Patient Flow Optimization in the Department of Medicine at MGH

by

Elizabeth Ugarph

B.S. Industrial Engineering, University of Wisconsin-Madison, 2008

Submitted to the MIT Sloan School of Management and the Institute for Data, Systems, and Society in partial fulfillment of the requirements for the degrees of

Master of Business Administration

and

Master of Science in Engineering Systems

in conjunction with the Leaders for Global Operations Program at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

Signature redacted

Author ....
MIT Sloan School of Management and the Institute for Data, Systems, and Society
May 12, 2017

Signature redacted

Certified by ....
Retsef Levi, Thesis Supervisor
J. Spencer Standish (1945) Professor of Operations Management

Signature redacted

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

Signature redacted

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

Signature redacted

Approved by ....
John N. Tsitsiklis
Clarence J. Lebel Professor of Electrical Engineering, IDSS Graduate Officer

THIS PAGE INTENTIONALLY LEFT BLANK

Patient Flow Optimization in the Department of Medicine at MGH

by

Elizabeth Ugarph

Submitted to the MIT Sloan School of Management and the Institute for Data, Systems, and Society on May 12, 2017, in partial fulfillment of the requirements for the degrees of
Master of Business Administration
and
Master of Science in Engineering Systems

Abstract

In 2015, there were approximately 17,000 General Medicine admissions in the Department of Medicine (DOM) at MGH. General Medicine patients regularly experience significant non-clinical delays caused by bed and care team unavailability, with approximately 25% of patients waiting ten hours or more for a bed. Delays in bed and care team assignments result in decreased patient satisfaction, congestion in the ED and ICUs, and increased overall hospital length-of-stay.

This project studies General Medicine patient flow, develops and evaluates interventions to improve this flow, and provides recommendations to hospital leadership. To this end, we construct a discrete-event simulation based on historical data. Intervention effectiveness is measured primarily based on patient-wait-for-bed, the time from when patient is medically ready for an inpatient bed until the bed is assigned to and ready for the patient.

We find that the simulation model accurately represents the wait times of General Medicine patients. We propose a new algorithm, which when implemented could reduce overall average patient-wait-for-bed by 9% from 7.36 to 6.67 hours. Implementation of additional capacity and reorganization of the physician care teams (known as the DOM redesign) is shown to result in a further 31% reduction in average wait time (from 6.67 to 4.59 hours). Other interventions tested such as early assignment of patients to care teams based on predicted discharges, and increased flexibility of care teams to cover different units are shown to have modest effects on overall patient-wait-for-bed.

Thesis Supervisor: Retsef Levi
J. Spencer Standish (1945) Professor of Operations Management

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

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgments

There are many individuals I would like to thank for the time and effort they dedicated in supporting my work on this project.

First, I would like to thank Professor Retsef Levi for his guidance and inspiring devotion to improving the quality of healthcare services delivery. I would also like to thank Professor David Simchi-Levi for his feedback and advice.

I am indebted to numerous individuals at Massachusetts General Hospital. Thank you to Dr. Peter Dunn, Bethany Daily, and Cecilia Zenteno for their support of the MGH-MIT collaboration, and this project in particular. Thank you to all of the members of the Capacity Task Force and the medicine sub-group. In particular, I would like to thank Rhodes Berube, Dr. Walter O'Donnell, and Dr. Kimi Kobayashi for their leadership of the team and their support of this project. Thank you to Gianna Wilkins, Allison Castagna, and Ed Morris for their help in facilitating and guiding this project. I thank Dr. Marjory Bravard, Dr. Chana Sacks, Dr. Melissa Mattison, and Lee Ann Tata R.N. for their help in arranging shadowing and generally improving my understanding of the Department of Medicine. I thank Ben Orcutt and Katie Turcotte from Admitting for their valuable feedback and for answering many questions about processes and data. A huge thank you to Stephen Buonopane and Kevin Murphy for their dedicated help in obtaining the data required to complete this project.

Thank you to Aleida Braaksma, a post-doctoral fellow in the Operations Management group at the MIT Sloan School of Management, who was an instrumental partner throughout this project. This thesis truly would not have been possible without her tireless support and technical guidance. Additionally, I thank Martin Copenhaver, a Ph.D. candidate at the MIT Operations Research Center (ORC), for his helpful tutoring that certainly accelerated my progress on this project and Andrew Vanden Berg, a Master's student at the ORC, for his help with data analysis.

To all of my LGO classmates and fellow MGH interns, thanks for your friendship and collaboration. And finally, thank you to my husband Matt and the rest of my family for their support and encouragement on this thesis and throughout my LGO journey.

THIS PAGE INTENTIONALLY LEFT BLANK

The author wishes to acknowledge the Leaders for Global Operations Program for its support of this work.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

List of Figures

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Background

This project aims to improve the flow of patients cared for by the Department of Medicine at Massachusetts General Hospital (MGH). The project has been conducted by the MGH-MIT collaboration under the guidance of the MGH Capacity Task Force. Background information on these organizations is provided in the following sub-sections.

1.1.1 Massachusetts General Hospital

MGH is the third oldest hospital in the United States and the oldest and largest hospital in New England. It is consistently ranked among the top hospitals on the U.S. News & World Report Best Hospitals Honor Roll, placing in the top three in every year since 2010. In 2016, US News & World Report ranked MGH third in the nation and first in New England based on its quality of care, patient safety, and reputation in 16 clinical specialties [1]. MGH is also a designated Magnet hospital, a recognition of nursing excellence [2].

In addition to excellence in patient care, MGH is the original and largest teaching hospital of Harvard Medical School and is known for the quality of its teaching program. The hospital admits approximately 48,000 inpatients per year to 999 inpatient beds, handles more than 1.5 million outpatient visits, records over 100,000 emergency room visits, and performs more than 42,000 surgeries [2].

MGH is a founding member of Partners Healthcare, which is an integrated healthcare system that offers patients a continuum of coordinated, high-quality care. The system includes primary care and specialty physicians, community hospitals, the two founding academic medical centers (MGH and Brigham and Womens), specialty facilities, community health centers, and other health-related entities [3].

1.1.2 The MGH-MIT Collaboration

The MGH-MIT Collaboration is a long-standing research partnership between MGH and the MIT Sloan School of Management. The Collaboration focuses on improving the operational effectiveness of the hospital through the application of Operations Research and continuous improvement methodologies. These efforts are driven by a team composed of MIT faculty, MGH leadership, post-doctoral fellows within the Operations Management Group at the MIT Sloan School of Management, and students within the MIT Leaders for Global Operations (LGO) program and Operations Research Center.

While the Collaboration initially focused on driving improvements within the hospital's Perioperative environment, including surgical scheduling, inpatient flow optimization, and surgical inventory management, it has since expanded in scope to include other departments and functions. Recent efforts include quantifying the effect that delayed transfers out of the ICU have on hospital LOS, evaluating Just-In-Time bed assignment in the Neurosciences, and analyzing the impact of discontinuities in care in the DOM.

This thesis and the research that underlies it was completed within the framework of an IRB-approved study[1]. This project builds upon previous work by proposing interventions that support and extend established best practices and applying them in a new environment. The new contributions of this thesis include the introduction and evaluation of early team assignment and the first use of discrete event simulation by the MGH-MIT Collaboration to analyze interventions in an environment containing a majority of semi-private rooms.

1.1.3 The MGH Capacity Task Force

By 2015, MGH leadership recognized that significant capacity challenges were impacting both staff and patients. This impact was most visible in the hospital's Emergency Depart-

[1]MIT Protocol #12010014856, MGH-MIT Collaboration: DOM Inpatient Flow, Principal Investigator: Retsef Levi. MGH Protocol #2011P001124.

ment (ED), where patients who were waiting for inpatient beds created physical congestion and consumed clinician time and attention, potentially detracting from their ability to care for newly arrived patients. On March 1, 2016 hospital leadership presented the following statistics to characterize the magnitude of the challenge and some of its causes [4]:

- Visits to the MGH Emergency Department increased by 18% in the past five years.
- In FY '15, 80% of all admitted patients boarded in the ED for varying periods of time, up from 74% in FY '14[2].
- Notifications that the hospital was approaching or had reached Code Help[3] or Capacity Disaster status[4] increased by 164% in FY '15 compared with FY '14.
- Inpatient operational occupancy typically hovers between 95% and 100% each day.

The symptoms of the capacity problem were described as "delays, dissatisfaction and sometimes even concerns about quality and safety" [4].

In late 2015, the Capacity Task Force was chartered by the President of the hospital to address these challenges with both immediate interventions and long-term strategic planning. Three areas of focus were chosen for the Task Force's efforts: avoidable ED admissions, preventable re-admissions, and delays related to patient placement and bed allocation. This work supports the third initiative by seeking to minimize patient wait times for beds and physician care, which has the systemic effect of relieving the burden on the ED.

1.1.4 The Department of Medicine

The Department of Medicine (DOM) is the largest department at MGH by inpatient volume, Admitting over 21,000 patients in 2015[5]. The DOM provides outpatient primary care, general

[2]An ED Boarder is a patient who waits in the ED for two hours or more following an inpatient bed request.

[3]Code Help is a state-mandated policy requiring hospitals to move all admitted patients out of the ED within a 30-minute period after the ED's maximum occupancy is reached or exceeded. Maximum occupancy is influenced by the number of patients present and their acuity/clinical complexity. A prolonged Code Help requires the hospital to make a report to government officials and might require the ED to divert patients to other hospitals.

[4]Capacity Disaster is defined as the state in which Code Help has been activated for two hours or more and there are still ED Boarders present.

[5]Statistic obtained from the Program Manager of Clinical Operations in the Department of Medicine.

inpatient care, intensive, and emergency medical services through a network of patient care units spanning ten clinical divisions[6], 15 primary care locations, 20 inpatient units, and over 400 inpatient beds [5].

This project is focused on inpatient care in the General Internal Medicine division (hereafter referred to as 'General Medicine'). The DOM, and General Medicine in particular, care for an inpatient population that is clinically heterogeneous, with many patients arriving to the hospital with complicated, multi-system ailments, and psychosocial complexities. The diagnostic uncertainty associated with newly admitted patients and the diversity of clinical needs result in an environment in which patient length-of-stay (LOS) can range from a day or two to several months. This environment can be sharply contrasted with surgical services where many admissions are elective and patients' diagnoses are largely known before admission; thus, paths of treatment are clear and LOS is more predictable [5].

At the time that this project began, General Medicine was allocated 230 beds; with many patients occupying additional 'non-regionalized' beds that are not allocated to General Medicine but to other services. These beds are called non-regionalized because General Medicine patients in them are cared for by physicians that are not local, or regional, to the floor. This is contrasted by regionalized beds that have local physician coverage. The placement of General Medicine patients in non-regionalized beds occurs because the demand for General Medicine beds frequently exceeds the allocated capacity, while other services have some degree of excess capacity on their assigned floors. Not all General Medicine patients are appropriate for placement in non-regionalized beds. Patients that are expected to require a high physician workload due to their acuity or complexity are referred to as 'Level 1' and must be placed in a regionalized setting where their physicians are consistently in close proximity. All other patients are 'Level 2' and may be placed in both regionalized or non-regionalized settings.

In an effort to mitigate the placement of General Medicine patients on non-regionalized units and improve patient flow, the DOM implemented a redesign in June 2016, which among

[6]The ten clinical divisions are: Cardiology, Endocrine, Gastroenterology, General Internal Medicine, Hematology & Oncology, Infectious Diseases, Nephrology, Palliative Care, Pulmonary and Critical Care Medicine, and Rheumatology, Allergy & Immunology.

other changes, increased the General Medicine allocated capacity to 240 beds (the redesign is discussed in detail in Section 3.3). Of the 230 beds allocated to General Medicine at the start of this project, approximately 84% were semi-private meaning that there are two beds that share a room and a bathroom and are separated by a curtain (the remaining 16% were private). Patients placed in semi-private beds must be compatible with their roommate in terms of sex and infection precautions, which imposes non-trivial constraints on the bed assignment process (details on bed assignment requirements can be found in Section 3.2.1).

General Medicine patients are cared for by two different types of physician teams. House staff teams are composed of resident physicians in training and their supervising Attending physicians. These teams only care for patients on regionalized units, and prior to the redesign, their capacity was sometimes less than the physical bed capacity of their assigned units (see Section 3.1.1 for details). The other teams belong to the Hospital Medicine Group (HMG) and are made up of only post-residency physicians. These teams cover some regionalized patients on their assigned units, but can also care for non-regionalized patients (see Section 3.1.2 for details).

1.2 Project Overview

1.2.1 Problem Statement

In 2015, there were nearly 17,000[7] patients admitted to the DOM at MGH. The majority of these patients (approximately 80%) are admitted through the MGH ED, an additional 9% come from other locations within MGH (ICU or other units), and the remaining 11% come from other medical institutions or systems (hospital transfers or outpatient referrals).

The DOM's General Medicine units have high occupancy rates, leading to patients regularly experiencing significant non-clinical delays. In 2015, approximately 25% of General Medicine patients waited ten hours or more for a bed[8]. In this highly-utilized system, finding a bed for a newly arrived patient is often contingent on another patient being discharged

[7]This number represents patient encounters, not distinct patients. A patient will be counted twice if they visit the hospital twice or if during one visit the patient moves from a General Medicine area to another area (such as an ICU) and back. Source: ADT and PEPL databases following procedure described in Section 4.1.2.

[8]This statistic is based on patient-wait-for-bed as defined in detail in Section 3.4.1.

from the hospital. At MGH there is a significant mismatch between the distribution of patient arrival times and that of discharge times. While the arrival rate is relatively steady throughout the day, discharges are concentrated in the afternoon and evening hours (see Section 3.2 for a detailed description of this effect). This mismatch in timing combined with the high occupancy rate results in long intraday wait times for many General Medicine patients, particularly those that arrive early in the day.

Extended wait times for beds may potentially have negative impacts on patients' experiences and clinical outcomes. Moreover, patients waiting in the ED can disrupt hospital operations. For example, if the ED is overly crowded, the hospital may be mandated to go into Code Help and cancel elective surgeries and other activities until the crowding in the ED is reduced. Activating Code Help can cause the hospital to prioritize moving patients out of the ED by placing patients on units other than those assigned to their service, and by delaying bed assignments for patients from other areas of the hospital (including the Post-Anesthesia Care Unit or PACU). The consequences of Code Help require significant management attention and can affect hospital operations for several days.

Currently, discharges and bed assignments for new admissions at MGH are managed with manual and relatively static processes, and without clear prioritization rules. This system, combined with the capacity constraints, cause long intraday wait times and interday delays for patients and may result in suboptimal bed assignments and reduced throughput.

Observations at the hospital and interviews with relevant stakeholders, backed by data analysis, revealed the following challenges that affect General Medicine along with many other clinical specialties and departments at the hospital:

- The bed assignment process is highly decentralized and not standardized. Patient placements are negotiated daily on a case-by-case basis between the Admitting department and clinical staff in the inpatient units. While Admitting is responsible for matching newly admitted patients to beds, the beds in the inpatient units are managed by the respective nursing staff.

- Bed-patient assignments are made without access to all relevant information, such as timing of discharges. No transparent guidelines exist on how to prioritize assignments.

As a result, assignments are often made prematurely (i.e., before the bed is truly available and/or the patient is ready to occupy the bed). Generally, beds are not reassigned adaptively when the bed or patient is not ready as expected, even when reassignment could reduce both patient wait times and bed idle times.

- Discharges generally do not occur until late in the day. Care providers prioritize teaching activities and inpatient care for newly admitted patients and patients likely to stay in the hospital throughout the morning over discharging patients. This contributes to a misalignment in the intraday timing of admissions and discharges.

- The majority of General Medicine beds are semi-private (two beds in one room), and therefore it can be very difficult to place patients that have infection precautions or behavioral issues that require them to be in a room alone. While some patients with similar infection precautions can be placed together in room (known as cohorting), many need to be alone. Patients who need to be alone must wait for a private room to become available or for a whole semi-private room to become available (i.e., there is no patient assigned to either bed). If the patient is placed in a semi-private room, the other bed will be closed. The staff on the unit works with Admitting to adjust the placement of patients within their unit to accommodate the maximum number of patients, a process referred to throughout this document as 'bed swaps'.

- General Medicine beds are divided between those that are regionalized and non-regionalized, based both on their physical unit location and the availability of a regionalized physician team to cover the bed. When a regionalized bed becomes available and there is not currently an appropriate Level 1 patient waiting, Admitting must determine whether the bed should be assigned to a waiting Level 2 patient or reserved for a Level 1 patient expected to arrive at a later time.

1.2.2 Goals

This project aims to evaluate and refine the redesign concept developed by the DOM and to propose and evaluate additional interventions to reduce patients' wait for beds and care teams. The additional interventions include the early assignment of ED patients to physician teams, based not on a particular bed being available, but instead on the high probability (based on historical discharge data) that a bed will become available for the patient in the appropriate unit before the end of the day. This allows the patient's inpatient care to begin before he reaches his bed and reduces workload on the ED physicians. Other interventions seek to optimize bed management through prioritization rules and adjustments to the scope of patients considered for bed swaps.

1.2.3 Approach

A discrete event simulation has been developed to assess these interventions on key performance criteria. The simulation and input data have been prepared based on careful study of the hospital's processes, including shadowing and process mapping. By adjusting features of the simulation and the input, the effects of the different interventions can be estimated, alone or in combination, without the disruption to operations associated with implementation. This way, the most impactful interventions can be selected for implementation and the expected benefits can be weighed against costs. This simulation has been developed with the intent of being generalizable to other areas of the hospital and for other interventions, giving future researchers the ability to adapt it as necessary. The specific interventions tested for this project are briefly described below and discussed in detail in Section 4.3:

- Patient assignment algorithms: Prioritize patients to available beds based on clear rules that are consistent with the current prioritization of the Admitting department. As previously stated, actual practices regarding patient prioritization are not standardized or codified. This intervention was developed to provide a standard practice for assignment in the other interventions and to demonstrate the opportunities for improvement over the current methodology. This algorithm is discussed in detail in Section 4.3.1. A methodology of first-come-first-served prioritization was also explored and is discussed in Section 4.3.4.

- The redesign: The redesign was initially conceived prior to the start of this project in February 2016 by administrators, physicians, and nurses within the DOM and was further developed in parallel with this project. The redesign aims to increase capacity and support teaching goals. Important features include better alignment of physician capacity with bed capacity, creation of pods where physician teams span two units, and the creation of an Oncology sub-unit. The redesign was implemented in June 2016 and is discussed in detail in Section 3.3.

- Early team assignment: Requires the assignment of a number of ED patients to General Medicine physician teams in the morning before they are matched with a particular bed. The General Medicine physician will then assume care for the patient while the patient waits in the ED, advancing the patient's care, removing workload from the ED staff, and incentivizing the medicine team to complete discharges that allow the patient to reach the unit.

- Increased flexibility within pods: Although the pod structure is part of the original redesign, the planned flexibility the pods provide is limited. Each pod consists of two units and three physician teams. In the original design, only the 'Flex' team in each pod covers patients on both floors and patients are not swapped across floors, a practice that in some cases could maximize available capacity by cohorting patients with compatible sex and isolation status that happen to be on different floors. Interventions analyzing increased flexibility relax these assumptions.

- Eliminate staffing closures: General Medicine units occasionally experience bed closures due to insufficient nurse staffing. It is believed that these bed closures could be eliminated through more pooling of resources between floors.

The following performance metrics are used to evaluate and compare the effectiveness of different process interventions:

1. Patient-wait-for-bed: Time from a patient's medical readiness to be assigned to a bed until the time a bed is cleaned and assigned to the patient. This metric does not include additional transfer processing wait times incurred after the patient is ready and a clean bed is assigned. The additional transfer processing time was omitted from this metric because the interventions evaluated are not intended to directly impact this time.

2. Patient-wait-for-team (ED only): Time from a patient's medical readiness to be admitted until a handoff occurs between ED and inpatient medicine physicians. This metric is tracked only for patients that originate in the ED since there is a record of the physician handoff that is not available for other populations. Additionally, for ED patients this time represents workload on ED physicians associated with patients that they have already determined require inpatient care.

By measuring the effects of the interventions on these performance metrics, it is possible to determine which combination of interventions is optimal, and what the impact is expected to be. These metrics also provide a basis to compare the base scenario of the simulation to the historical data and therefore validate the simulation itself.

This project focuses on methods to reduce the wait times experienced by patients and not on quantifying the effects of reducing these wait times. For example, if a patient could arrive on the appropriate unit six hours earlier due to the proposed interventions, he may be seen by the attending physician the first day of his stay rather than waiting until the next morning. Prior work has suggested that this could shorten the patient's hospital LOS [5]. Previous work has also shown that delays in patients leaving ICUs for general care units lead to increased overall hospital LOS [6].

1.3 Results

The simulation of patient flow is validated by comparing the base scenario to historical data. Difference of means tests are used to validate that the mean patient-wait-for-bed and patient-wait-for-team are consistent with history for all subgroups of interest. Additional comparison of the quantiles supports the finding that simulation results are consistent with the historical data[9].

