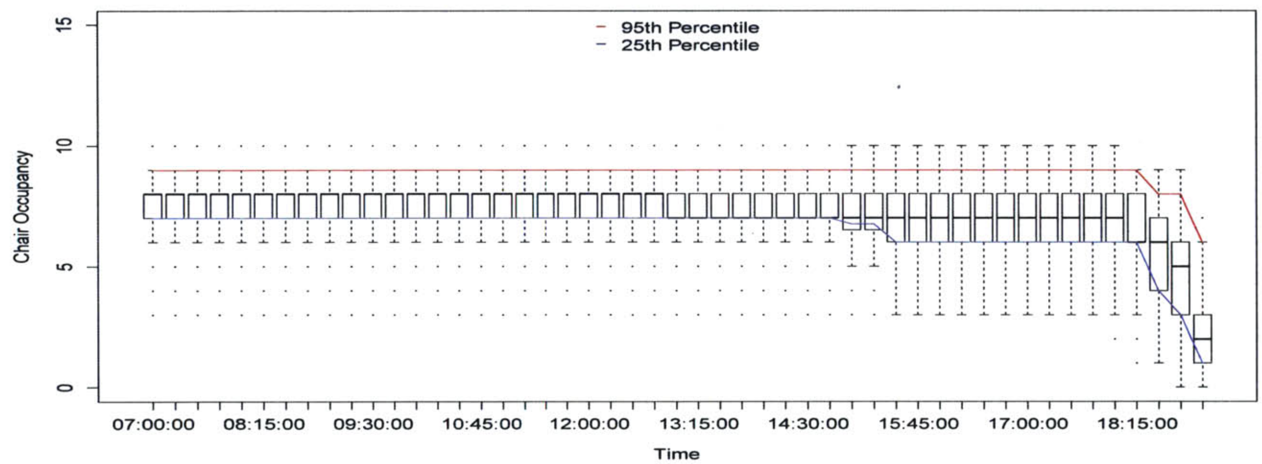

Figure D-2: Stagger Rule Impact, Chair Requirements with Removal

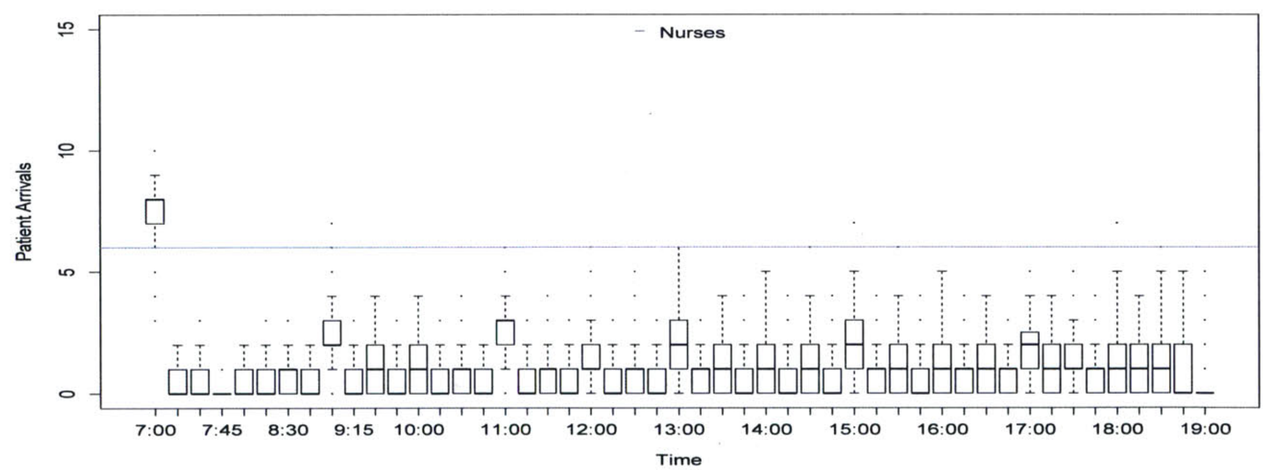

Figure D-3: Stagger Rule Impact, Appointment Starts with Removal

Appendix E

Sample Infusion Clinic Interview Questions

1. Patient access

    - How long must patients wait for an appointment (by treatment)?
    - What are the drivers for this wait time?
    - Are any alternative treatments (example injection) not in use today due to lack of access? For example, if there was extra capacity would there be demand for patients to receive injection at a multi specialty infusion center?

2. Supply and Demand

    - What treatments are administered and by whom?
    - What is the demand for each treatment (patient volume)?
    - What are the considerations for each treatment type (example: appointment length)?
    - How many chairs or stations are available for infusion?
    - Are patients referred to other clinics for infusion?
    - Hours of operation?
    - What operational constraints exist?