Implementation of the patient assignment algorithm in the pre-redesign system is found to reduce overall average patient-wait-for-bed time by 9% (from 7.36 hours to 6.67 hours), increase the median and 75th percentile patient-wait-for-bed time by 1% (from 2.95 to 2.97 hours) and 7% (from 9.33 to 10.00 hours) respectively, and to reduce the 95th percentile

[9]The differences in historical patient-wait-for-bed minus simulation patient-wait-for-bed were as follows for the 5th, 25th, 50th, 75th, and 95th percentiles respectively: -2, -4, 10, 7, and 0 minutes.

wait time by 12% (from 24.72 to 21.78 hours). This improvement in the overall mean comes at the cost of a 4% increase in median patient-wait-for-bed time for ED patients (2.68 hours to 2.78 hours) and results in large decreases for other populations. ICU patient-wait-for-bed on average drops from 24.09 to 13.59 hours, at the median drops from 10.38 to 5.65 hours, at the 75th percentile drops from 30.51 to 18.05 hours, and at the 95th for ICU patients drops from 84.95 to 56.20 hours. Hospital transfer patients see similar decreases, from 20.17 to 13.57 hours on average, 6.28 to 3.99 at the median, 24.92 to 21.82 at the 75th percentile, and 81.47 to 52.55 at the 95th percentile. Based on the improvement in overall patient wait time and the reduction in excessive waits for these populations, we recommend that the DOM and Admitting adopt such an algorithm for patient placement.

The redesign is shown to further decrease overall patient-wait-for-bed by 31% on average (from 6.67 hours to 4.59 hours), 47% at the median (from 2.97 to 1.58 hours), 38% at the 75th percentile (from 10.00 to 6.20 hours), and 18% at the 95th percentile (from 21.78 to 17.77 hours). Due to the features of the redesign, this effect is particularly pronounced for Level 1 patients who see their average wait drop 40% from 8.09 to 4.86 hours and their median drop 57% from 4.33 to 1.88 hours, and for Oncology patients whose average wait drops 52% from 6.91 to 3.29 hours and 67% at the median from 3.50 to 1.17. These findings confirm that the implementation of the redesign as originally conceived reduces patient wait time, but suggest that there may be an opportunity to refine the capacity and flexibility of the Oncology sub-unit.

We found that the implementation of early team assignment in the simulation could reduce the patient-wait-for-team for those patients selected by 27% on average (from 9.93 to 7.28 hours), 25% at the median (from 9.93 to 7.45 hours), 25% at the 75th percentile (from 14.40 to 10.77 hours), and 25% at the 95th percentile (from 19.31 to 14.41 hours). However, there is a small increase in overall patient-wait-for-bed for the entire population with the implementation of early team assignment (from 4.58 to 4.74 hours on average and from 1.60 to 1.67 at the median). We believe that there are key benefits of the early-team-assignment intervention that are not captured in the simulation, namely the likelihood of accelerated discharges due to the realignment of physician incentives, and that these benefits may be enough to offset the observed increase in overall patient-wait-for-bed. Because of the encouraging results for those patients who are selected for early assignment and the potential for further improvement with early discharges, we recommend that the DOM perform a pilot study of this intervention prior to implementation.

1.4 Thesis Outline

The thesis is organized as follows. Chapter 2 provides a review of relevant studies in the existing literature, including previously published research of the MGH-MIT Collaboration. Chapter 3 provides description and analysis of the DOM at MGH including the house staff (teaching) teams and units (Section 3.1.1) and Hospital Medicine Group teams and units (Section 3.1.2). Section 3.2.1 describes the processes for assigning patients to the beds and physician teams. In Section 3.3 the redesign is described in detail. Sections 3.4.1 and 3.4.2 introduce the performance metrics and evaluate the baseline state of the system on these metrics. An analysis of the current state and a summary of the key challenges is presented in Section 3.5. Chapter 4 describes the model used for the data-driven patient flow simulation. It discusses the data sources used in the simulation (Section 4.1), the modeling framework (Section 4.2), and the model implementation of different process interventions (Section 4.3). Chapter 4 also presents the combinations of interventions that were selected for testing and the rationale behind these selections (Section 4.4). Chapter 5 discusses the results of the base model (Section 5.1) and the intervention models (Sections 5.2-5.4). Chapter 6 makes recommendations for the implementation and prioritization of different interventions based on the simulation results (Section 6.1), and provides ideas for further study (Section 6.2).

Chapter 2

Literature Review

There is a large volume of literature related to the implementation of Operations Research methods in the field of healthcare and a growing recognition of the potential improvements to be realized through the application of these methodologies. In particular, Green states that there has never been a more opportune time for people with analytical skills, particularly in Operations Research, to provide decision-making guidance to improve the healthcare delivery system [7].

As mentioned previously, this thesis does not seek to explore the relationship between delays in bed assignment and patient outcomes, but it is motivated by the belief that such delays are detrimental to patient satisfaction and care. Several publications support this hypothesis, particularly for patients originating through the ED, who make up 80% of the MGH General Medicine population. In a 2007 study of 187 acute hospitals in California, Sun et al. found that admitted patients arriving on days of high ED crowding had 5% greater odds of inpatient death (95% confidence interval [CI] 2% to 8%) [8]. In a 1999 Australian study, patients delayed in the ED (defined as spending at least 8 hours in the ED) were found to have a mean inpatient LOS of 4.9 days (95% CI, 4.7 to 5.1), compared with 4.1 days for non-delayed patients (95% CI, 4.0 to 4.2; $P<0.0001$) [9]. Other researchers have explored the relationship between delayed discharges and ED boarding, and have found that ED boarding and crowding can be significantly alleviated by prioritizing timely discharges [10][11]. This intervention will not be explored directly in this thesis but is believed to be complementary to those that are proposed.

This thesis explores algorithms for assigning patients to beds and physician teams in real-time. The majority of the literature in the field of bed assignment focuses on non-real-time methods, where patient admissions are known at the start of a period and all assignments are performed simultaneously. In particular, Bachouch et al. (2012) investigate the management of hospital bed planning and propose a decision support tool based on an integer linear program [12]. Thomas et al. (2013) use a mixed-integer goal-programming approach to develop a prototype bed-assignment solution, which periodically recommends bed-patient assignments based on analytical decision support tools with embedded mathematical models [13]. Of the studies that propose methods for real-time bed assignment, one focuses on assignment to units only [14], while the other assumes known patient discharge times [15]. In contrast with these studies, this thesis proposes straightforward and easily explained algorithms that do not rely on optimization methods and are evaluated through simulation.

Previous MGH-MIT Collaboration research projects have developed valuable insights into patient flow dynamics and operational processes at MGH that inform the approaches and interventions presented in this thesis. Christensen (2012) studies transfer delays for patients in the surgical intensive care unit (SICU) at MGH using multivariable regression and detailed discrete event simulation. He finds that one of the most promising solutions is transferring patients as soon as possible after medical readiness, eliminating the current practice of de-prioritizing ICU movers [16]. Dolcetti (2015) builds on the work of Christensen by using data from all six of MGH's adult ICUs over a four-year period, providing statistical validation of the model and results, and including additional explanatory variables. She finds that patients spend approximately the same amount of time in the hospital after their ICU stay (i.e., in a general care unit), regardless of whether they were delayed in the ICU [6]. These findings inform our proposed bed assignment algorithms and our treatment of ICU patients in the simulation as described in Section 4.2.

Hiltrop (2014) and McNichols (2015) develop and analyze a simulation that represents MGH's neuroscience units (two general care units and one ICU) [17][18]. Hiltrop finds that a combination of accelerated discharges and Just-In-Time (JIT) bed assignment can significantly decrease average wait times. Hiltrop defines the JIT methodology as a procedure to allocate available inpatient beds adaptively and only to patients who are ready to move to the bed. In practice this means waiting to assign patients until they are medically ready, which for PACU patients may be hours after their bed request, and only assigning patients to beds that are empty (not forecast to be empty at some point in the future). McNichols refines the JIT algorithm to account for variations in acceptable wait length (AWL) for patients from different origins. When combined with discharges earlier in the day, he finds that his algorithm increases the proportion of patients that are placed within their AWL.

This thesis uses simulation methods that are similar to those of McNichols and Hiltrop, but there are differences in the patient populations and unit structures that complicate the simulation of General Medicine units. Namely, while all beds in the neuroscience units are private, most General Medicine beds are semi-private and therefore cohorting requirements must be considered. General Medicine bed assignments are also impacted by physician teams and triage levels (described in Section 3.2.1), complications that were not introduced in previous work. The interventions proposed in this thesis differ from those proposed by McNichols and Hiltrop. JIT bed assignment was not explored in this setting since a large majority of General Medicine patients are assumed to be medically ready at the time of their bed request. Earlier discharges were not proposed as an intervention in this work due to challenges with implementing the recommendations of McNichols in practice at MGH.

Johnston (2016) examines the impact of clinical care team discontinuities on inpatient LOS and admission wait time within the DOM at MGH, with a focus on the house staff teams [5]. He finds that patients admitted to a floor two days before an end-of-rotation Attending physician handoff spend an average of 0.8 days longer in the hospital than otherwise similar patients. He hypothesizes that the discovery and diagnosis-focused activities at the start of a patient's stay are highly sensitive to disruption. This insight informed the development of the early team assignment intervention and its implementation in the simulation.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Current State

This section provides a description of the General Medicine units and provider teams at MGH (Section 3.1) and the associated bed management processes (Section 3.2). The analysis was conducted through numerous interviews with administrative and clinical stakeholders, and through the observation and documentation of work processes in different departments (i.e., the ED, the Admitting department, and General Medicine units).

We corroborated these observations with the analysis of historical data and through discussion with a multidisciplinary working group. This approach allowed us to map the existing bed management processes at the hospital, including their limitations. The findings of the current state analysis inform the modeling approach for the simulation developed in Chapter 4, and the key challenges of the current process presented in Section 3.6 motivate the design of the process interventions that are tested using the simulation.

This project began in February 2016 with the goal of both evaluating the impact of the redesign that was planned for June 2016 and developing complementary interventions. The data used for the simulation are from the calendar year 2015, before the implementation of the redesign. As such, Sections 3.1.1 to 3.2.4 describe the pre-redesign system that we sought to replicate in the base scenario of the simulation and Section 3.3 describes the redesign in detail.

3.1 MGH General Medicine

As previously mentioned, General Medicine is a sub-specialty within the DOM. Prior to the redesign, General Medicine had 230 allocated beds on regionalized units. Additionally, many General Medicine patients occupy 'non-regionalized' beds that are not allocated specifically to General Medicine and are located on floors assigned to other services, with the majority being on surgical floors. General Medicine at MGH is broadly divided into two types of practices, house staff or teaching teams and the Hospital Medicine Group (HMG). These teams and their associated units are described in detail in the following sections.

3.1.1 House Staff Units

House staff teams consist of resident physicians who are still completing their medical training and senior physicians who supervise them. Some of these teams also include a nurse practitioner who works with the physicians, takes responsibility for stable patients near the end of their stays, and helps to coordinate discharges. Prior to the redesign there were five different house staff teams that exhibited several variations in their staffing and the number of patients that they could cover. One thing that these teams had in common was that they only cared for patients in a regionalized setting, meaning that the team only had patients in one physical floor. The number of patients the teams covered varied by the composition of the team, and the team's patient cap did not always correspond with the physical number of beds on the unit. The teams' capacities and units are summarized in Table 3.1.

One effect of this mismatch between team capacity and physical unit capacity is that beds on these units were sometimes assigned to Level 2 patients (who do not require regionalized beds). Such placements could be due to the regionalized team reaching their capacity and therefore beds on the unit only being appropriate for Level 2 patients, or they could be due to a decision made by Admitting to place a Level 2 patient in a bed that would also be appropriate for a Level 1 patient. Overall, approximately 10% of the bed occupancy on the house staff units was by Level 2 patients, with the other 90% by Level 1 patients.

As Table 3.1 indicates, Bigelow teams A, B D, and E had a lower capacity for patients than there were beds on their assigned units. Most beds in these units are semi-private

Team	Patient Capacity	Unit	Unit Bed Capacity
Bigelow A	24	White 8	26
Bigelow B	20	White 9	25
Bigelow C	20	White 10	20
Bigelow D	20	White 11	24
Bigelow E	24	Bigelow 11	25
Team 1	36	Ellison 16	36

Table 3.1: Pre-redesign house staff teams and units

rooms, which means that the patients occupying the two beds in a given room must be compatible. This means that the patients must be of the same sex and either be free of infection precautions or have similar precautions that allow them to be cohorted.

It is common practice at MGH that patients with Methicillin-resistant Staphylococcus aureus (MRSA) or Vancomycin-resistant enterococci (VRE) precautions can be cohorted. Patients with Influenza precautions can be cohorted only if they are known to have the same strain of virus and were infected around the same time. All other infectious precautions require the patient to be alone in a room. When a patient needs to be alone in a room, he can either be placed in a private room or in a semi-private room with the other bed closed. Such patients make up approximately 23% of the General Medicine population and are referred to as 'non-cohortable' in this thesis. Approximately 69% of General Medicine patients have no infection precautions and are referred to as 'clean'.

In addition to infection precautions, there are a number of behavioral and situational reasons that a patient may need to be alone in a room, including disruptive behavior and end-of-life care. Taken together, these issues and infection precautions are referred to throughout this work as isolation status. These requirements for placing patients mean that even though a unit, for example White 11, may have 24 beds, it is not uncommon for only 20 of the beds to be utilized at a time. When this is true the regionalized house staff team would be able to cover all the patients on the unit. When there are more patients placed on the unit than the house staff can cover, the additional patients must be covered by a non-regionalized HMG team (described in the next section). This condition of having more patients on the unit than the regionalized house staff team could cover was very common for White 9 (where it occurred 77.4% of the time) and White 11 (occurring 74.2% of the time) since these units had regionalized team capacity that was significantly less than the number of beds. It happened less frequently on White 8 and Bigelow 11 (18.5% and 7.8% of the time respectively), where regionalized team capacity was more closely matched with bed capacity. This condition did not occur at all on White 10 and Ellison 16, where regionalized team capacity was exactly equal to bed capacity.

3.1.2 Hospital Medicine Group (HMG) Units

HMG teams are different from house staff teams in that they consist of board-certified internal medicine physicians (also known as Hospitalists)[10]. In the calendar year 2015 there were several changes to the structure of the HMG teams at MGH, with the most notable being variations in the amount of capacity available to care for non-regionalized patients. The team names, capacities, and coverage model presented in this section are representative of how patients were cared for throughout the year, despite the slight variations, and were used for the entire duration of the simulation.

Each HMG physician covers nine to eleven patients at a time and rotates between the HMG teams on a regular basis. In contrast with the house staff teams where residents provide most of the direct patient care and complete record-keeping tasks with senior physicians supervising, HMG physicians complete all these tasks personally. HMG teams also differ from house staff in that they are not limited to covering patients on only regionalized floors. Although some of the teams are assigned to a regionalized floor, there are other teams that cover non-regionalized patients exclusively. Whether they have their own regionalized floor or not, all HMG teams can cover non-regionalized patients. These non-regionalized patients could be on a house staff floor on which the patient population has exceeded the regionalized team's cap, or they could be placed on a unit belonging to a different service entirely. In 2015, General Medicine patients were placed on 13 different non-regionalized units[11], most

[10]While the majority of the HMG staff is made up of physicians, there are a few Nurse Practitioners that work on the Orange team. The Orange team is a non-regionalized team that works only with patients with special social needs. For simplicity, the capacities of the Orange team and the HMG lead (known as the White team) have been combined with the Green team for the remainder of this thesis.

[11]These units are Blake 6, Ellison 6, Ellison 7, Ellison 8, Ellison 11, Ellison 14, Phillips 21, Phillips 22, White 6, White 7, Bigelow 14, Lunder 7, Lunder 8.

of which are assigned to surgical services. The reason for this non-regionalized placement is that the volume of General Medicine patients routinely exceeds the effective capacity of the regionalized units, while other services have some degree of available capacity. Criteria for placing patients in non-regionalized beds are provided in Section 3.2.1.

Table 3.2 describes the HMG coverage structure used in the simulation model. It is important to note that the number of patients covered by the non-regionalized Green team was highly variable and dependent on the capacity situation of the hospital, the staffing level of HMG, and the acuity of the patients being covered. Figure 3-1 presents a graphical representation of the variation in non-regionalized census throughout the year[12]. In this context, non-regionalized refers only to the beds that are on floors allocated to services other than General Medicine, and not those that are on units allocated to General Medicine but covered by non-regionalized teams. The maximum 7 am non-regionalized census is found to be 75 patients and the minimum 29, with an average of 44.7 patients. Starting in June, HMG leadership began to reduce and enforce physician team caps, and the effects of this can be observed in Figure 3-1. Section 4.2.4 discusses how the simulation accounted for this variability.

3.1.3 Nursing Teams

Nurses form an important part of the patient care team. In the context of General Medicine, nursing teams are always unit-based. This means that for a patient on White 9, regardless of whether they are covered by a house staff team or an HMG team, they will be cared for by the White 9 nursing team. This is also true for patients placed on non-regionalized units that are assigned to services other than General Medicine, they will be cared for by the nursing team that is local to that unit.

There are a variety of roles within the nursing team for each unit, including the staff nurses who directly care for the patients, an attending nurse who assists with discharges and removing roadblocks to patient care, and a nurse supervisor who oversees the operations of the unit. One role that is of particular interest in the context of this project is the Resource

[12] The census is a count of the number of patients present at a certain time of day for each day of the time period.

Team	Patient Capacity	Regionalized Unit(s)	Unit Bed Capacity	Notes
Red	10	Bigelow 9 and White 9	12	Red team covers White 9 patients not covered by the Bigelow B and is regionalized on both White 9 and Bigelow 9
Blue	36	Ellison 12	36	
Yellow	34	Ellison 19 and Phillips 20	30	Ellison 19 is shared with Thoracic surgery and the average number of beds available to General Medicine is 10. In history there were as many as 14 patients on Ellison 19 at once, so the team cap is set to accommodate this (plus 20 patients on Phillips 20).
Green	Varies	Non-regionalized	Varies	Green team covers only non-regionalized patients

Table 3.2: Pre-redesign Hospital Medicine Group (HMG) teams and units

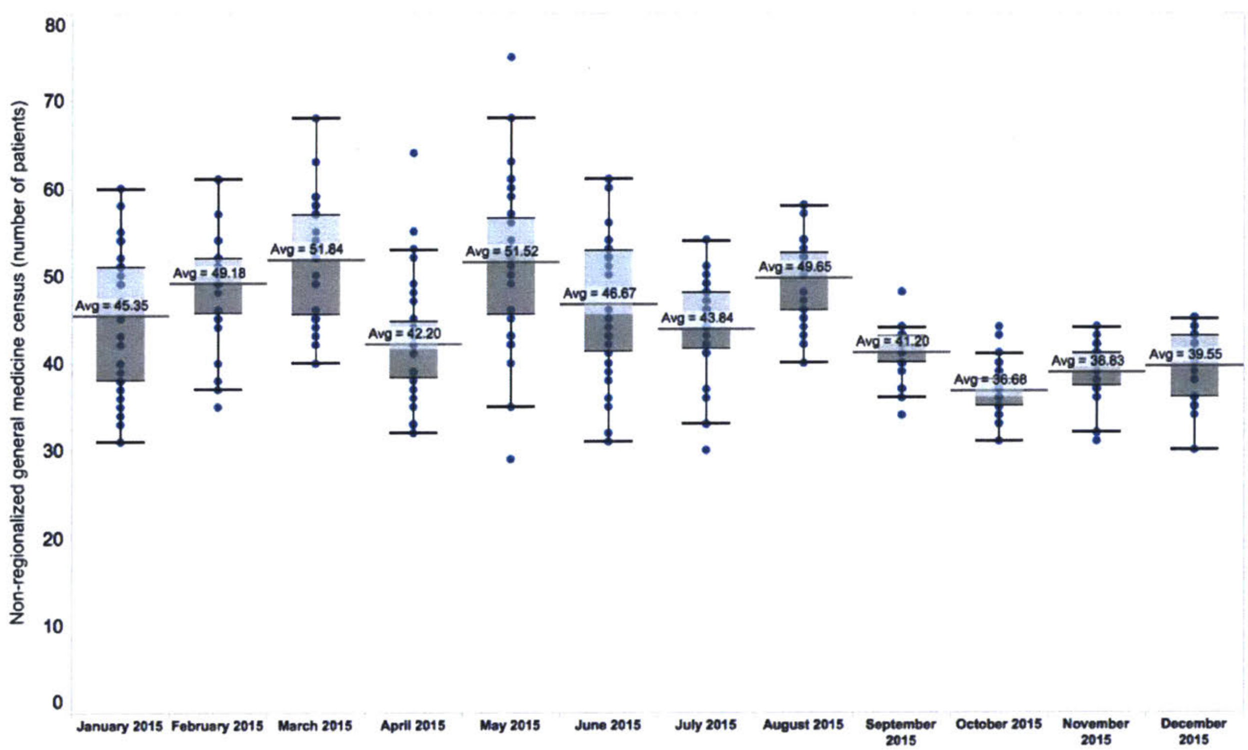

Figure 3-1: Historical non-regionalized General Medicine patient census

Data sources: ADT (tbl EDW_PtADTDept), PEPL (Tbl_Encounter_PEPLCareUnitFact_MGH) Timeframe: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients.

Note: In this context, non-regionalized refers only to beds on floors allocated to services other than General Medicine.

nurse, who is responsible for reviewing and approving patient placements on the unit. This process is discussed in detail in Section 3.2.1.

The nurse supervisor is responsible for creating nurse staffing plans to make sure that all patients on the unit get the nursing care they need. However, on occasion the nurse staffing is not sufficient to cover all the beds in the unit. This can occur when there are an unexpected number of nurses absent or when the acuity of the patients on the unit is unusually high (or a combination of the two). When this occurs, the nurse supervisor will communicate the situation to Admitting who will close an appropriate number of beds until more staffing is available or the units' acuity decreases.

3.1.4 ED Boarder Service

Due to the significant delays in placing patients in General Medicine beds, there was an initiative in 2015 to start covering some of the patients boarding in the ED with General Medicine physicians while they awaited a bed. Patients are defined by MGH as boarders when two hours have passed since their bed request and they are still waiting in the ED. During 2015, General Medicine physicians were covering up to ten ED boarders. The boarder program was first staffed by the house staff teams and then passed over to HMG. It was discontinued in June of 2015 due to a lack of staffing to support it and was later restarted in 2016 when physician staffing levels increased.

While it operated in 2015, the boarder service worked as a bridge from ED care to the General Medicine unit care. A handoff was performed from the ED physician to the boarder team physician and then another handoff occurred between the boarder physician and the physician who cared for the patient on the unit.

3.2 Bed Management

On an average weekday in 2015, the DOM at MGH admits about 64 patients, 49 of whom are General Medicine[13]. As previously mentioned, the ED is the main source of these admissions. All of these patients have to be matched to appropriate hospital beds in a timely manner.

[13]DOM statistic obtained from the Program Manager of Clinical Operations in the Department of Medicine, General Medicine statistic calculated from ADT data.

For General Medicine, the situation is complicated by the distinction between Level 1 and Level 2 patients, the provider team patient caps, the need to cohort patients in semi-private rooms, which place additional constraints on the possible bed-patient assignments.

As previously mentioned, there is a considerable mismatch between the pattern of patient arrivals (bed requests) throughout the day and the pattern of discharges. In 2015, the average time for beds to be requested was 1:50 pm and the average time for discharges was 2:59 pm. In addition to discharges occurring over an hour later on average than bed requests, they are far more time concentrated than the requests with a standard deviation of only 3.16 hours compared to 6.47 hours for bed requests. This difference in timing of bed requests and discharges is shown in the histograms in Figure 3-2, which represent the number of bed requests and discharges at each hour of the day in 2015.

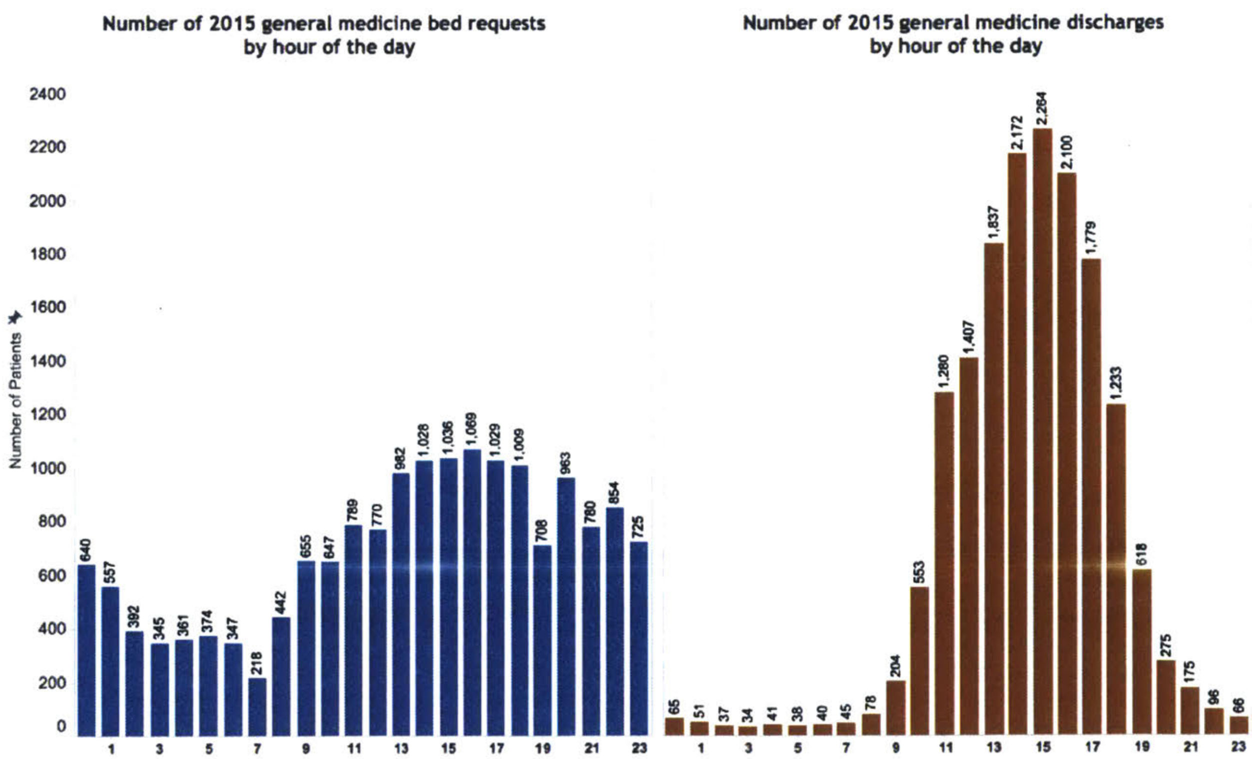

Figure 3-2: Histograms of 2015 bed requests and discharges by hour

Data sources: ADT (tbl EDW_PtADTDept), EPIC (MGH_BED_PENDING) Timeframe: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients.

3.2.1 The Bed Assignment Process

The Admitting department is responsible for bed management, working closely with representatives from the units. The individuals within Admitting with the most hands-on role in this process are the bed managers. On a given day there are four bed managers on duty, one for the surgical units, one for medicine, one for the lower volume services including pediatrics, urology and orthopedics, and one specifically handling hospital transfers[14]. The primary goal of the bed managers is to efficiently match waiting patients with available beds.

The main tool that supports the bed assignment process is EPIC, the hospital's operations and electronic medical record system. EPIC provides a graphical user interface with detailed information about the current status of each bed. The system distinguishes between the following bed statuses:

- Occupied: A patient is currently occupying the bed.
- Dirty: The bed is not occupied by a patient and needs to be cleaned.
- In cleaning: The bed is not occupied by a patient and currently being cleaned.
- Clean: The bed is not occupied by a patient and clean (i.e., ready for occupancy).
- Closed: The bed is unavailable due to neighbor isolation status, staffing shortages, repairs, maintenance, or other reasons.

Each status in EPIC can be combined with the following additional information:

- Assignment: The bed can be assigned to a waiting patient. This applies to beds that are available (i.e., clean) or expected to become available in the foreseeable future.
- Pending discharge: If the patient that is currently occupying the bed has been identified to be discharged that day, the unit is expected to update EPIC to show that the patient is now 'pending discharge'. This causes an icon to appear next to the patient's name and allows the Admitting team to assign the next patient to the bed. The unit can also provide an expected time of discharge, but this is rarely populated.
- Precautions: Medical precautions for the current occupant of the bed can be specified.

[14]Hospital transfers are patients who come to a MGH unit from another inpatient facility. When the patient is assigned to General Medicine, these transfers must be approved by a medicine senior resident before Admitting agrees to place them.

- Additional information: The reason for why a bed is closed, specific bed features (e.g. negative atmospheric pressure, ADA accessibility), and other information about the bed.

The information in EPIC is updated by bed managers in Admitting, Operations Associates (OAs) in the different inpatient units, and bed cleaners. Bed managers continually monitor the EPIC interface throughout the day, looking for new patients to assign and beds that have become available for assignment. A bed is available for assignment when the patient that previously occupied it leaves or is identified as pending discharge, or when a bed closure ends.

When a bed manager assigns a patient to a unit, she generally provides a note with the specific bed that she has identified for the patient. The Resource nurse confirms this assignment based on her research into the patient's medical record and works with the OA to finalize the assignment in EPIC. In addition, OAs enter information about (expected) discharges from their units into EPIC. They enter pending discharges when they are notified by the Resource nurse and they also enter actual patient departures. Bed cleaners update the bed status when cleanings are started and completed using a phone dial-in system that is linked to EPIC.

Bed managers follow a general set of guidelines to prioritize bed assignments among patients admitted to the same medical service and level of care. For General Medicine, patients arriving through the ED or the PACU receive the highest priority, followed by those being admitted directly to the hospital from an outpatient setting (known as 'front door'), then patients transferring from other general care units or ICUs, and lastly hospital transfers. This prioritization order is not documented and is subject to exceptions based on the particular situation of the patient and the hospital. The logic behind this prioritization is that patients who are not currently in an inpatient setting and are using valuable MGH resources (ED and PACU space) are prioritized first, followed by those who are not receiving inpatient care and are not using MGH resources (front door), then those that are receiving inpatient care and using MGH resources (floor and ICU), and lastly those that are receiving inpatient care and are not consuming MGH resources (hospital transfers). This prioritization therefore seeks to optimize for both patient care and MGH capacity. The specifics of the bed assignment process for the various patient origins are discussed in the following sections and summary statistics on volume of completed bed requests by origin are presented in Table 3.3. It should be noted that the patient's origin always refers to the location of the patient immediately prior to his arrival on the General Medicine unit. For example, if a patient presented at the ED before being admitted to the ICU and was later transferred to a General Medicine unit, he would be referred to as an ICU patient for the purposes of this analysis.

Origin	**Level 1 Admissions**	**Level 2 Admissions**	**Total**	**% of Total Admissions**	**% Level 1**
ED	**7298**	**6050**	**13348**	**80%**	**55%**
Direct Admission	**1237**	**609**	**1846**	**11%**	**67%**
Front door	933	569	1502	9%	62%
Hospital transfer	304	40	344	2%	88%
MGH internal	**1371**	**155**	**1526**	**9%**	**90%**
ICU	892	115	1007	6%	89%
Floor	171	194	365	2%	47%
PACU	62	92	154	1%	40%
Total	**9906**	**6814**	**16720**	**100%**	**59%**

Table 3.3: Summary statistics for 2015 General Medicine admissions by origin
Data sources: ADT (Tbl_Encounter_ADT_MGH), PEPL (Tbl_Encounter_PEPLCareUnitFact_MGH), EDIS (EDtbl_ED_InpatientBedRequestData), EPIC (MGH_BED_PENDING). Filtered for General Medicine patients with bed requests created in 2015.

For General Medicine patients, a distinction is made between patients that are 'Level 1' and those that are 'Level 2'. This attribute is known within the DOM as the triage level. These levels describe the clinical workload associated with caring for the patient. Level 1 patients are expected to require more physician resources, due to their acuity and/or the complicated nature of their medical condition. Level 2 patients can be placed in non-regionalized settings, whereas Level 1 patients must be cared for by a regionalized physician team[15]. Admitting is notified of a patient's triage level at the time of the bed request and looks for a bed that will be appropriate for that patient. Due to the large number of Level 2 patients that are placed on non-regionalized units, these bed requests are sometimes assigned

[15]Even though Ellison 12 was covered by a regionalized HMG team for much of 2015, assignments to this unit were primarily limited to Level 2 patients. In the base scenario of the simulation we have maintained this practice to better mimic the historical state of the system.

to the bed manager who is working with the surgical units.

3.2.2 Bed Assignment for ED Patients

As previously mentioned, approximately 80% of General Medicine patients originate from the ED. When a patient first arrives in the ED, the ED physicians work to stabilize and diagnose the patient. When they decide that the patient should be admitted as an inpatient to General Medicine, they contact a Senior medicine resident (known as the 'med senior') who screens the patient to verify that he is appropriate to admit to medicine. At this time the med senior and the ED physicians also agree to the triage level for the patient. The information about the patient's triage status is then passed on to Admitting as part of the bed request.

Bed requests for ED patients are largely prioritized on a first-come-first-served basis. This is not strictly true, of course, since differences in bed availability for patients of different genders, isolation statuses, and triage levels affect the order in which patients are assigned. ED General Medicine patients are generally prioritized over those from other sources due to concerns of overcrowding in the ED. Once a bed is found for an ED patient, a handoff must be performed between the ED and medicine nurse and physician teams before the patient can be transported to the General Medicine unit.

3.2.3 Bed Assignment for Direct Admissions to the Hospital

Patients that are directly admitted to General Medicine (not coming from another area of MGH) represented approximately 11% of 2015 admissions. For the purpose of this simulation these patients were broadly split into two groups, front door admissions and hospital transfers. Front door refers to patients who were referred from a doctor's office or other outpatient facility and were not in an inpatient setting prior to their bed request. These are the majority of direct admission patients. The priority given to front door patients is largely based on the severity of their condition and their physical location. Many of these patients are given instructions to wait at home until there is a bed available for them, but some come to the hospital to wait for a bed. Discussion with bed managers suggest that patients waiting in the hospital are likely to get a bed more quickly than if they are waiting at home.

Hospital transfers are patients who are transferring from another inpatient setting to MGH. This is usually because they require a higher level of care than is available at their current hospital. General Medicine hospital transfers are facilitated by the dedicated bed manager. Hospital transfer patients are generally de-prioritized because they are currently receiving inpatient care and are not causing congestion at MGH. Direct admissions are subject to the same triage procedures as ED patients, however in 2015 there was not a clear place to register the triage level in the bed request. The majority of these patients (approximately 67%) were determined to be Level 1 in 2015.

3.2.4 Bed Assignment for within Hospital Movers

Patients that are transferring to General Medicine from non-ED locations within MGH account for roughly 9% of the 2015 admissions and are categorized as being from the ICU, floor, or PACU. ICU patients can originate from any of the six adult ICUs at MGH. Once a patient's care team in the ICU has determined that he no longer requires ICU-level care, they will submit a request for a bed on a General Medicine unit. When a General Medicine bed is ready for the patient, the physician and nursing teams in the ICU need to handoff to the General Medicine teams. By submitting this bed request, the ICU team asserts that the patient is medically ready to move on to a lower level of care, however there are instances where a patient is in fact not ready when a bed is available for him on the general care unit. Instances like this tend to stick in the bed managers' minds and may contribute to a systemic de-prioritization of ICU patients for General Medicine bed assignments. However, the primary reason articulated by MGH staff that these patients are de-prioritized is that they are already receiving excellent care in the ICU and it is often the case that there is more congestion in the general care units than the ICUs, meaning that there is no urgency to move them out. As discussed in Chapter 2, this approach may cause unintended problems since delays leaving the ICU have been shown to contribute to increases in the overall hospital LOS [6]. Discussions with MGH staff suggest that the primary impetus to move a patient out of the ICU is when her bed is required for a new ICU patient. This practice could result in increased wait times for patients that require ICU beds, which was not measured as part of this analysis.

In addition to the ICU, patients can also come to General Medicine units or from other general care floors. This situation would likely occur when a patient needed to change clinical service due to a change in condition or diagnosis. The majority of these patients came from floors assigned to services other than General Medicine, with only 97 instances of patient moving between General Medicine units observed in the data. Transfers between general care units are handled in much the same way described above, but these patients are generally given a higher priority than the ICU patients due to congestion in the general care units.

The remaining source for patients arriving in General Medicine units from within MGH is the PACU, which is the area where patients recover from surgeries and other procedures requiring anesthesia. Most medicine patients do not require surgery, and therefore the population of PACU patients is very small (less than 1% of total General Medicine admissions in 2015). PACU patients have their bed requests generated when they check in for their procedure in the morning and are prioritized based on the expected end time of their procedure. Unlike patients from other sources, PACU patients are not actually ready to occupy a General Medicine bed at the time of their bed request. Section 3.4.1 describes the procedure for estimating when PACU patients become ready to occupy their requested beds.

3.3 The Department of Medicine Redesign

As previously mentioned, the redesign was conceived prior to the start of this project by administrators, physicians, and nurses within the DOM. The redesign aimed to increase the capacity of the department and therefore shorten patient wait times and reduce boarding in the ED, while also supporting key teaching goals. The elements of the redesign are discussed in detail in the following sections and the changes are summarized in Figure 3-3.

3.3.1 Additional Regionalized Capacity

The redesign added 10 allocated General Medicine beds on Bigelow 9. Pre-redesign, the Bigelow 9 unit was split between General Medicine and the respiratory acute care unit (RACU). This meant that General Medicine patients could be placed in only 8 out of the 18 beds. With the redesign and the relocation of the RACU to another area of the hospital, all 18 beds are now allocated to General Medicine and covered by a regionalized team.

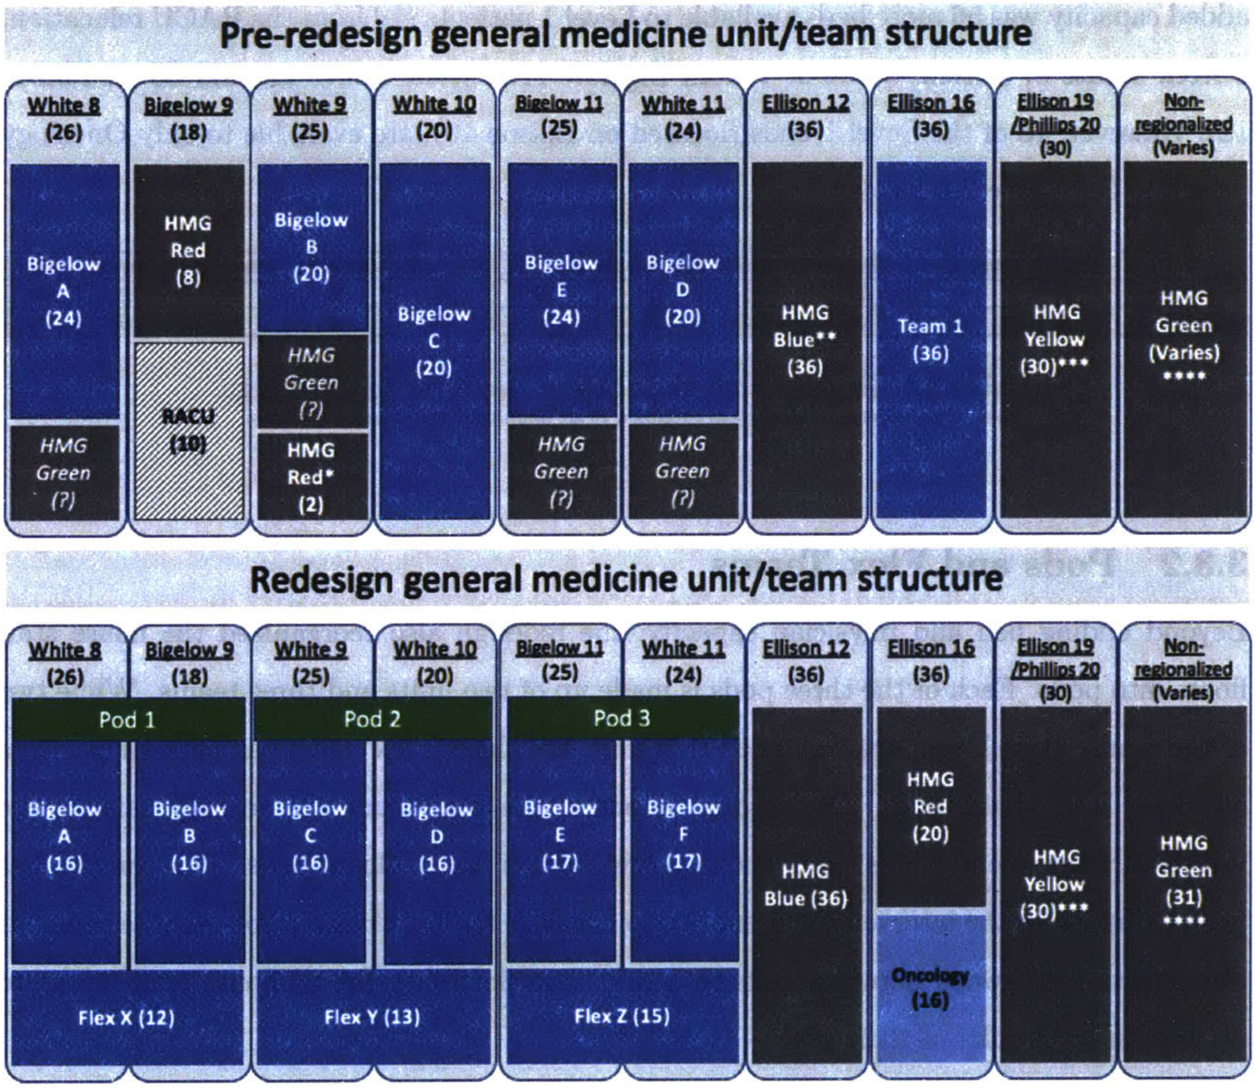

Figure 3-3: Summary of redesign changes

Notes: *Pre-redesign, the HMG Red team covered Level 1 patients on White 9. **Pre-redesign, all beds on Ellison 12 were Level 2 even though the Blue team was regionalized. *** The allocated number of General Medicine beds on Ellison 19 was 10, but the actual number utilized varied based on demand for beds by Thoracic surgery. **** The Green team's capacity includes the Orange and White teams as mentioned in Section 3.1.2.

In addition to adding bed capacity, the redesign added physician capacity and better aligned teams with units to create more regionalized beds that are appropriate for Level 1 patients. Whereas White 8, White 9, Bigelow 11, and White 11 previously had regionalized teams that could not cover all the physical beds on the unit, the redesign added capacity to the house staff teams and eliminated this problem. Furthermore, while Ellison 12 had previously been limited to Level 2 patients due to established practice, the redesign included training and communication that this unit was now fully regionalized and that Level 1 patients could and should be assigned to Ellison 12. The net result of these changes and the added capacity was 56 more beds available to Level 1 patients (10 from the RACU relocation, 2 from White 8, 3 from White 9, 1 from Bigelow 11, 4 from White 11, and 36 on Ellison 12). However, 16 of the Level 1 beds (located on Ellison 16), are available to only Oncology patients. See Table 3.4 for details.