- What operational challenges exist?
- What concerns, if any, exist for centralization?

3. Scheduling and prior authorizations
   - How are prior authorizations processed?
   - Are appointments scheduled before or after obtaining authorization?
   - What are the processing times for prior authorizations?

4. Resources and staffing
   - How many nurses perform infusions?
   - How many administrators perform scheduling and prior authorizations?

5. Finances and inpatient demand
   - What financial data is available describing the profitability of the infusion clinic?
   - What is the number of inpatients admitted per year primarily for infusion purposes?
   - Could any inpatient admissions be avoided if there was more capacity for infusions (through centralization)? If so, how many?
   - Are there any trends in infusion demand either by new drugs (e.g., injections; personalized drugs) or policies imposed by insurers?

Bibliography

[1] *Hospital Overview.* URL: http://www.massgeneral.org/about/overview.aspx (visited on 02/18/2016).

[2] Trevor A. Schwartz. "Improving surgical patient flow in a congested recovery area". eng. Thesis. Massachusetts Institute of Technology, 2012.

[3] Ashleigh Royalty Range. "Improving surgical patient flow through simulation of scheduling heuristics". eng. Thesis. Massachusetts Institute of Technology, 2013.

[4] Jonas Hiltrop. "Modeling neuroscience patient flow and inpatient bed management". eng. Thesis. Massachusetts Institute of Technology, 2014.

[5] Sean T. McNichols. "Reducing intraday patient wait times through just-in-time bed assignment". eng. Thesis. Massachusetts Institute of Technology, 2015.

[6] Thomas Daniel Sanderson. "Pooling and segmentation to improve primary care prescription management". eng. Thesis. Massachusetts Institute of Technology, 2014.

[7] Wendi Rieb. "Increasing patient throughput in the MGH Cancer Center Infusion Unit". eng. Thesis. Massachusetts Institute of Technology, 2015.

[8] C. Daniel Mullins et al. "Variability And Growth In Spending For Outpatient Specialty Pharmaceuticals". en. In: *Health Affairs* 24.4 (July 2005), pp. 1117–1127. ISSN: 0278-2715, 1544-5208. DOI: 10.1377/hlthaff.24.4.1117.

[9] Alan M. Lotvin et al. "Specialty Medications: Traditional And Novel Tools Can Address Rising Spending On These Costly Drugs". en. In: *Health Affairs* 33.10 (Oct. 2014), pp. 1736–1744. ISSN: 0278-2715, 1544-5208.

[10] Bijal Nitin Patel and Patricia R. Audet. "A review of approaches for the management of specialty pharmaceuticals in the United States". eng. In: *PharmacoEconomics* 32.11 (Nov. 2014), pp. 1105–1114. ISSN: 1179-2027. DOI: 10.1007/s40273-014-0196-0.

[11] John F. Foley and Anne M. Dunne. "Successful Management of a Neurology Infusion Practice". In: *International Journal of MS Care* 13.2 (July 2011), pp. 95–104. ISSN: 1537-2073.

[12] Sanjar K. Rao. "Trends in market access for specialty biologics: Challenges & promises". In: *Journal of Commercial Biotechnology* 21.2 (Apr. 2015), pp. 10–19. ISSN: 14628732. DOI: 10.5912/jcb669.

[13] Barbara E Ostrov, Kristine Reynolds, and Lisabeth V Scalzi. "Patient preferences and satisfaction in a multispecialty infusion center". In: *Patient preference and adherence* 8 (May 2014), pp. 755–761. ISSN: 1177-889X. DOI: 10.2147/PPA.S63214.

[14] Facmpe Edward N. Freeman. "Addressing Costs and Continuity of Care through Innovative Solutions for Infused Therapies: A Collaborative Experience with Infliximab". In: (Feb. 2011).

[15] Ken Terry. "Electronic prior authorization. The solution to physicians' headaches?" In: *Medical Economics* 92.1 (Jan. 2015), p. 26. ISSN: 0025-7206.

[16] Howard Bell. "The prior authorization burden. The process is frustrating, time-consuming and costly". In: *Minnesota Medicine* 97.11-12 (Dec. 2014), pp. 18–25. ISSN: 0026-556X.

[17] Bendix J. "The prior authorization predicament." eng. In: *Medical economics* 91.13 (July 2014), pp. 29–30, 32, 34–5. ISSN: 0025-7206.

[18] Robert Leinss Jr, Todd Karpinski, and Binita Patel. "Implementation of a comprehensive medication prior-authorization service". In: *American Journal Of Health-System Pharmacy: AJHP: Official Journal Of The American Society Of Health-System Pharmacists* 72.2 (Jan. 2015), pp. 159–163. ISSN: 1535-2900. DOI: 10.2146/ajhp130786.