	White 8	Bigelow 9	White 9	White 10	Bigelow 11	White 11	Ellison 12	Ellison 16	Ellison 19	Phillips 20	Total
Pre-redesign	24	8	22	20	24	20	0	36	10	20	184
Redesign	26	18	25	20	25	24	36	36	10	20	240

Table 3.4: Regionalized beds before and after the Redesign
Note: Post-redesign, 16 of the beds on Ellison 16 are available to Oncology patients only.

3.3.2 Pods and Flex Teams

Beyond adding bed and physician capacity, the redesign also reorganized the house staff floors into pods. Each of the three pods is made up of two units and three teams. While two of the teams in each pod are dedicated to a unit, the third 'Flex' team can care for patients on either unit of the pod. The original intent of the Flex team was to help balance the workload when one unit was experiencing a low volume due to closed beds. This concept has been extended in the intervention scenarios to consider the implications for early assignment of patients, bed swaps between units in a pod, and flexibility for all house staff teams to operate in both units of their pod.

3.3.3 Oncology Sub-unit

The last major aspect of the redesign is the creation of a sub-unit for Oncology patients. While there are dedicated units for Oncology patients that are separate from General Medicine, these units are frequently at capacity and therefore patients who are receiving chemotherapy during their hospital stay are prioritized for placement on these floors. Oncology patients who are not due to receive chemotherapy are often placed on General Medicine units where they will be cared for by General Medicine teams in addition to their oncologist. Pre-redesign, these patients were placed throughout the General Medicine floors. To support the educational goal of providing residents with exposure to Oncology patients during one of their rotations, the decision was made to create a sub-unit for these Oncology patients on Ellison 16. In the redesign Ellison 16 is covered by two different teams, one is the aforementioned Ellison Oncology teaching team, which can cover up to 16 patients, the other is the Red team that can cover the maximum remaining load of 20 patients.

3.4 Performance Metrics

We used two key performance metrics to analyze the current state and the anticipated impact of the interventions: patient-wait-for-bed, and patient-wait-for-team. These metrics were first calculated for the 2015 data to better describe the historical performance of the system, enable model validation, and evaluate the effectiveness of the interventions. Descriptions of the metrics follow below and analysis of the historical data appears in Section 3.5.

3.4.1 Patient-wait-for-bed

Patient-wait-for-bed is defined as the amount of time that passes from when the patient is medically ready to go to her General Medicine bed to the time at which the bed is ready for that patient. Patients from all origins except the PACU are considered to be medically ready at the time of their bed requests. It should be noted that for ED patients, the bed request is likely to be placed several hours or more after arrival in the ED and this wait time does not capture any of the time prior to bed request. Since PACU patients have their bed requests automatically generated in the morning before their procedures, they are considered to be medically ready two hours after physically arriving in the PACU or upon their 'ready to depart PACU' timestamp, whichever is earlier[16]. A bed is determined to be ready for the patient once it is open, empty and clean, the patient has been assigned to the bed, and the patient is medically ready.

Patient-wait-for-bed does not capture any of the transfer processing time associated with physically getting the patient to the ready bed. Transfer processing time includes nurse and physician handoffs, wait for transportation personnel, and the actual transport of the patient. The reason that this time was omitted from patient-wait-for-bed is that none of the interventions studied are intended to directly impact it. Transfer processing time took an

[16]The reason that the 'ready to depart PACU' timestamp was not used exclusively is that there was a belief among the process experts that this timestamp is often entered when the patient is about to leave the PACU and is not an accurate representation of medical readiness. There are a small number of patients who are intentionally kept in the PACU overnight for observation, for these patients the 'ready to depart PACU' timestamp was used exclusively when determining medical readiness.

average of 2.1 hours for General Medicine patients in 2015.

Figure 3-4 presents several likely sequences of events with patient-wait-for-bed defined for each. This metric is calculated for all General Medicine patients who requested and reached their beds in 2015, a total of 16,720 patients[17].

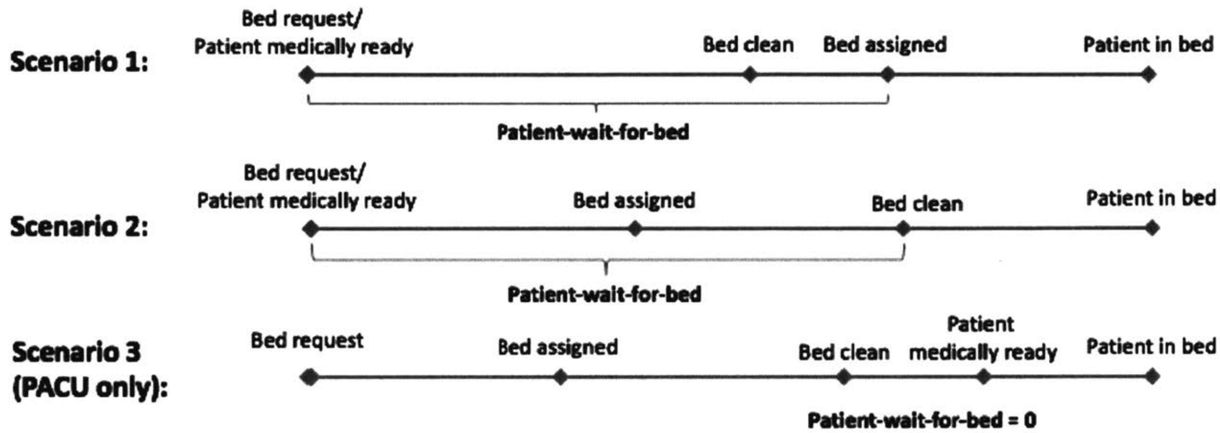

Figure 3-4: Illustration of patient-wait-for-bed

3.4.2 Patient-wait-for-team

Patient-wait-for-team is calculated for ED patients only and is intended to illustrate the impact of the early team assignment intervention. Like patient-wait-for-bed, it starts when the patient is medically ready (which for ED patients is always at the time of bed request) and continues until the physician handoff from the ED to the General Medicine team is complete. It is generally longer than patient-wait-for-bed since handoffs normally take place after the bed is ready for the patient[18]. This metric is calculated for all ED patients who reached their General Medicine beds in 2015 and had a physician handoff timestamp in EDIS

[17]Note that 'patients' refers to patient stays on a General Medicine unit. A single patient could be counted multiple times if he or she had multiple stays on General Medicine units. For example, if a patient was admitted through the ED to White 9 then went to an ICU and then to Ellison 16, he or she would count as both an ED patient on White 9 and an ICU patient on Ellison 16.

[18]Historical practice regarding when physician handoffs were completed was not totally consistent. The large majority (88%) of recorded handoffs took place after the bed was clean and ready for the patient, but in some instances the handoff took place before. In the simulation, all handoffs take place after the bed is ready for the patient and historical metrics have been adjusted so that the patient-wait-for-team cannot be shorter than the patient-wait-for-bed.

(the ED's primary tracking system), a total of 13,302 patients[19].

3.5 Analysis of the Current State

3.5.1 Patient-wait-for-bed Analysis

In 2015, the average patient-wait-for-bed across all patient origins was 7.44 hours, with a median of 3.14 hours and a 25th percentile and 75th percentile of 0.80 and 9.47 hours respectively. Overall, Level 1 patients waited longer than Level 2 patients. Level 1 patients experienced an average wait of 8.45 hours and a median wait of 3.47 hours. Level 2 patients experienced an average wait of 6.05 hours and a median wait of 2.73 hours. Patients with no precautions or issues (referred to as 'clean') waited the shortest amount of time (6.18 hours on average), while patients with a status other than clean waited 10.28 hours on average. The mean wait times for these populations along with the quantiles are summarized in Table 3.5.

	All Patients	Level 1 Patients	Level 2 Patients	Clean Patients	Other than Clean Patients
Number of Records	16,720	9,660	7,060	11,537	5,183
Mean Patient-wait-for-bed (hours)	7.44	8.45	6.05	6.16	10.28
P5 Patient-wait-for-bed (hours)	0.07	0.08	0.07	0.07	0.08
P25 Patient-wait-for-bed (hours)	0.80	0.90	0.68	0.68	1.28
Median Patient-wait-for-bed (hours)	3.14	3.47	2.73	2.67	4.55
P75 Patient-wait-for-bed (hours)	9.47	10.40	8.17	7.68	13.85
P95 Patient-wait-for-bed (hours)	24.77	28.33	21.00	21.42	31.36

Table 3.5: Historical patient-wait-for-bed hours by patient characteristics
Data sources: ADT (Tbl_Encounter_ADT_MGH) , EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

There was a great deal of variation in wait time by patient origin. As discussed in Section 3.2, ED patients are generally prioritized and, as such, they experienced the second lowest wait time with an average of 5.71 hours. PACU patients experienced the shortest waits since they are often assigned to beds before they are medically ready (in this case patient-wait-for-bed is zero). However, recall that only 1% of General Medicine patients originated in the PACU. In contrast, ICU and hospital transfer patients wait an average of over 20 hours for their beds. The maximum wait times for floor, ICU, and hospital transfer patients

[19]There were 13,348 total General Medicine patients originating from the ED in the population, 42 were excluded from this calculation due to the absence of a physician handoff timestamp.

all exceeded 200 hours. These wait times were calculated using the procedure described in Section 3.4.1 and the data described in Section 4.1. Although they are technically correct based on the definitions used, it is likely that there were some extenuating circumstances that led to these patients being delayed. Therefore, it is more informative to consider the 95th percentiles as measures of excessive wait instead of the maximums. See Figure 3-5 for a graphical and tabular representation of patient-wait-for-bed by patient origin.

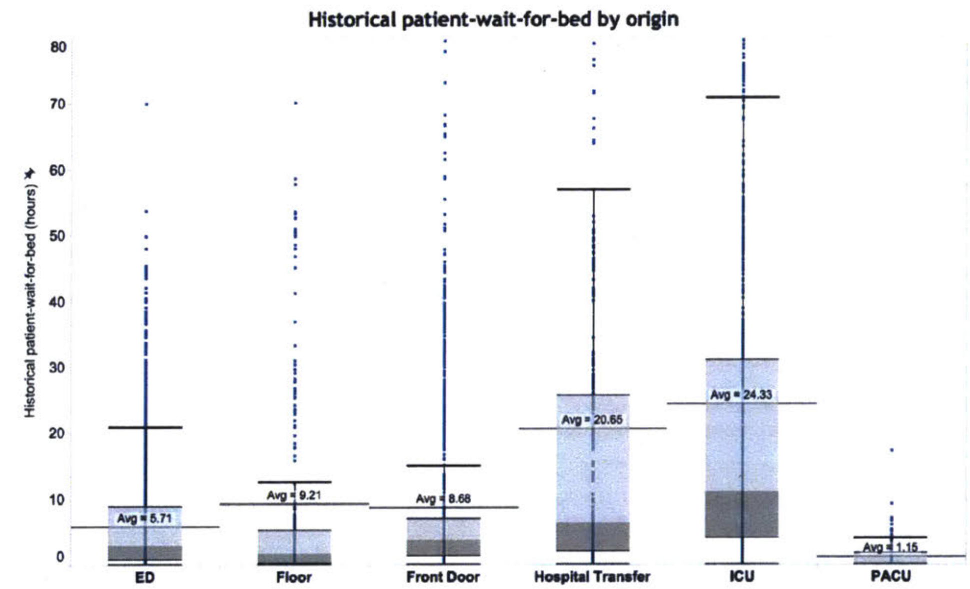

	ED	Floor	Front Door	Hospital Transfer	ICU	PACU	Grand Total
Number of patients	13,348.00	365.00	1,502.00	344.00	1,007.00	154.00	16,720.00
Average	5.71	9.21	8.68	20.65	24.33	1.15	7.44
P5 Historical	0.08	0.00	0.03	0.08	0.45	0.00	0.07
P25 Historical	0.72	0.17	1.34	1.98	4.04	0.00	0.80
Median Historical	2.85	1.78	3.69	6.33	11.07	0.00	3.14
P75 Historical	8.77	5.13	6.97	25.69	30.91	1.66	9.47
P95 Historical	19.45	50.52	39.23	81.01	84.28	4.62	24.77
Maximum	70.00	222.35	152.60	214.53	293.32	17.28	293.32

Figure 3-5: Historical patient-wait-for-bed hours by origin
Note, vertical axis of graph truncated for readability, maximum values shown in table
Data sources: ADT (Tbl_Encounter_ADT_MGH) , EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, 2015. Filtered for General Medicine patients with bed requests created and completed in 2015.

3.5.2 Patient-wait-for-team Analysis

As expected, at 6.99 hours the average patient-wait-for team for ED patients was longer than their patient-wait-for bed which averaged 5.71 hours. This is because the physician handoff does not normally take place until after the bed is ready for the patient and this step can be subject to delays. Again, Level 1 patients wait slightly longer on average than Level 2 patients (7.08 vs. 6.88 hours), and again 'clean' patients have the shortest wait (6.28 hours on average for 'clean' vs. 8.70 hours for other isolation statuses).

3.5.3 Occupancy and Utilization Metrics

In 2015, the average 7 am census in allocated General Medicine beds was 208.3 patients. This implies an average 7 am utilization of 90% for a capacity of 230 beds. This figure does not tell the entire story, however, due to the challenges of cohorting patients in semi-private rooms. As mentioned in Section 3.1.2, the census of patients placed on units that were not allocated to General Medicine was highly variable in 2015 with a 7 am average of 44.7 patients. Since there is not a set level of capacity for General Medicine patients on these units, it is not possible to calculate the utilization.

Figure 3-6 provides a summary of the average weekday 7 am General Medicine (regionalized and non-regionalized) census by month in 2015. This figure shows considerable variability by month, which is largely due to the changes in HMG physician capacity and the variable non-regionalized census discussed in Section 3.1.2. In the figure, the blue line shows the physician team caps in place each month, the green line represents the regionalized bed capacity, and the orange shows the average effective regionalized beds (accounting for closures). Note that as the physician caps started to decrease in June, patient census occasionally exceeded the established cap.

Another metric that is informative when considering the historical performance of the system is the number of ED boarders. The overall average number of ED boarders at 7 am was 8.19, with a median of 7 and a maximum of 28. Figure 3-7 shows how the weekday 7 am census of ED boarders changed by month and suggests that the reduction in HMG team caps near the end of the year may have led to an increase in boarders.

3.5.4 Discharge Communication

Current bed management practices delay the assignment of patients to beds and teams until a discharge is communicated to Admitting. This discharge can be expected (a pending discharge) or can have already occurred. The General Medicine units do not consistently communicate expected discharges to Admitting, which limits the bed managers' ability to

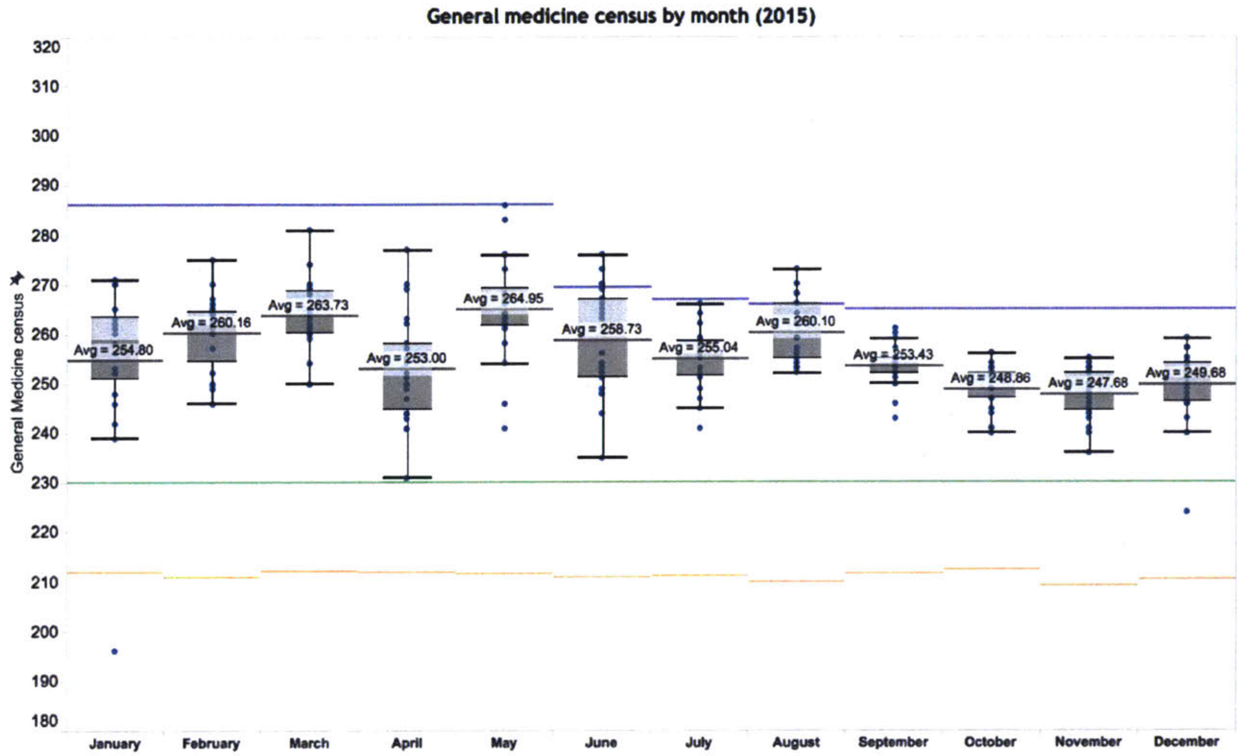

	January	February	March	April	May	June	July	August	September	October	November	December
Number of Records	20.0	19.0	22.0	22.0	20.0	22.0	23.0	21.0	21.0	22.0	19.0	22.0
Avg. General Medicine census	254.8	260.2	263.7	253.0	265.0	258.7	255.0	260.1	253.4	248.9	247.7	249.7
Median General Medicine census	259.0	262.0	264.0	251.5	264.0	259.5	256.0	259.0	254.0	249.0	248.0	251.0
P10 of General Medicine census	241.7	249.8	254.5	241.2	253.2	248.0	247.4	253.0	250.0	241.3	240.8	243.3
Avg. General Medicine phys. cap	286.0	286.0	286.0	286.0	286.0	269.5	267.0	266.0	265.0	265.0	265.0	265.0
Avg. General Medicine beds	230.0	230.0	230.0	230.0	230.0	230.0	230.0	230.0	230.0	230.0	230.0	230.0
Avg. Effective General Medicine beds	212.0	210.9	212.1	211.8	211.8	210.8	211.2	210.0	211.5	212.4	209.2	210.4

Figure 3-6: General Medicine patient 7 am census, 2015

Data source: ADT (tbl EDW_PtADTDept), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING), PEPL (Tbl_Encounter_PEPLCareUnitFact_MGH), Provider team caps supplied by M. Bravard Time frame: Jan 1, 2015 – Dec 31, 2015 holidays and weekends excluded, Filtered for General Medicine patients.

plan the day and assess the available capacity. In 2015 Admitting received advance communication of only 49% of the General Medicine discharges and when pending discharges were communicated the average time between notification and discharge was only 2.5 hours.

3.6 Key Findings

The analysis of the system and the historical data reveal three key findings that guide the development of the proposed interventions.

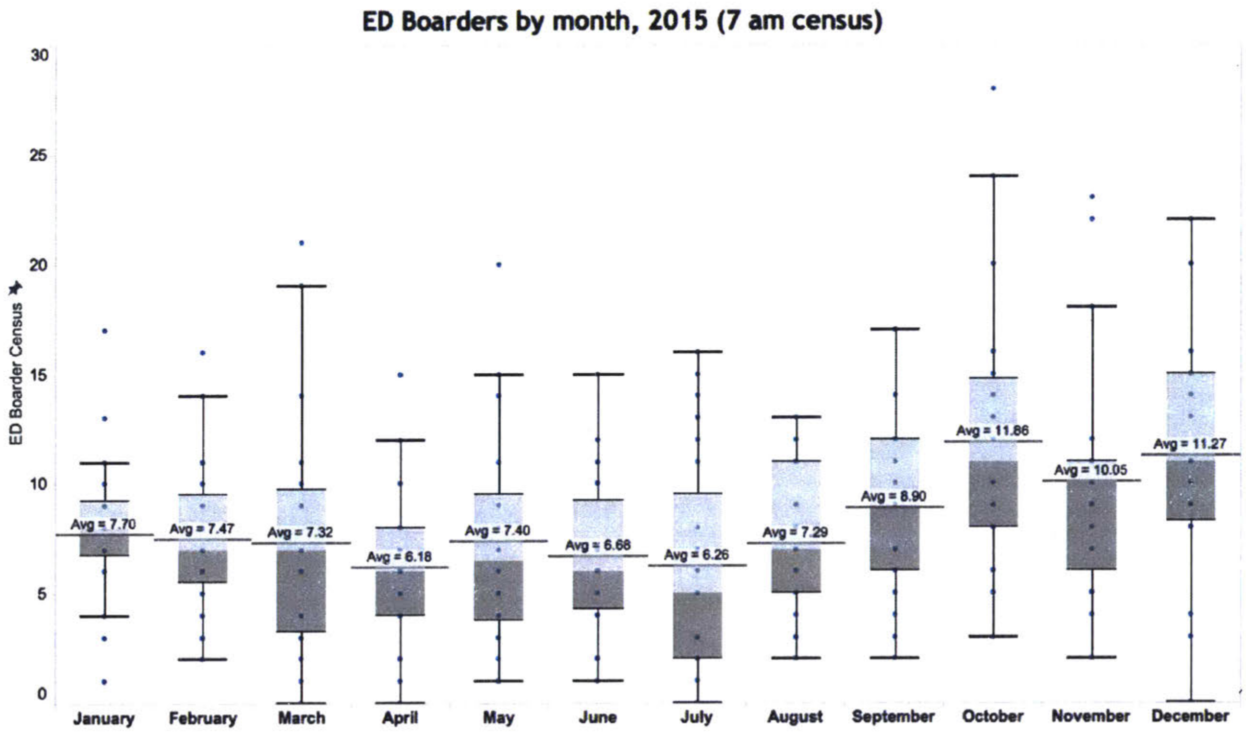

	January	February	March	April	May	June	July	August	September	October	November	December
Number of Records	20.00	19.00	22.00	22.00	20.00	22.00	23.00	21.00	21.00	22.00	19.00	22.00
Avg. ED Boarder Census	7.70	7.47	7.32	6.18	7.40	6.68	6.26	7.29	8.90	11.86	10.05	11.27
Median ED Boarder Census	8.00	7.00	7.00	6.00	6.50	6.00	5.00	7.00	9.00	11.00	10.00	11.00

Figure 3-7: ED Boarder 7 am census, 2015

Data source: ADT (tbl EDW_PtADTDept), EDIS (EDtbl_ED_InpatientBedRequestData, and EDObstbl_ED_InpatientBedRequestData), Time frame: Jan 1, 2015 – Dec 31, 2015 holidays and weekends excluded, Filtered for General Medicine patients from the ED and had ED Boarder flag in EDIS and bed request 2 hours or more before census time.

3.6.1 Prioritization of Bed Assignments not Based on Medical Readiness

As Figure 3-5 clearly shows, wait times vary widely for patients from different origins. An average ICU mover waits more than four times as long as an average ED admission. Although there are valid reasons for this prioritization in the current setting, notably the desire to ease ED congestion, the results indicate that a change in the prioritization methodology could significantly decrease average wait for ICU and hospital transfer patients, and potentially the overall average patient-wait-for-bed time.

3.6.2 Incentives Discourage Timely Discharges

In this highly-constrained system, the units generally operate with no slack capacity. This means that a new admission is only possible when a patient is discharged. Since the highest workloads for the nursing and physician teams occur at admission and discharge, the decision to discharge a patient brings with it a peak in demands on the staff. This phenomenon may lead physicians and nurses to prioritize other activities or to err on the side of caution when discharging patients, which could lead to increased wait times. The data represented in Figure 3-2 show that discharges are heavily concentrated in the afternoon hours, while bed requests arrive more evenly throughout the day, which leads to intraday wait times. This work proposes the intervention of early team assignment, which is aimed to begin to address this incentive misalignment.

3.6.3 Semi-private Rooms Complicate Bed Assignment

This work differs from previous MGH-MIT Collaboration studies of bed assignment processes since the units involved are largely made up of semi-private beds. Analysis of patient-wait-for-bed by isolation status indicates that those patients with a status other than clean, and particularly those patients who are non-cohortable, wait longer for a bed in this environment. This finding suggests that interventions that introduce more flexibility into the pods may help to decrease patient wait times by providing more efficient cohorting situations.

Chapter 4

Modeling Methodology and Experimental Design

This chapter discusses the discrete event simulation used to model the current state of General Medicine patient flow and to predict the impact of the potential interventions. The simulation is implemented using the Simpy 3.0.10 in Python. It uses historical timestamps from a variety of data sources to create a model of the General Medicine patient flow in 2015. Section 4.1 discusses data sources and preparation. Section 4.2 describes the modeling framework. Section 4.3 defines the different interventions that are analyzed using this model and Section 4.4 discusses the experimental design used to evaluate these interventions.

4.1 Data Sources and Preparation

The simulation model uses data from a variety of sources at MGH, with the most important being the hospital-wide patient flow timestamps captured in EPIC, which also includes bed request and discharge events, infection precautions, and bed characteristics. The historical timestamps recorded in EPIC and other data sources are used to recreate the transitions of patients from one process or queue to another, and to determine the statuses of beds in the General Medicine units. The following sections provide descriptions of the primary data sources, the data cleaning processes, and a comprehensive list of the variables used to prepare the simulation.

4.1.1 PEPL Data

PEPL is a database that contains, among other things, the physician care team that was assigned to a patient at various times during her hospital stay. This information was used in combination with EPIC information to identify patients in the General Medicine population as described in the following section.

4.1.2 EPIC Data

EPIC is an electronic medical record system that has been implemented at MGH over the course of several years. Starting in July 2014, EPIC was in use for revenue management, but its clinical functions were not implemented until April 2016. The following sections describe the diverse types of data that were retrieved from EPIC.

Admission-Discharge-Transfer Data

The simulation is meant to capture the flow of all General Medicine patients who stayed on inpatient units in 2015. Since General Medicine patients can be non-regionalized, the procedure to identify these patients was non-trivial. The procedure used is described below.

1. Using the ADT dataset, identify and include all patient stays on fully regionalized General Medicine units[20].

2. Using the ADT dataset, identify and include all patient stays in the eight Bigelow 9 beds that were allocated to General Medicine[21].

3. Using the ADT dataset along with the PEPL dataset containing physician team information, identify patients that stayed on non-regionalized, non-cardiac floors[22], had 'Medicine' as their service description in ADT, and had their primary care team from PEPL either populated as one of the General Medicine teams or not populated at all. A list of all the General Medicine teams is presented in Appendix A.

[20]The fully regionalized General Medicine units are White 8, White 9, White 10, White 11, Bigelow 11, Ellison 12, Ellison 16, and Phillips 20.

[21]These eight beds are G0934A, G0936A, G0938A, G0940A, G0942A, G0944A, G0946A, and G0948A.

[22]These floors are Blake 6, Ellison 6, Ellison 7, Ellison 14, Ellison 19, Phillips 21, Phillips 22, White 6, White 7, Bigelow 14, Lunder 7, Lunder 8. Even though Ellison 19 has 10 allocated General Medicine beds, the beds themselves aren't fixed so the non-regionalized approach to patient identification was used.

4. Using the ADT dataset along with the PEPL dataset containing physician team information, identify patients that stayed on non-regionalized, cardiac floors[23], had 'Medicine' as their service description, and had their primary care team from PEPL populated as one of the General Medicine teams.

Once these filters were applied, 21,790 patient stays in beds were identified in scope. These stays in beds were then transformed to contain one entry for each patient stay in a unit as shown in Table 4.1. Each stay in a unit was treated independently for the purposes of the simulation. For instance, if a single patient originated in the ED then went from a unit in scope to an ICU and back to a unit in scope (whether it was the same or different unit), this would be treated as two different stays. The first stay would have the patient origin 'ED' and the second would have the origin 'ICU'. We would report on wait time for both these stays. Table 4.1 below illustrates this example. This procedure resulted in 16,957 patient encounters being in scope, with 237 of these being patients that were already present in the units of interest at the beginning of 2015 and therefore did not have their wait times calculated or reported. Throughout this thesis, the word patient is used to represent patient stays and the word discharge will be used to describe a patient leaving a unit, regardless of whether they leave the hospital or transfer internally.

Bed request information

EPIC captures a comprehensive picture of each bed request from beginning to end. We used this information to determine when the request was made, when Admitting assigned the patient to the unit, and when the patient became pending discharge. The bed requests that were considered for this simulation were those that were not cancelled and were linked to a patient's initial placement on a unit. Since each of these bed requests can contain many instances of each type of timestamp (e.g., created, requested, pre-assigned, assigned, approved), a condensed dataframe was created that showed only the first and last instance of each type of timestamp.

The bed request time was taken to be the first RequestedDTS when it was populated

[23]These floors are Ellison 8 and Ellison 11. For the cardiology floors, only patients with a General Medicine team populated in PEPL were selected because the true cardiology patients were likely to have their service populated as medicine.

Original Data			
Unit	Bed	Time In	Time Out
ED		1/22/15 6:58	1/22/15 19:01
MGH BIGELOW 11 MED	G1126 A	1/22/15 19:01	1/24/15 18:32
MGH BIGELOW 11 MED	G1138 A	1/24/15 18:32	1/28/15 17:08
MGH BIGELOW 11 MED	G1116 B	1/28/15 17:08	1/30/15 15:04
ICU		1/30/15 15:04	2/13/15 15:53
MGH WHITE 11 MEDICINE	W1122 A	2/13/15 15:53	2/16/15 16:18
MGH WHITE 11 MEDICINE	W1116 A	2/16/15 16:18	2/19/15 16:26
MGH WHITE 11 MEDICINE	W1122 A	2/19/15 16:27	2/20/15 14:04

Transformed Data		
Unit	Time In	Time Out
ED	1/22/15 6:58	1/22/15 19:01
MGH BIGELOW 11 MED	1/22/15 19:01	1/30/15 15:04
ICU	1/30/15 15:04	2/13/15 15:53
MGH WHITE 11 MED	2/13/15 15:53	2/20/15 14:04

Table 4.1: Illustration of patient encounter transformation

and the first EventTypeDTS on the rare occasion that RequestedDTS was not available. To determine when Admitting assigned the patient to the unit, we first identified all the Admitting users active in 2015 and created a subset of all the actions performed by these users. Then, for each encounter, we identified the first time that an Admitting user assigned the patient to the unit that they ultimately went to (if they were assigned to their ultimate unit, then another unit, and then again to their ultimate unit we took the timestamp from the last time they were assigned to their ultimate unit).

For pending discharge, we looked at the time that a discharge event that ended in a completion (rather than a cancellation) was first created for the patient. If a discharge was created and completed simultaneously, the pending discharge time would be the same as the leave time. The simulation does not change bed assignments once they are made and therefore we have adjusted any pending discharges that occurred in history more than twelve hours before the actual discharge to occur exactly twelve hours before discharge in the simulation input. These instances likely represent mistakes in the data, and our intervention is intended to keep patients in the simulation from waiting for an excessive amount of time for their beds to become ready once they are assigned.

Bed closure information

A report is available in EPIC that provides information on all bed closures recorded including the bed impacted, start time, end time, and reason. This information was processed to remove and combine overlapping closures as shown in Figure 4-1. Closure types were then categorized and all closures with categories of staffing or maintenance were replicated in the simulation.

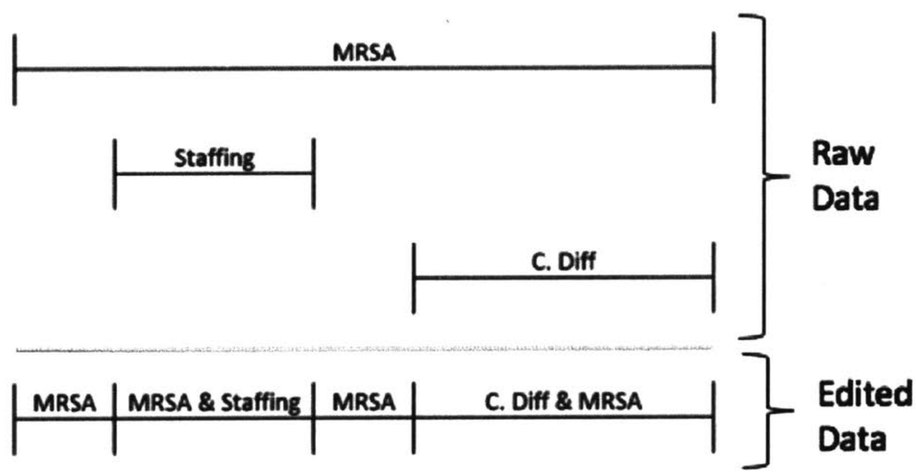

Figure 4-1: Illustration of closure overlap removal

Infection status information

A key input to the simulation is the isolation status of each patient over time. This information is gleaned from three different sources. The first is patient stays in private rooms. Since private rooms are so highly demanded, we assumed that patients only stay in them when absolutely required. Therefore, for any time that a patient was in a private room in history we considered her to be non-cohortable. This information is available through ADT. It is important to note that in the simulation only non-cohortable patients can stay in private rooms. The second source of isolation information is bed closures. If a patient's neighboring bed was closed for a patient-related reason, we consider that closure in determining the patient's isolation status. If the neighboring bed is closed for MRSA or VRE the patient is given the appropriate status. Any other patient-related closure results in the patient being classified as non-cohortable. The last source of isolation information is EPIC. In 2015 clinicians had the ability to alert other users to patients' isolation statuses by adding an infection flag in the system. We used EPIC data to see what infection precautions each patient had and when they were added and removed. We know based on interviews with clinicians that this dataset is not comprehensive for 2015 because providers did not have to enter clinical information in EPIC, however when populated it provides additional insight that cannot be found in the patient placement data. By combining this data with the stays in private rooms and the bed closures, we created a comprehensive isolation status history for each patient from the time of their bed request until discharge. In the process of validating the model we found that the isolation status data was not fully consistent. We identified over 750 instances where the isolation status that we had assigned to the patients was not consistent with their historical bed placement (e.g. we had designated a patient as non-cohortable at a time that he was sharing a semi-private room). Approximately 100 of these instances were sustained for more than four hours and we have corrected these by evaluating the details of the patients' placement before, during, and after the time where their isolation status was inconsistent. The details of the procedure used for these corrections are presented in Appendix E.

Bed cleaning information

EPIC provides a bed event report that tracks all the cleaning events associated with each bed. These data are gathered by the bed cleaners making phone calls when they start and end cleaning. As such, there are discrepancies in the data that we addressed by eliminating any sequence that does not end in clean, is totally overlapped by another sequence for the same bed, or has multiple starting timestamps. Based on expert input we then split then cleaning times into five different categories: regionalized private beds, non-regionalized private beds, regionalized semi-private beds, non-regionalized semi-private beds, and Phillips/Lunder[24] private beds. We then removed outliers from each category based on minimum and maximum estimates of realistic cleaning time. The cleaning times used in the simulation are sampled from the appropriate distribution based on the type of bed that is becoming available.

[24]Private rooms in Phillips and Lunder units have more furniture and take more time to clean than private rooms in other units.

4.1.3 EDIS and Perioperative Data

Two of the areas that interface with the General Medicine units had their own data systems in 2015: the ED and the Perioperative department. The ED used EDIS to track encounters within the ED before April, 2016. EDIS provides several data points that can also be found in other data systems, including the time of the bed request and the time of departure from the ED (which is not directly available in any other system but can be approximated from the arrival on the unit from ADT). When there were other sources for data, EDIS data was not used as its collection is more manual than the other systems and therefore thought to be less reliable. There were a limited number of data points that are available only through EDIS. The first is the ED triage assessment, which the ED began capturing in the 'AdmittingTeam' field of EDIS partway through 2015. The other EDIS data-point that we used was the 'PhysPassOffTime', which is manually populated to show that the ED physician performed handoff to the unit before the patient leaves the ED.

In 2015, the Perioperative department utilized a home-grown system called 'PRISM' to track patients in the OR environment, including the PACU. We were interested in the timestamps that indicated medical readiness of patients to enter their General Medicine beds. Based on interviews with Perioperative staff, we found three timestamps that were useful for this. The first is the 'ICUCare' timestamp. When present, this indicates that the patient was intentionally kept in the PACU overnight with ICU-level care. These patients were assumed to be ready at their 'patient_ready_to_depart_pacu' timestamp. For patients that were not intentionally kept in the PACU overnight, their medical readiness was determined to be the earlier of two hours after they entered the PACU per the timestamp 'patient_enters_pacu' and the time of 'patient_ready_to_depart_pacu'. The reason for this calculation was that staff indicated that the 'patient_ready_to_depart_pacu' was often recorded concurrently with the patient leaving.

4.1.4 List of Variables and Sources

Please see Table 4.2 for a comprehensive list of the key variables that were used to develop the input to the simulation and to calculate the historical metrics.

Source	Variable Name	Description
ADT	BedID	Identifier for the bed the patient stayed in
ADT	DepartmentDSC	Name of the unit the bed was located on
ADT	InDTS	Time that the patient arrived in the bed
ADT	OutDTS	Time that the patient left the bed
ADT	patServiceDSC	The clinical service responsible for the patient during the stay in bed
EDIS	AdmittingTeam	The triage level selected by the ED
EDIS	PhysPassOffTime	The timestamp associated with ED to unit physician passoff
EPIC - Bed cleaning	StatusDSC	The cleaning status of a bed at a particular timestamp
EPIC - Bed cleaning	StatusUpdateDTS	The timestamp associated with a cleaning event
EPIC - Bed closures	ADDED.AT	Timestamp when closure added
EPIC - Bed closures	BLOCK_TYPE	Determine closure type
EPIC - Bed closures	REMOVED.AT	Timestamp when closure removed
EPIC - Infection status	ADDED_DATTIME	Time precaution was added
EPIC - Infection status	EVENT_TYPE	Type of infection precaution
EPIC - Infection status	LEVEL	Determine if precaution is patient-level or encounter specific
EPIC - Infection status	REMOVED_DATETIME	Time precaution was removed
EPIC - Pending Events	EVENT_TYPE	Identifies the type of event in the sequence. Admissions == 1, Discharges ==2, Transfers == 3.
EPIC - Pending Events	CreatedPendingDischargeDTS	Used to determine when patient became pending discharge
EPIC - Pending Events	PreassignedDTS	Used to determine when admitting assigned the patient to the unit (compared with list of admitting users)
EPIC - Pending Events	RequestedDTS	Used to determine when a bed request was created
PEPL	PrimaryTeamNM	The name of the team that cared for the patient
PEPL	RunDTS	The time the team was recorded as caring for the patient
Peri-operative data	ICUCare	If populated indicates that patient was intentional PACU overnighter
Peri-operative data	patient_enters_pacu	Time that the patient enters the PACU
Peri-operative data	patient_ready_to_depart_pacu	Time the patient is ready to depart the PACU

Table 4.2: List of relevant variables with sources and descriptions

4.1.5 Determination of Triage Level

The practice of capturing the triage level for ED patients directly in the bed request was introduced in May of 2015. Prior to this date, the process of determining the appropriate placement for a patient was not standardized and the triage level was not captured in the data. Approximately 60% of ED patients admitted from May to the end of 2015 had the triage level populated directly in the bed request. In the 2015 data, the triage level is not captured for patients with origins other than the ED. We addressed this lack of data on patient triage level by using the following procedure to determine the triage level for 2015 patients:

- Find the patient's triage level as assigned by the ED (if available).
- Find the patient's triage level as suggested by his placement on a unit and team.
  - Assign all patients on non-regionalized units and Ellison 12 Level 2 status[25].
  - Assign patients on other units Level 2 status unless they are covered by the regionalized team (as shown in Tables 3.1 and 3.2), in which case assign Level 1.
- If the patient's ED triage level is Level 2, categorize the patient as Level 2.
- Otherwise, categorize the patient based on the placement triage[26].

Triage levels are used in the simulation to determine appropriate placement of patients. Because of the lack of available data, there are likely to be some inconsistencies between how we have categorized patients and how they were treated in history. Through the validation procedure described in Section 5.1, we show that these differences are not significant enough to impact the performance of the simulation.

4.2 Modeling Framework

4.2.1 General Patient Flow

The patient flow model starts with the arrival of each patient at his historical bed request time. Since the historical rules for prioritization are not standardized or documented, in the base scenario the model prioritizes patients for assignment based on the order that they were historically assigned. To accomplish this, each patient that is not already in a bed at the beginning of 2015 is given a ready-to-assign time. The ready-to-assign time must be after the patient's bed request and after the ready-to-assign times of all patients that were assigned before that patient in history. It will be the earliest possible time that meets both criteria.

[25]Although Ellison 12 had a regionalized HMG team for much of the year, interviews with the Admitting, DOM administrators, and HMG physicians revealed that primarily Level 2 patients were assigned to Ellison 12 in 2015.

[26]The reason that placement triage is used over ED triage is that there were instances in the data where patients were categorized as Level 1 by the ED and then placed in non-regionalized settings, indicating that the ED triage is not reliable if the two are in conflict.

Once a patient reaches his ready-to-assign time, the simulation will begin to look for a bed available for assignment that matches his needs in terms of gender, isolation status, and triage level. In the base scenario, the patient must go to the same unit to which he was historically assigned. This requirement is relaxed in the intervention scenarios. The simulation processes patients on a first-in-first-out basis, so if two patients have the same requirements, the one with the earlier ready-to-assign time will be assigned first. In addition to finding a bed with matching characteristics, there must also be a team with capacity to cover the bed. Each bed in the simulation has characteristics that say which team(s) can cover it for Level 1 patients and which teams(s) can cover it for Level 2 patients. Summaries of these teams for the base, redesign, and fully flexible scenarios are provided in Appendix B. The patient can only be assigned to the bed if one or more teams at the appropriate level are under its patient capacity limit.

The simulation includes an administrative delay to account for the fact that while assignments are made instantaneously when they become feasible in the simulation, in the real world it is a manual process that is not instantaneous. To determine a distribution for administrative delays, we looked at the historical timestamps for assignments and beds becoming available in a unit and then matched each assignment with the bed that became available before it occurred, starting from the latest assignment and the bed that became available immediately before that assignment. The time between the bed becoming available for assignment and the assignment occurring is the administrative delay. We used this method instead of finding the time that patient's assigned bed became available because bed swaps within the unit and team caps mean that the patient's assignment could be facilitated by a different bed becoming available. In the simulation, administrative delays are sampled from the historical data by patient source.

Once a patient is assigned to a bed, he must wait until becoming medically ready to continue the process. 99% of patients in the simulation become medically ready at the time of their bed request, but PACU patients may need to wait at this point. Once the patient is medically ready he starts waiting for his bed to become available to occupy. This means that the patient or closure that was occupying the bed previously must leave and the bed must be cleaned. This wait can range from zero, when the patient was assigned to a bed that was already ready, to over twelve hours, in the case where a pending discharge was entered far in advance of the actual discharge[27].