[19] Sungchul Park et al. "Decentralization and centralization of healthcare resources: Investigating the associations of hospital competition and number of cardiologists per hospital with mortality and resource utilization in Japan". In: *Health policy* 113 (Nov. 2013), pp. 100–109. ISSN: 0168-8510. DOI: 10.1016/j.healthpol.2013.06.005.

[20] Jamie Young et al. "Centralization of services and reduction of adverse events in pancreatic cancer surgery". In: *World Journal Of Surgery* 37.9 (Sept. 2013), pp. 2229–2233. ISSN: 1432-2323. DOI: 10.1007/s00268-013-2108-4.

[21] Christopher T. Aquina et al. "Colon and Rectal Surgery: High volume improves outcomes: The argument for centralization of rectal cancer surgery". In: *Surgery* 159 (Mar. 2016), pp. 736–748. ISSN: 0039-6060. DOI: 10.1016/j.surg.2015.09.021.

[22] Karyn B Stitzenberg and Neal J Meropol. "Trends in centralization of cancer surgery". In: *Annals Of Surgical Oncology* 17.11 (Nov. 2010), pp. 2824–2831. ISSN: 1534-4681. DOI: 10.1245/s10434-010-1159-0.

[23] W. Earl Sasser. "Match supply and demand in service industries". In: *Harvard Business Review* 54.6 (Dec. 1976), pp. 133–140. ISSN: 00178012.

[24] Wallace J. Hopp and Mark L. Spearman. *Factory physics : foundations of manufacturing management.* Boston : Irwin/McGraw-Hill, c2001., 2001. ISBN: 0-256-24795-1.

[25] Peter T. Vanberkel et al. "Efficiency evaluation for pooling resources in health care". en. In: *OR Spectrum* 34.2 (Sept. 2010), pp. 371–390. ISSN: 0171-6468, 1436-6304. DOI: 10.1007/s00291-010-0228-x.

[26] Diwakar Gupta and Brian Denton. "Appointment scheduling in health care: Challenges and opportunities". In: *IIE Transactions* 40.9 (July 2008), pp. 800–819. ISSN: 0740-817X. DOI: 10.1080/07408170802165880.

[27] Yasin Gocgun and Martin Puterman. "Dynamic scheduling with due dates and time windows: an application to chemotherapy patient appointment booking". In: *Health Care Management Science* 17.1 (Mar. 2014), pp. 60–76. ISSN: 13869620. DOI: 10.1007/s10729-013-9253-z.

[28] Tugba Cayirli and Emre Veral. "Outpatient Scheduling in Health Care: A Review of Literature". en. In: *Production and Operations Management* 12.4 (Dec. 2003), pp. 519–549. ISSN: 1937-5956. DOI: 10.1111/j.1937-5956.2003.tb00218.x.

[29] ROBERT D. DIETRICH and SIDNEY J. YAKOWITZ. "A rule-based approach to the trim-loss problem". In: *International Journal of Production Research* 29.2 (Feb. 1991), pp. 401–415. ISSN: 0020-7543. DOI: 10.1080/00207549108930079.

[30] Marc P. Renault, Adi Rosén, and Rob van Stee. "Online algorithms with advice for bin packing and scheduling problems". In: *Theoretical Computer Science* 600 (Oct. 2015), pp. 155–170. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2015.07.050.

[31] Pedro Castro and Ignacio Grossmann. "From Time Representation in Scheduling to the Solution of Strip Packing Problems". In: *Computers and Chemical Engineering* (Nov. 2011), pp. 45–57. DOI: 10.1016/j.compchemeng.2012.05.002.

[32] J.l. Hurink and J.j. Paulus. "Improved online algorithms for parallel job scheduling and strip packing". In: *Theoretical Computer Science* 412.7 (Feb. 2011), pp. 583–593. DOI: 10.1016/j.tcs.2009.05.033.

[33] Marc P. Renault, Adi Rosén, and Rob van Stee. "Online algorithms with advice for bin packing and scheduling problems". In: *Theoretical Computer Science* 600 (Oct. 2015), pp. 155–170. ISSN: 0304-3975. DOI: 10.1016/j.tcs.2015.07.050.

[34] Takahiro Tanaka. *[Infusion Chair Scheduling Algorithms Based on Bin-Packing Heuristics .]* English. Proquest, Umi Dissertation Publishing, July 2012.

[35] Bharathwaj Vijayakumar et al. "A Dual Bin-Packing Approach to Scheduling Surgical Cases at a Publically-Funded Hospital". In: *European Journal of Operational Research* (Feb. 2013), pp. 583–591. DOI: 10.1016/j.ejor.2012.09.010.

[36] *Medical Infusion Center Planning. Update to the MGH/MIT Leadership Council.* Internal MGH Presentation. Feb. 2016.