Cleaning times are sampled from the historical data. Once the bed becomes available the patient-wait-for-bed time ends. The patient then waits for a transfer processing time before he occupies his bed. This transfer processing time is also drawn from the historical data. For ED patients, we also record their physician handoff time. This is assumed to always follow the bed being ready and the amount of time it takes is sampled from the historical data[28]. The completion of the physician handoff marks the end of patient-wait-for-team.

The patient then stays in his bed until his historical pending discharge time (unless he is subject to isolation status changes or bed swaps, which are described in the following sections). At the time of the historical pending discharge, the patient's bed becomes available for assignment. Before the bed becomes available for assignment to patients, the simulation first checks whether there is a closure for staffing or maintenance that needs to be implemented on that unit. If there is such a waiting closure, the bed will be closed for the historical duration of the closure. It is important to note that the characteristics of the bed that is now available reflect the characteristics of the roommate (if the bed is in a semi-private room). This guarantees that the patient matched to the bed will be compatible with the existing roommate. If the neighboring bed is not assigned at the time that the bed becomes available for assignment, it can be assigned to any waiting patient.

In the base case, the patient then waits until his historical departure time from the unit before relinquishing his actual bed for cleaning. The decision was made to use the historical departure time instead of a departure time generated on the arrival time in the bed and the historical LOS because of the belief that small changes in a patient's intraday wait time will not affect his eventual discharge time. However, in the intervention scenarios, where patients are prioritized differently than they were in history, there is the opportunity for individual patients to experience larger differences in their wait times. This is particularly true for

[27]For the purposes of the simulation pending discharges were allowed to be a maximum of twelve hours in advance of the actual discharge. In the historical data there were rare occasions where the delay was greater.

[28]There were instances in the historical data where a physician handoff was recorded before the bed became ready, however these instances were rare and were excluded from both the distribution that was sampled in the simulation and the historical performance metrics.

ICU patients and hospital transfer patients who were historically de-prioritized. Based on insights from previous MGH-MIT collaboration work, in the intervention scenarios we have recalculated the departure time for ICU patients based on their simulation arrival time and historical LOS [6]. The general patient flow described in this section is also represented graphically in Appendix D.

4.2.2 Bed Swaps

As previously mentioned, cohorting requirements have an effect on the capacity of the General Medicine units. Therefore, in practice, patients are regularly swapped between beds on their unit to improve the cohorting situation. We have replicated this in the simulation using the following procedure.

1. When a semi-private bed becomes available for assignment (either because of a closure ending or a pending discharge) and there is no waiting patient that is appropriate for the bed, the simulation checks whether there is a patient on the unit that is currently in a room alone and matches the characteristics of the bed that is now available for assignment.

2. If such a patient exists, he will be swapped into the bed that the first patient is leaving and his bed will be made available to waiting patients instead of the original bed. This allows more flexibility in the patients that can be accommodated since the bed that becomes available to waiting patients is now suitable for any patient.

Bed swaps can also be initiated when patients change isolation status as discussed in the next section.

4.2.3 Isolation Status Changes

As described in Section 3.1.1, the possible isolation statuses in the simulation are clean, MRSA, VRE, MRSA & VRE, Influenza[29], and non-cohortable. While a patient with one

[29] As previously mentioned, Influenza patients can be cohorted only when they have the same strain and were infected around the same time. While making the corrections to isolation status discussed in Section 4.1.2, we found instances where patients we had designated as non-cohortable were cohorted together. Regardless of these patients' actual isolation issues, we have designated them as Influenza. The number of patients is very small (15 out of 16,957) and our assumption is that by giving them the status of Influenza, they will be cohorted together as in history.

of the first five statuses listed can be placed in a semi-private room with another patient with the same sex and status, non-cohortable patients must always be alone in a room. As described in Section 4.1, a patient's isolation status can change during her encounter. When this occurs, the following procedures are followed.

If the patient is in a semi-private room with a roommate (who will no longer match isolation status):

1. Check to see if there is a room on the unit that is available for the patient or her roommate to move into with another patient that they now match. If so, execute this move.

2. If there is no room with a matching patient, see if there is an empty room available to move the patient or roommate to. If so, execute this move.

3. If this does not work, leave the patient together with her roommate until another bed on the unit becomes available. Any time a bed becomes available check to see if either the patient or his roommate can be moved into it.

If the patient is in a semi-private room with no roommate:

1. Check to see if there is now an opportunity to cohort this patient with another patient that matches her new isolation status. If the new isolation status is non-cohortable, look for a private room for this patient.

2. If a cohorting situation or private room is not found, leave the patient in the semi-private room.

If the patient is in a private room and is no longer non-cohortable:

1. Check to see if the patient can now be cohorted with another patient on the unit.

2. If not, look for an empty semi-private room for this patient.

3. If neither 1 nor 2 are successful, leave the patient in the private room for the time being.

4.2.4 Bed Closures

There are three types of bed closures that occur in the simulation. The first is the bed closures that occur due to patient isolation status. Any time that a patient with an isolation status other than clean is in a semi-private room alone, the other bed will be closed based on the isolation status. In the case of MRSA, VRE, MRSA & VRE, or Influenza, the bed can be opened by finding an appropriate patient to cohort. For non-cohortable patients, the neighbor bed will remain closed as long as they are in the room with status non-cohortable.

The second type of closure is historical closures for non-patient related reasons. Beds can be closed for staffing when there are not enough nurses to cover all the beds in the unit. Beds are also periodically closed for maintenance (both planned and unplanned). These closures have been identified in the historical data and recreated in the simulation. At the historical time of a closure start, the simulation will look for a bed to close on the appropriate unit. If no bed is available at that time, the next bed to become available will be closed. The closure will last for the historical duration of the closure.

The third class of closures have been created to account for the variability in the number of non-regionalized beds occupied by General Medicine patients[30]. The input to the simulation contains the maximum number of non-regionalized beds that were ever occupied simultaneously by General Medicine patients and then closures are used to reduce the capacity to levels that mirror the realized capacity over time. These closures are given characteristics that match the sex and isolation status of the other (not General Medicine) patients that were present on the non-regionalized units at the time. This is intended to produce a base scenario that is as close as possible to the historical bed availability.

4.3 Interventions

The goal of this project is to evaluate the impact of interventions on the key metrics of patient wait. In addition to the redesign, the current state analysis suggests that a patient assignment algorithm, early team assignment, first-come-first-served prioritization, and in-

[30]This applies to the non-regionalized units that are assigned to other services and to Ellison 19, which General Medicine shares with Thoracic surgery. The number of beds on Ellison 19 available for medicine is approximately ten, but the actual occupancy varies significantly.

creased flexibility within pods could yield significant wait time reductions. This section develops these six interventions in detail.

The discussion focuses on describing the implementation of the interventions in the simulation. The goal of this project is to analyze the effectiveness of different interventions to present to hospital stakeholders who can determine the most cost-efficient and least disruptive ways to implement those operational changes that yield the most significant performance improvements in the simulation.

4.3.1 Patient Assignment Algorithm

As previously mentioned, the base scenario of the simulation was accomplished by prioritizing patients in the same order that they were historically assigned and assigning them to their historical units. This method is very effective at replicating past results as shown in Section 5.1, however it is not adequate for evaluating interventions that change the capacity of the system since the constraints that contributed to the historical prioritization are now relaxed. In addition to addressing the challenge of prioritizing patients in the intervention scenarios, we also sought to test an assignment algorithm that was standardized and clearly communicated the prioritization that was latent in the historical assignment method.

The algorithm implemented in the intervention scenarios is employed whenever a bed becomes available for assignment and does not need to be utilized for a closure, a patient who has changed isolation status, or a patient that has been early assigned to a team and is waiting for a bed assignment (this procedure is described in detail in Section 4.3.3). When such a bed becomes available, the list of patients who have requested a bed and are medically ready[31], but have not yet been assigned is searched sequentially for each category of patients shown in Table 4.3. Once a matching patient is found, the subsequent categories need not be searched. Within each category, patients are identified in a first-come-first-served manner, meaning that if there are two matching patients within the same category the one who had an earlier bed request will be selected. For a patient to be selected he must match the available bed's sex and isolation characteristics as described in Section 3.1.1. When a waiting patient

[31]PACU patients must be within an hour of becoming medically ready before they are considered for assignment in this algorithm.

is found, the bed is assigned to that patient. If there is no matching patient waiting at the time that the bed is ready to be assigned, the bed remains available until a new, matching, patient arrives or until the characteristics of the bed change (due to the neighbor leaving or changing isolation status) such that it now matches a waiting patient. To read Table 4.3, first select a bed type from the top and then look down the list at the patient categories. When a bed of the selected type becomes available, the categories are searched in order. For example, if a semi-private bed that can accommodate a Level 1 patient becomes available for assignment, the simulation first looks for an appropriate, cohortable (i.e., clean, MRSA, VRE, MRSA & VRE, or Influenza), Level 1 patient from the ED or PACU to assign to the bed. If there is such a patient waiting, he will be assigned to the bed even if there is another appropriate patient with a different origin or triage level who has been waiting longer. If there is no such patient waiting, the simulation will look for a patient in the next category (non-cohortable, Level 1, from the ED or PACU). Note that the only way that a non-cohortable patient can be placed in a semi-private bed is if the other bed is empty.

This algorithm was developed to capture the historical prioritization by source described in Section 3.2.1 and shown in the data in Figure 3-4. Additionally, this algorithm seeks to first assign Level 1 beds to Level 1 patients, since their placement requirements are more restrictive. Lastly, the algorithm prioritizes the placement of cohortable patients in semi-private rooms over non-cohortable patients. Note that a non-cohortable patient could only be placed in a semi-private room when both beds are currently unassigned. In this case, we would still prefer to place a cohortable patient in the room because there is the possibility that another patient could later be assigned to the neighboring bed.

4.3.2 The Redesign

The goal of the redesign is to improve the DOM in four key ways: increase capacity, improve regionalization of teams, create pods to increase effective capacity through more flexibility by covering teams, and establish an Oncology sub-unit to meet teaching goals. The increased capacity was implemented in the model by adding ten additional beds to Bigelow 9 in the bed input. Team changes in capacity and coverage were similarly implemented by updating the caps and the Level 1/2 coverage for each bed. A summary of the updated team coverage

	If the bed can accommodate a Level 1 patient:		If the bed can only accommodate a Level 2 patient:	
	If semi-private:	If Private:	If semi-private:	If Private:
Category 1	Cohortable Level 1 from ED or PACU	Non-cohortable Level 1 from ED or PACU	Cohortable Level 2 from ED or PACU	Non-cohortable Level 2 from ED or PACU
Category 2	Non-cohortable Level 1 from ED or PACU	Non-cohortable Level 1 from Front Door or Floor	Non-cohortable Level 2 from ED or PACU	Non-cohortable Level 2 from Front Door or Floor
Category 3	Cohortable Level 1 from Front Door or Floor	Non-cohortable Level 2 from ED or PACU	Cohortable Level 2 from Front Door or Floor	Non-cohortable Level 2 from ICU or Hospital Transfer
Category 4	Non-cohortable Level 1 from Front Door or Floor	Non-cohortable Level 1 from ICU or Hospital Transfer	Non-cohortable Level 2 from Front Door or Floor	
Category 5	Cohortable Level 2 from ED or PACU			
Category 6	Non-cohortable Level 2 from ED or PACU			
Category 7	Cohortable Level 1 from ICU or Hospital Transfer	Non-cohortable Level 2 from Front Door or Floor	Cohortable Level 2 from ICU or Hospital Transfer	
Category 8	Non-cohortable Level 1 from ICU or Hospital Transfer	Non-cohortable Level 2 from ICU or Hospital Transfer	Non-cohortable Level 2 from ICU or Hospital Transfer	
Category 9	Cohortable Level 2 from Front Door or Floor			
Category 10	Non-cohortable Level 2 from Front Door or Floor			
Category 11	Cohortable Level 2 from ICU or Hospital Transfer			
Category 12	Non-cohortable Level 2 from ICU or Hospital Transfer			

Table 4.3: Patient assignment algorithm

is shown in Table 4.4. In the table, the team identified as 'Level 1 Team 1' would be the preferred team to cover a Level 1 patient on that unit. If that team is not available the 'Level 1 Team 2' can cover the patient. With the implementation of the redesign, all regionalized units have the same teams covering Level 1 and Level 2 patients with the same priority. It is still the case that only Level 2 patients can be assigned to non-regionalized beds. Whereas in the base scenario the cap of the Green team was set to be high enough to cover all the available beds, in the redesign the cap is set at 31. This is done because of the increased regionalized capacity and because it is the intent of the DOM to enforce team caps more strictly in the future. In addition to the team constraints, non-regionalized beds are still constrained by their historical availability. The addition of a Flex team that covers beds at Level 1 on both units of the pods can also be seen in Table 4.4. The changes in

Unit	Beds	Level 1 Team 1	Level 1 Team 2	Level 2 Team 1	Level 2 Team 2	Level 2 Team 3	Level 2 Team 4
White 8	26	Bigelow A	Flex X	Bigelow A	Flex X		
Bigelow 9	18	Bigelow B	Flex X	Bigelow B	Flex X		
White 9	25	Bigelow C	Flex Y	Bigelow C	Flex Y		
White 10	20	Bigelow D	Flex Y	Bigelow D	Flex Y		
Bigelow 11	25	Bigelow E	Flex Z	Bigelow E	Flex Z		
White 11	24	Bigelow F	Flex Z	Bigelow F	Flex Z		
Ellison 12	36	Blue		Blue			
Ellison 16	36	Red	Oncology	Red	Oncology		
Ellison 19/ Phillips 20	30	Yellow		Yellow			
Non-regionalized	Varies			Green	Red	Blue	Yellow

Table 4.4: Post-redesign team coverage

team structure also support the creation of the Oncology sub-unit. By creating a dedicated Oncology team with capacity of 16 and the Red General Medicine team with capacity of 20, the unit is effectively divided between medicine and Oncology without having to dedicate specific beds. In the redesign scenario, only patients that are designated as Oncology[32] can be assigned to the Oncology team and the bed assignment mechanism within the simulation will assign Oncology patients to Ellison 16 Oncology beds whenever they are available. If an Oncology patient is assigned to a bed on Ellison 16 and covered by the Red team (because the Oncology team was full when he was assigned), he will continue to be covered by the Red team even when the Oncology team discharges a patient. If an Ellison 16 bed is not available, Oncology patients can be assigned to any other unit in the simulation.

[32]Oncology patients are defined in the historical data as those that were cared for by one of the Oncology Attending physicians while they were on a General Medicine floor. This information was obtained from H. Martinson in the MGPO Performance Analysis and Improvement group. These patients would have also been covered by one of the General Medicine teams (house staff or HMG) as their responding clinician, but their care was supervised by the Oncology specialty.

4.3.3 Early Team Assignment

The intervention of early team assignment seeks to reduce patient-wait-for-team for ED patients by pre-assigning them to a Flex team based on the high probability that a bed on the correct pod will become available during the day. To determine a reasonable number of patients to select for early team assignment, we looked at historical data on discharges and found that by pooling the beds in two units, we could be more confident that discharges would occur that would allow early assignment patients to reach beds in a timely fashion. Figure 4-2 shows the analysis performed for White 9 and White 10, which make up Pod 2 in the redesign. The 5th percentiles in the table indicate that while on most weekdays we could be confident that there would be at least one or two discharges from each unit, when pooled together we can expect four to seven discharges from each pod depending on the day of the week. This pooling effect and the fact that the Flex teams can cover beds on either unit form the impetus for exploring early team assignment.

The time for early assignment was set to be 8:00 am on every non-holiday weekday. At this time, the simulation reviews the list of non-oncology, clean, Level 1 patients originating in the ED that have submitted their bed request but have not yet been assigned a bed. The daily goal is to select six patients for early assignment (two for each pod). Once the simulation identifies the patients, their team is set to the appropriate Flex team. The time until physician handoff is then sampled from the historical data and once the handoff occurs the patient-wait-for-team stops while the patient-wait-for-bed continues until an appropriate bed is ready on the appropriate pod. When beds on the pods become available for assignment, early assignment patients are prioritized over other waiting patients. If fewer than six patients are identified at 8 am, new arrivals can be selected for early assignment up until 4 pm or until the target of six total early assignment patients is reached. As mentioned in Section 3.6.2, it is our hypothesis that the implementation of early team assignment will re-align the incentives for physicians regarding discharges and that discharges may occur earlier in the day. Unfortunately, since we did not have a detailed prediction about the effect of this incentive change, we did not include any acceleration of discharges in this intervention.

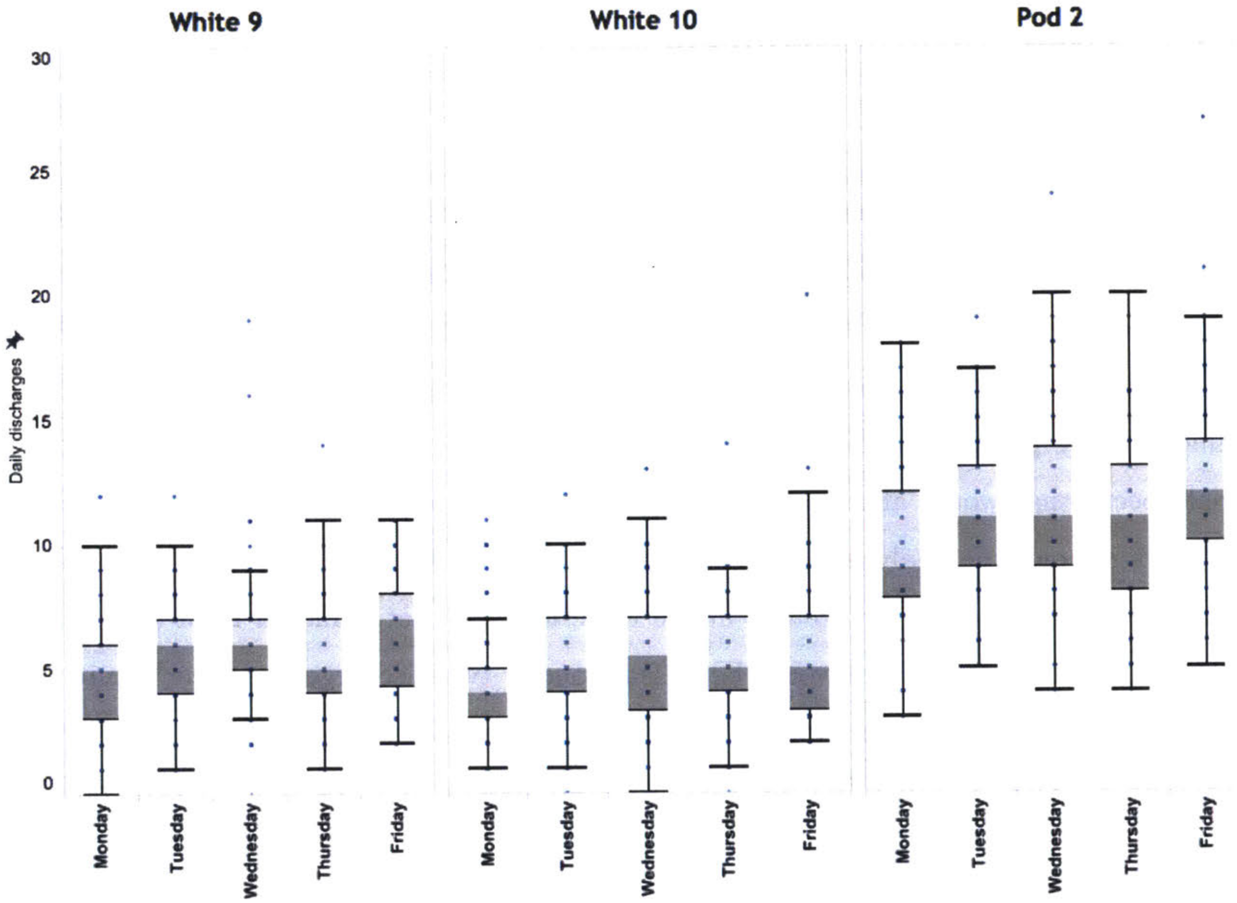

	Mon	Tue	Wed	Thu	Fri	Mon	Tue	Wed	Thu	Fri	Mon	Tue	Wed	Thu	Fri
5th %ile	1.000	2.000	2.000	2.000	3.000	2.000	1.000	2.000	2.000	2.000	4.000	6.000	7.000	5.000	6.450
10th %ile	2.000	3.200	3.000	2.000	3.000	2.000	2.000	2.000	2.000	2.000	5.400	6.400	7.300	6.000	7.000
Average	5.015	5.795	6.311	5.486	6.443	4.574	5.178	5.527	5.314	5.629	9.588	10.973	11.836	10.800	12.071

Figure 4-2: Historical weekday discharges from White 9, White 10, and Pod 2

Data sources: ADT (Tbl_Encounter_ADT_MGH) Time frame: Jan 1, 2015 – Dec 31 weekends and holidays excluded. Discharges include all patients who left the unit, whether they left the hospital or not.

4.3.4 First-come-first-served Prioritization

This intervention requires that patients be prioritized by their medical readiness and not their origin. This is accomplished by modifying the algorithm presented in Section 4.3.1 to only consider triage level and isolation status, not source, when assigning patients to beds (as shown in Table 4.5). For example if there were two clean males waiting for bed assignments where the one who arrived first originated from the ICU and was triage Level 1 and the second originated in the ED and was triage Level 2, the algorithm described in Section 4.3.1 would assign the bed to the ED patient, whereas the first-come-first-served algorithm would assign the bed to the ICU patient.

	If the bed can accommodate a Level 1 patient:		If the bed can only accommodate a Level 2 patient:	
	If semi-private:	If Private:	If semi-private:	If Private:
Category 1	Cohortable Level 1	Non-cohortable Level 1	Cohortable Level 2	Non-cohortable Level 2
Category 2	Non-cohortable Level 1	Non-cohortable Level 2	Non-cohortable Level 2	
Category 3	Cohortable Level 2			
Category 4	Non-cohortable Level 2			

Table 4.5: First-come-first-served algorithm for patient assignment

4.3.5 Increased Flexibility within Pods

This intervention has two levels. The first is allowing patients that are assigned to the Flex team for each pod to be swapped across units for better cohorting. In future discussion, this level of flexibility will be referred to as 'swaps within pods'. The second level relaxes the requirement that only Flex teams cover patients on both units and allows all teams in the pod to cover patients on both units. This sets the stage to allow all patients to transfer across units within the pod for optimal cohorting. This scenario will be referred to as 'full flexibility' for the remainder of this thesis. It is important to note that the current concept of the redesign maintains the nursing team structure that is associated with the unit (not the pod) and therefore proposals to swap patients between units in a pod will affect nursing coverage and these consequences should be considered prior to implementation.

4.3.6 Eliminate Staffing Closures

This intervention is very straightforward to implement in the simulation and merely requires that staffing closures not be introduced.

4.4 Design Strategy and Combinations of Interventions

The strategy used to evaluate these interventions is to limit the number of combinations by focusing on the most feasible. The first scenario allows us to test the effectiveness of the patient assignment algorithm with the pre-redesign capacity and therefore compare its performance to the base scenario. The second is the evaluation of the redesign as it was originally conceived (with the patient assignment algorithm in place). Interventions are then layered in the order seen to be the most feasible. Table 4.6 provides an overview of the scenarios and introduces the naming conventions that will be used for the remainder of the document.

Scenario Name	Historic Unit and Order	Patient assignment algorithm	Redesign	Swaps Within Pods	Early Team Assignment	Full Flexibility	FCFS Patient prioritization	Eliminate staffing closures
Base simulation	X							
Patient assignment algorithm		X						
Redesign		X	X					
Redesign, swaps within pods		X	X	X				
Redesign, swaps within pods, early assignment		X	X	X	X			
Redesign, fully flexible, early assignment		X	X		X	X		
Redesign, fully flexible, early assignment, FCFS			X		X	X	X	
Redesign, fully flexible, early assignment, FCFS, staffing			X		X	X	X	X

Table 4.6: Intervention combination strategy

Chapter 5

Results

This chapter presents the results for both the base scenario model and the different operational interventions. Section 5.1 explains the model validation that was performed to ensure that the baseline results accurately reflect the current realities at the hospital. Sections 5.2 through 5.4 present results and analysis of the interventions introduced in Section 4.3. The simulation was implemented in the SimPy 3.0.10 environment of Python 2.7 and tested on an Intel 2.2 GHz Macbook Air with 4 GB of RAM.

5.1 Model Validation

The validation of the model is performed by comparing the wait times generated in the simulation's base case to the historical actual wait times. As discussed in Section 4.2, in the base case the patients are prioritized in the same order that they were historically assigned and are constrained to be assigned to the same unit that they were in history. As described in Chapter 4, there are several elements in the simulation that are sampled from distributions and therefore subject to random variation; thus, we produced 20 replications of the simulation results (this was done for the base scenario and the intervention scenarios described in the following sections).

The main basis of comparison used is the difference between the historical mean patient-wait-for-bed times and the simulation-produced mean patient-wait-for-bed times. For ED patients, we also compared patient-wait-for-team times. All differences are expressed as historical wait minus simulation wait, hence a positive number indicates that patients wait a shorter amount of time in the simulation than in history (and vice-versa for a negative number). For each subset of interest we found this difference and then calculated a 95% confidence interval (CI) for the true difference in means. If this interval contains zero we can assume that the historical and simulation data have the same mean. The procedure used for this calculation is described in Section 10-3.4 of Montgomery and Runger's book Applied Statistics and Probability for Engineers [19].

In addition to the difference of means, we compared the 5th percentiles, 25th percentiles, medians, 75th percentiles, and 95th percentiles of the historical and simulated wait times for the subsets of interest. Comparisons of the quantiles were not subject to the same statistical testing as the means, rather we considered whether the results were roughly consistent across the distribution. Using these methods, we find that the base simulation produces results that closely reflect history across all subsets. Detailed results for the total General Medicine patient population and each subset are discussed in the following sections.

5.1.1 Overall Validation

When the total General Medicine patient population is considered, the mean patient-wait-for-bed time from the simulation results is 442 minutes (7.37 hours) compared to the historical mean wait of 446 minutes (7.43 hours), a difference of 4 minutes 95% CI [-8, 17 minutes]. Comparison of the quantiles reveals slightly larger differences than the means, but no difference that exceeds 12 minutes in magnitude. Figure 5-1 shows these quantiles in hours. The differences in minutes are as follows: 5th percentile -2 minutes, 25th percentile -3 minutes, median 11 minutes, 75th percentile 8 minutes, 95th percentile 3 minutes.

Other metrics that suggest that the base simulation functions in a way that is consistent with history are the number of bed swaps and the average non-regionalized census. We found that the average number of beds each patient stayed in across all 20 runs of the base simulation was 1.34. This is only 4% higher than the historical average of 1.29. It is logical that this number is slightly higher since the simulation automatically completes all feasible and productive bed swaps, whereas historically the execution of these swaps relied on an individual in Admitting or on the unit noticing the opportunity. The average non-regionalized census was also found to be consistent with history at 42.9 patients in the base simulation versus 43.0 patients in history. The small variation can be attributed to slight differences in processing times and other dimensions that are sampled.

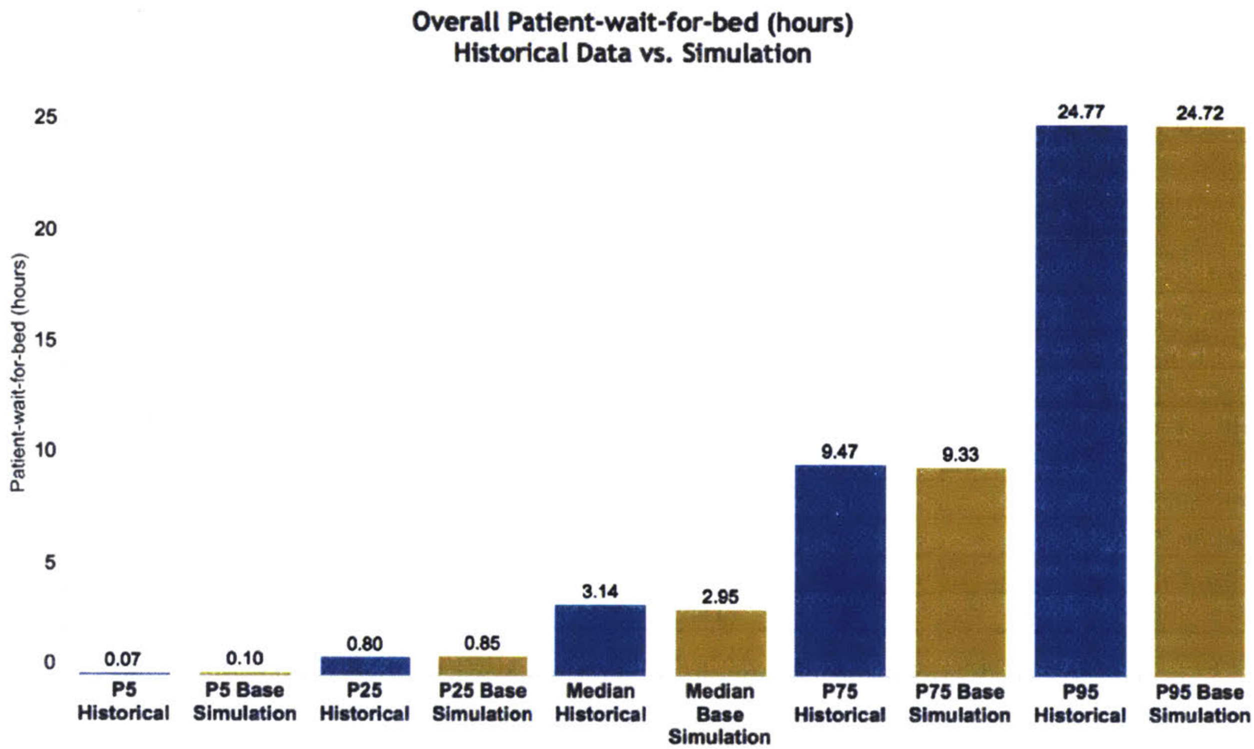

Figure 5-1: Comparison of patient-wait-for-bed quantiles, validation
Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

5.1.2 Validation by Patient Origin

Approximately 80% of General Medicine patients originate from the ED. We found that for these patients the mean patient-wait-for-bed time from the simulation results is 340 minutes (5.67 hours) compared to the historical mean wait of 342 minutes (5.70 hours), a difference of 2 minutes 95% CI [-5, 9 minutes]. Comparison of the quantiles reveals slightly larger differences than the means, but no difference that exceeds 10 minutes in magnitude. Figure 5-2 shows these quantiles in hours. The differences in minutes are as follows: 5th percentile -2 minutes, 25th percentile -4 minutes, median 10 minutes, 75th percentile 7 minutes, 95th percentile 0 minutes. Table 5.1 provides an overview of the results for patients by origin including the number of patients originating from each location. Origins with fewer patients and larger mean wait times tend to have larger differences of means (up to 29 minutes for hospital transfer patients), but based on the confidence intervals, validation is achieved for all origins.

	ED	Floor	Front Door	Hospital Transfer	ICU	PACU
Number of Observations	13348	365	1502	344	1007	154
Hist. Mean	342	552	521	1239	1460	69
Sim. Mean	340	537	512	1210	1445	66
Difference of means	2	16	9	29	14	3
CI Lower Bound	-5	-131	-37	-173	-116	-18
CI Upper Bound	9	162	56	231	145	24
Hist. P5	5	0	2	5	27	0
Sim. P5	7	1	13	7	31	0
Hist. P25	43	10	80	119	243	0
Sim. P25	47	11	78	93	231	0
Hist. Median	171	107	222	380	664	0
Sim. Median	161	96	201	377	623	0
Hist. P75	526	308	418	1541	1855	100
Sim. P75	519	268	420	1495	1830	97
Hist. P95	1167	3031	2354	4861	5057	277
Sim. P95	1167	2983	2345	4888	5097	324

Table 5.1: Key validation metrics on patient-wait for bed by patient origin
Note: All times given in minutes (24 hours = 1440 minutes)
Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

For ED patients, it is also important to validate the patient-wait-for-team metric since this is used to evaluate the early team assignment intervention. We found that the mean historical patient-wait-for-team time was 420 minutes (7.00 hours) and the mean simulation wait was 424 minutes (7.07 hours), a difference of -4 minutes 95% CI [-11, 2 minutes]. Since the confidence interval includes zero, the simulation validates on this dimension in addition to patient-wait-for-bed.

5.1.3 Validation by Triage Level

Level 1 patients were found to have a mean patient-wait-for-bed of 502 minutes (8.37 hours) in the simulation compared to the historical mean wait of 507 minutes (8.45 hours), a difference of 5 minutes 95% CI [-14, 24 minutes]. Level 2 patients were found to have a mean patient-wait-for-bed of 359 minutes (5.98 hours) in the simulation compared to the historical mean

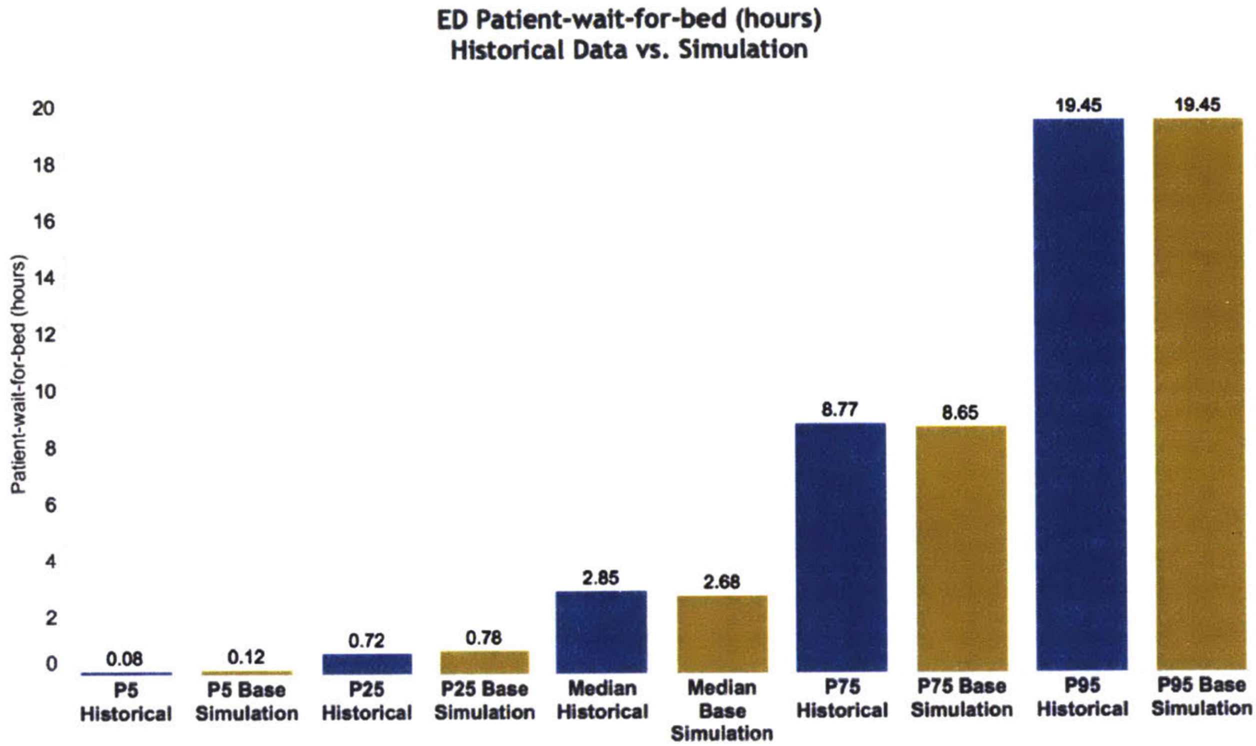

Figure 5-2: Comparison of patient-wait-for-bed quantiles, validation, ED patients
Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

wait of 363 minutes (6.03 hours), a difference of 4 minutes 95% CI [-9, 17 minutes]. These results suggest that the simulation accurately reproduces the wait times by triage level, and the quantile results shown in Figure 5-3 do not indicate any major discrepancies in the distributions of the wait times.

5.1.4 Validation by Unit

Since patients in the base scenario were constrained to go to the unit that they historically occupied, it is important to confirm validation for each unit. Table 5.2 provides an overview of each unit and indicates that the mean simulation patient-wait-for-bed wait time is consistent with history for each unit. Note that for the validation (and the interventions) all non-regionalized beds are pooled together even though patients in history occupied beds on many different units.

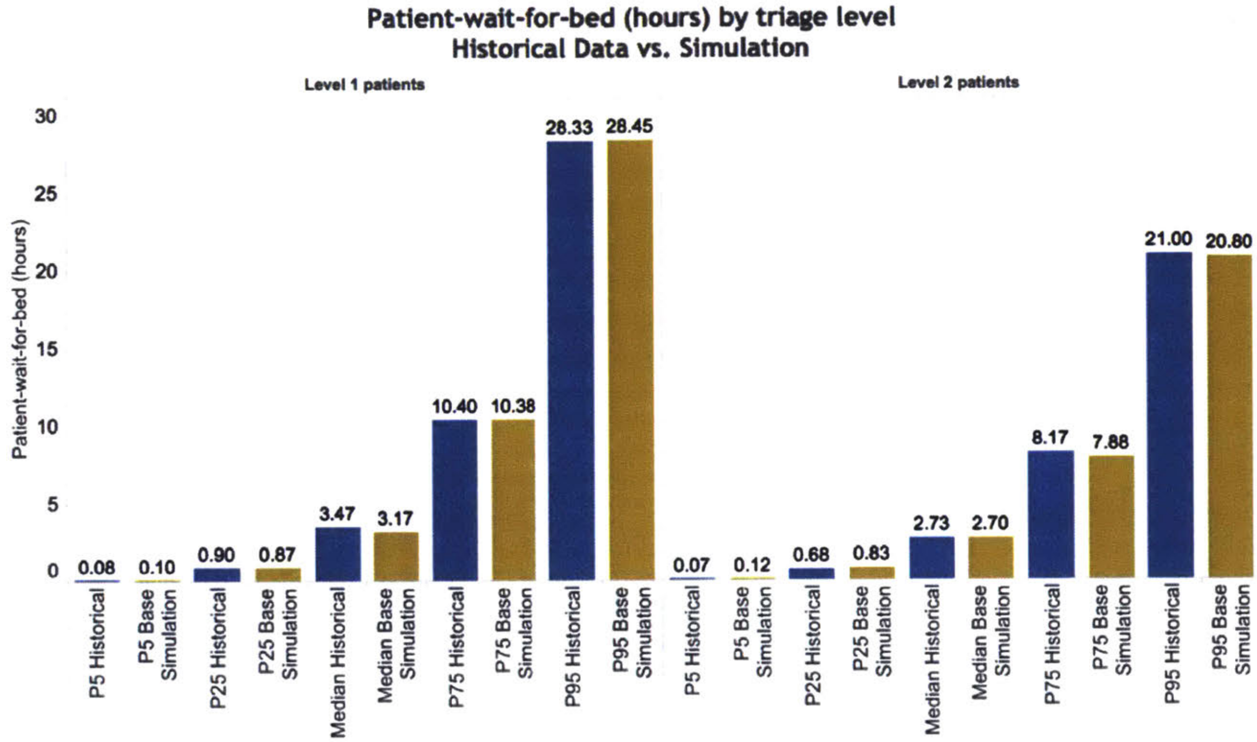

Figure 5-3: Comparison of patient-wait-for-bed quantiles, validation, separated by triage level

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

	Ellison 12	Ellison 16	Ellison 19	Bigelow 9	Bigelow 11	Phillips 20	White 8	White 9	White 10	White 11	Non-regionalized
Number of Observations	2226	2034	613	458	1567	947	1561	1451	1230	1435	3198
Hist. Mean	344	422	385	460	478	873	502	466	505	449	341
Sim. Mean	346	389	399	460	470	852	492	483	510	455	338
Difference of means	-2	32	-14	0	8	20	11	-17	-5	-5	3
CI Lower Bound	-21	-4	-71	-55	-41	-65	-36	-56	-64	-51	-16
CI Upper Bound	17	68	44	56	58	106	57	22	53	41	23
Hist. P5	3	6	4	7	5	8	6	6	4	4	4
Sim. P5	5	4	10	6	6	10	7	6	5	5	12
Hist. P25	42	50	45	68	47	155	50	52	53	41	38
Sim. P25	47	37	65	56	48	137	58	54	44	45	56
Hist. Median	152	177	158	235	185	459	229	195	204	184	156
Sim. Median	163	128	171	214	161	440	202	209	192	173	158
Hist. P75	497	447	422	703	617	1079	640	630	646	542	427
Sim. P75	487	412	437	717	589	1072	623	671	638	537	434
Hist. P95	1213	1498	1350	1390	1593	2859	1652	1560	1638	1497	1233
Sim. P95	1222	1452	1366	1460	1593	2838	1688	1570	1649	1567	1187

Table 5.2: Key validation metrics on patient-wait for bed by unit
Note: All times given in minutes (24 hours = 1440 minutes)

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 - Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

5.1.5 Validation by Isolation Status

Lastly, we viewed the validation statistics by patient isolation status. Table 5.3 provides an overview of each isolation status and indicates that the mean simulation patient-wait-for-bed wait time is consistent with history for each status.

	PACU	Clean	Influenza	MRSA	MRSA & VRE	Non-cohortable	VRE
Number of Observations	154	11537	15	376	104	3905	783
Hist. Mean	69	370	767	528	610	649	494
Sim. Mean	66	368	726	510	634	639	476
Difference of means	3	1	41	18	-24	10	17
CI Lower Bound	-18	-10	-354	-86	-200	-26	-59
CI Upper Bound	24	13	436	122	153	47	94
Hist. P5	0	4	78	6	5	4	5
Sim. P5	0	6	2	4	8	10	6
Hist. P25	0	41	202	55	61	91	36
Sim. P25	0	45	136	40	74	86	36
Hist. Median	0	160	511	210	313	304	165
Sim. Median	0	151	464	179	314	288	157
Hist. P75	100	461	1159	636	743	896	541
Sim. P75	97	458	1127	592	942	884	514
Hist. P95	277	1285	2002	1705	1912	1902	1643
Sim. P95	324	1298	2309	1807	2198	1884	1598

Table 5.3: Key validation metrics on patient-wait for bed by isolation status
Note: All times given in minutes (24 hours = 1440 minutes)
Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

5.2 Results

We first present an overall look at how the interventions impacted patient-wait-for-bed for all patients, Level 1 ED patients, and Level 2 ED patients before more closely examining the effects of each intervention on relevant populations in the following section. Figure 5-4 shows the mean patient-wait-for-bed time for all patients, Level 1 ED patients, and Level 2 ED patients under each scenario. Important observations from this Figure include the 9% decrease in overall average patient-wait-for-bed with the implementation of the patient assignment algorithm (from 7.36 hours in the base simulation to 6.67 hours with the patient assignment algorithm). With this same intervention however, ED Level 1 patients actually wait longer on average than in the base scenario. The detailed effects of this intervention are discussed in Section 5.3.1. The implementation of the redesign in addition to the patient assignment algorithm results in a further 31% decrease in mean overall patient-wait-for bed. In this case the implementation of the intervention results in decreases in average waits for both Level 1 ED patients and Level 2 ED patients. The detailed effects of this intervention are discussed in Section 5.3.2. Figure 5-4 also reveals that the remaining interventions have only modest effects on the overall mean patient-wait-for-bed time and that the results on the means for the different populations (Level 1 ED patients and Level 2 ED) are not always consistent (e.g. the implementation of early team assignment causes an increase in the mean wait time for Level 1 ED patients and a decrease for Level 2 ED patients). Table 5.4 provides a summary of the results for the overall population under each combination of interventions. Explanations for these varying effects along with detailed analysis of affected populations are presented in the following sections.

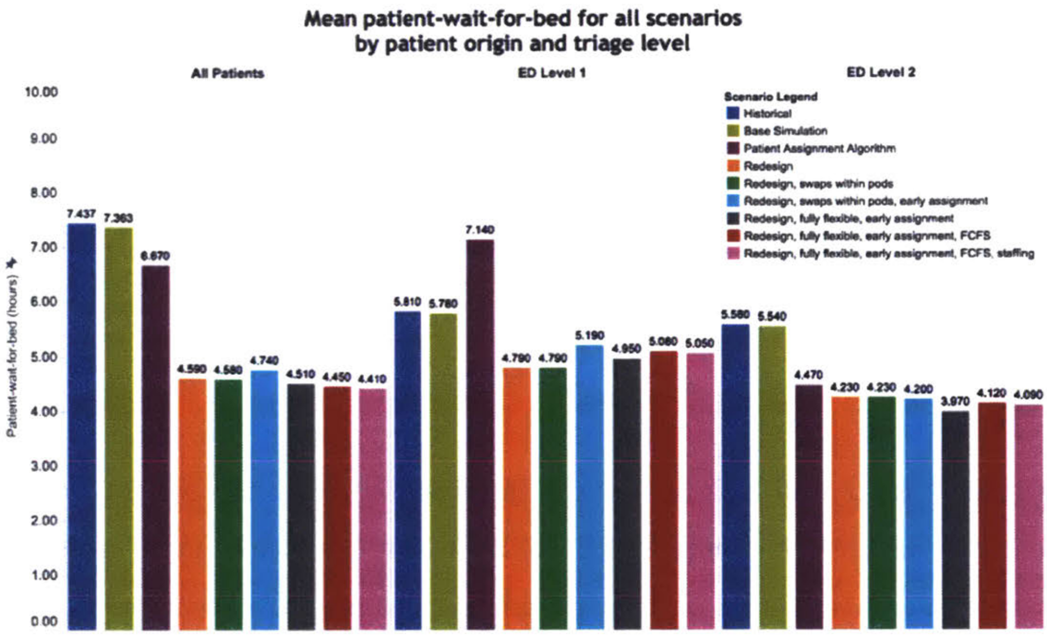

Figure 5-4: Comparison of patient-wait-for-bed means, simulation scenarios and historical data

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

In addition to analyzing the effects of the interventions on patient-wait-for-bed time for the populations of interest, in Figures 5-5, 5-6, and 5-7 we present the effects on the 5th percentile, 25th percentile, median, 75th percentile, and 95th percentile of patient-wait-for-bed. The effects on the quantiles are generally consistent with the direction and magnitude of the effects presented in Figure 5-4.

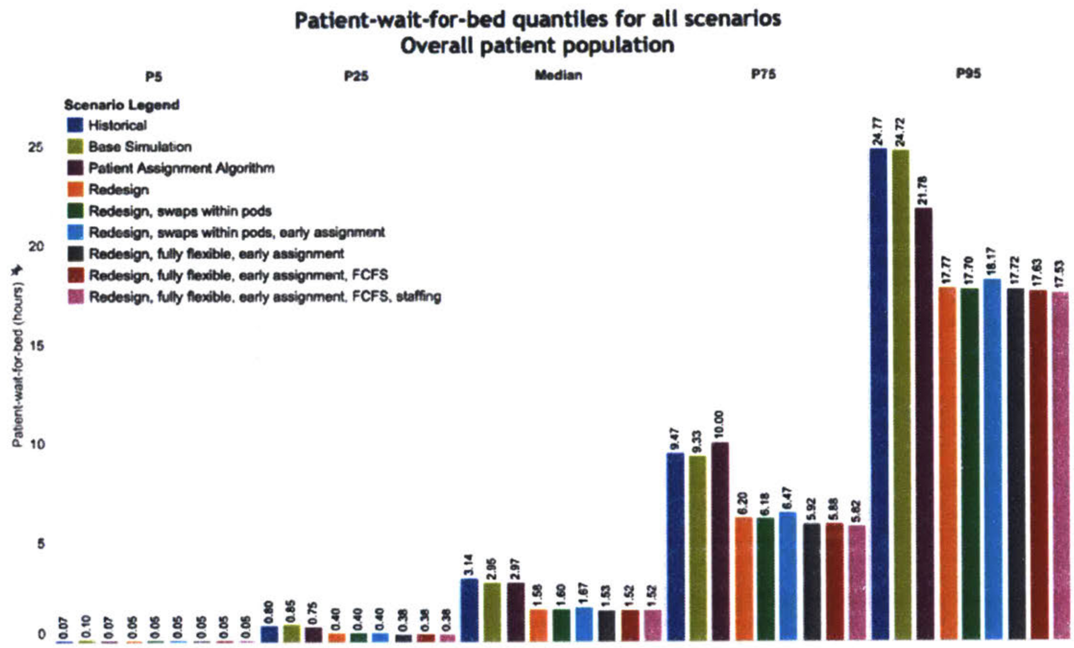

Figure 5-5: Comparison of patient-wait-for-bed quantiles for all patients, simulation scenarios and historical data

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

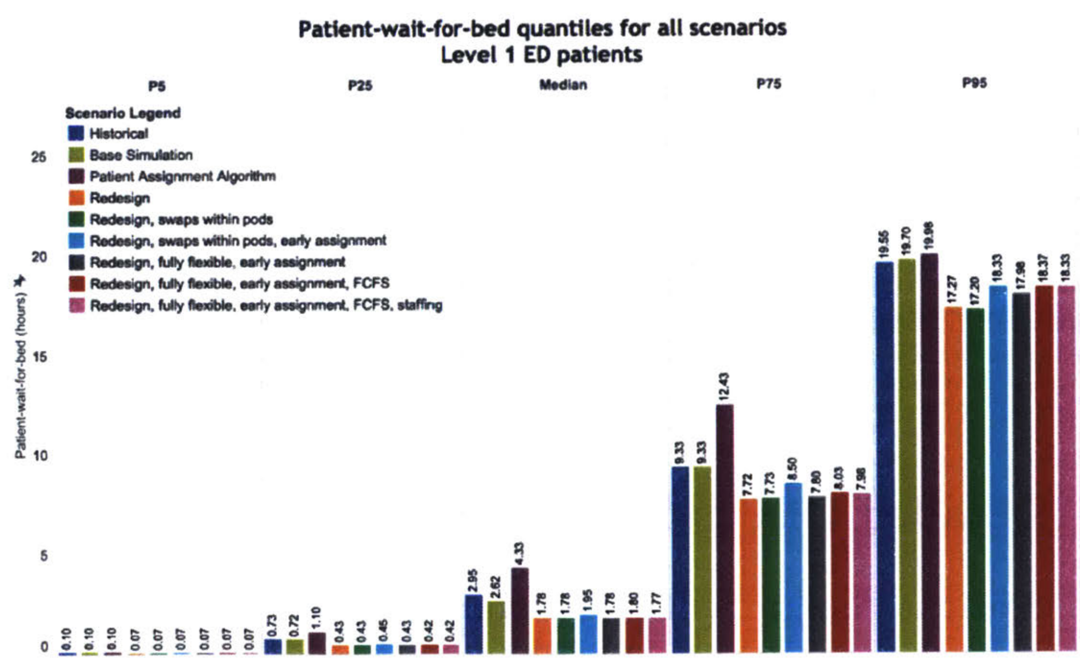

Figure 5-6: Comparison of patient-wait-for-bed quantiles for ED Level 1 patients, simulation scenarios and historical data

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

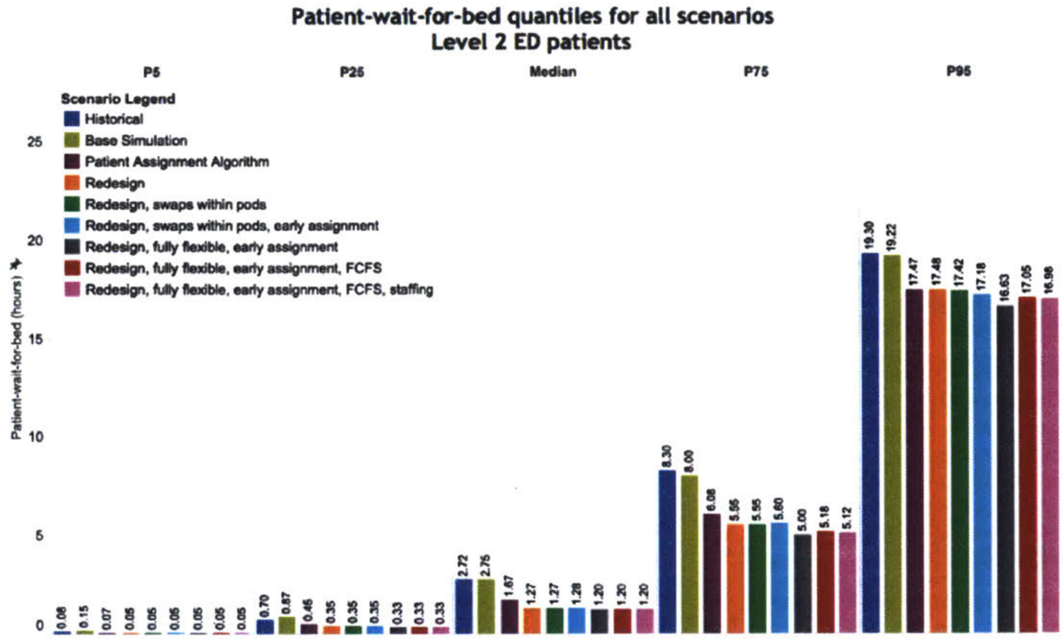

Figure 5-7: Comparison of patient-wait-for-bed quantiles for ED Level 2 patients, simulation scenarios and historical data

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

	Historical	Base Simulation	Patient Assignment Algorithm	Redesign	Redesign, swaps within pods	Redesign, swaps within pods, early assignment	Redesign, fully flexible, early assignment	Redesign, fully flexible, early assignment, FCFS	Redesign, fully flexible, early assignment, FCFS, staffing
Number of Records	16,720	334,400	334,400	334,400	334,400	334,400	334,400	334,400	334,400
Mean Patient-wait-for-bed	7.44	7.36	6.67	4.59	4.58	4.74	4.51	4.45	4.41
P5 Patient-wait-for-bed	0.07	0.10	0.07	0.05	0.05	0.05	0.05	0.05	0.05
P25 Patient-wait-for-bed	0.80	0.85	0.75	0.40	0.40	0.40	0.38	0.38	0.38
Median Patient-wait-for-bed	3.14	2.95	2.97	1.58	1.60	1.67	1.53	1.52	1.52
P75 Patient-wait-for-bed	9.47	9.33	10.00	6.20	6.18	6.47	5.92	5.88	5.82
P95 Patient-wait-for-bed	24.77	24.72	21.78	17.77	17.70	18.17	17.72	17.63	17.53

Table 5.4: Table comparing of patient-wait-for-bed in hours for all patients, simulation scenarios and historical data

Data sources: Simulation results, ADT (Tbl_Encounter_ADT_MGH), EPIC (Tbl_Clinical_BedEventData_MGH, MGH_BED_BLOCKING) Time frame: Jan 1, 2015 – Dec 31, Filtered for General Medicine patients with bed requests created and completed in 2015.

5.3 Detailed Discussion of Results

5.3.1 Patient Assignment Algorithm

The first intervention applied to the simulation is the patient assignment algorithm. This algorithm prioritizes patients for assignment based on their bed request time, triage level, isolation status, and origin as described in Section 4.3.1. Unlike the base simulation, patients are no longer prioritized in the order that they were in history or required to go to their historical unit. We find that when this algorithm is applied in the simulation the overall mean patient-wait-for-bed time falls from 7.36 hours to 6.67 hours, a change of 0.69 hours 95% CI [0.64, 0.75 hours]. An examination of the quantiles shown in Figure 5-8 reveals decreases occur at the 5th, 25th and 95th percentiles.

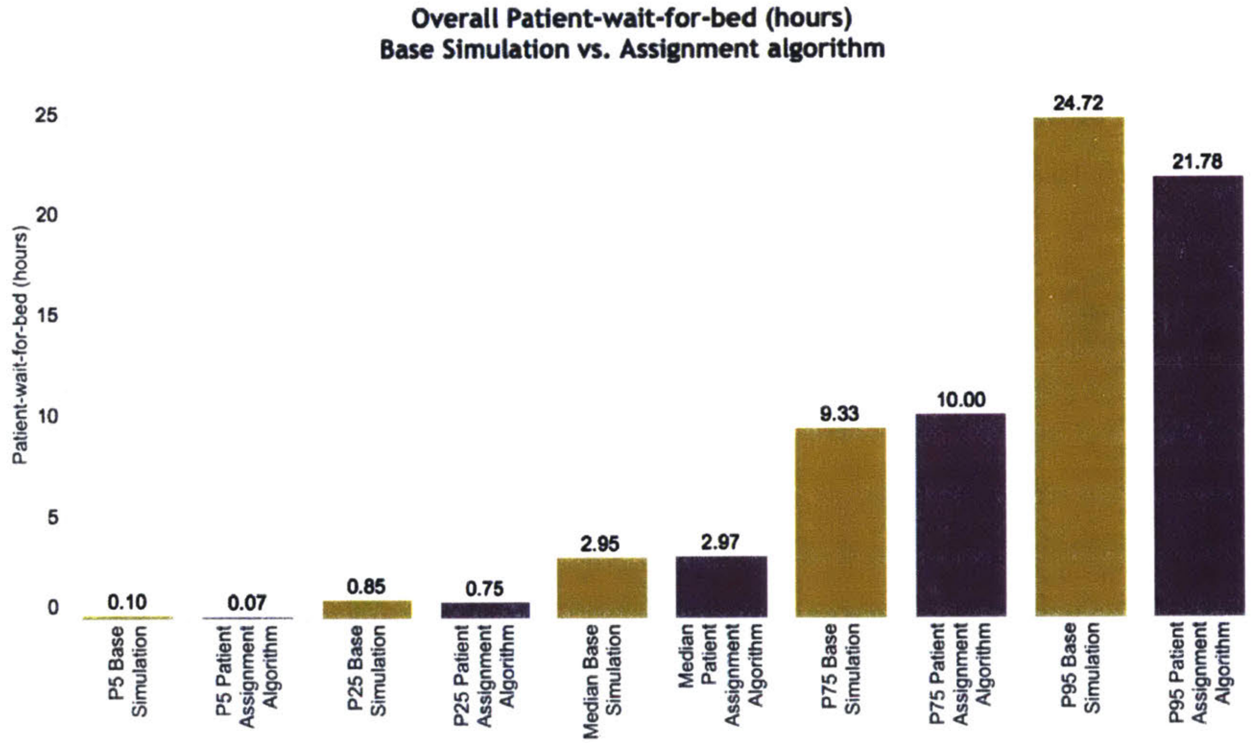

Figure 5-8: Comparison of patient-wait-for-bed quantiles, base simulation vs. assignment algorithm

Data sources: Simulation results.

These differences can be better understood by examining the results by source shown in Figure 5-9, which reveals that despite continuing to prioritize patients by origin, the assignment algorithm results in less disparity in wait time by origin than the base simulation. This is most evident in the modest increase in patient-wait-for-bed for ED patients (from a mean of 5.67 hours to 5.93 hours) and the large decreases for hospital transfer and ICU patients (from 20.17 to 13.57 and 24.09 to 13.59 hours respectively). These changes occur despite Level 1 and Level 2 ED patients being prioritized over Level 1 ICU and hospital transfer patients for Level 1 beds (see Table 4.3 for an overview of the assignment algorithm). In fact, we performed experimentation with the algorithm in an attempt to more closely match the historical and base simulation waits by origin and found that the only way to achieve this was to make ICU and hospital transfer patients ineligible to be assigned to a bed at the time of their bed requests, essentially introducing an artificial delay for these patients [17]. This suggests that in practice, available beds were not assigned to appropriate, waiting ICU and hospital transfer patients, even when there were no other matching patients waiting. For the evaluation of the other interventions we have used the algorithm presented in Section 4.4.1 (rather than one that artificially delays ICU and hospital transfer patients) because it provides an overall decrease in average patient-wait-for-bed time and because we believe that this benefit outweighs the slight increase in average wait for ED patients.

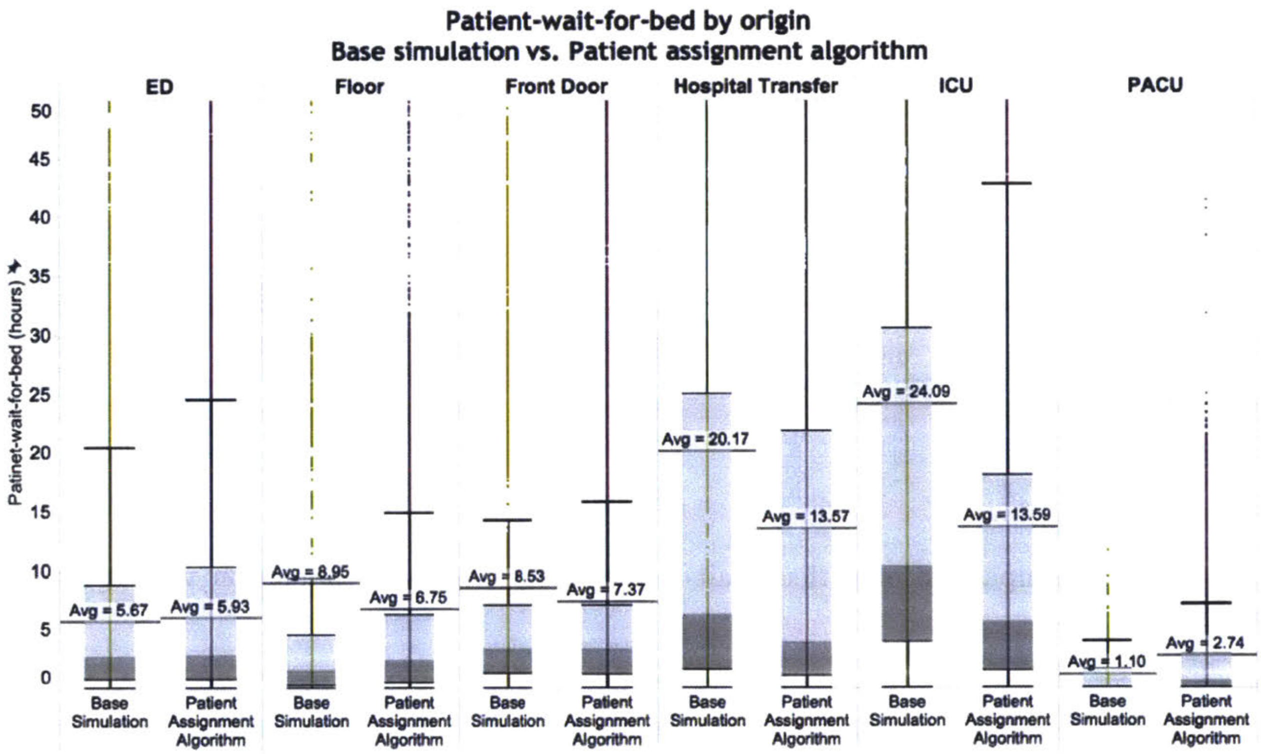

Figure 5-9: Comparison of patient-wait-for-bed times by origin, base simulation vs. assignment algorithm

Note: Number of observations across 20 runs is as follows: ED = 266,960, Floor = 7300, Front door = 30,040, Hospital transfer = 6,880, ICU = 20,140, PACU = 3080. Maximum values are truncated on this graph for legibility. Maximum values for the base scenario (in hours) are as follows: ED 68.45, Floor 220.45, Front door 152.92, Hospital transfer 214.98, ICU 288.1, PACU 11.68. Maximum values for the assignment algorithm scenario (in hours) are as follows: ED 94.27, Floor 165.48, Front door 148.48, hospital transfer 219.63, ICU 238.13, PACU 41.43. Data sources: Simulation results.

5.3.2 DOM Redesign

The results show that the implementation of the DOM redesign has a large effect on overall patient wait time. For all patients, average patient-wait-for-bed time decreases from 6.67

hours with the patient assignment algorithm to 4.59 hours once the redesign is implemented. This decrease of 2.08 hours (95% CI [2.02, 2.12 hours]) represents a change of 31% from the pre-redesign results. In addition to improving the average patient-wait-for-bed, all quantiles decrease as shown in Figure 5-10. Another effect of the redesign is to decrease the average census of non-regionalized patients from 42.9 in the base case to 38.4 in the redesign.

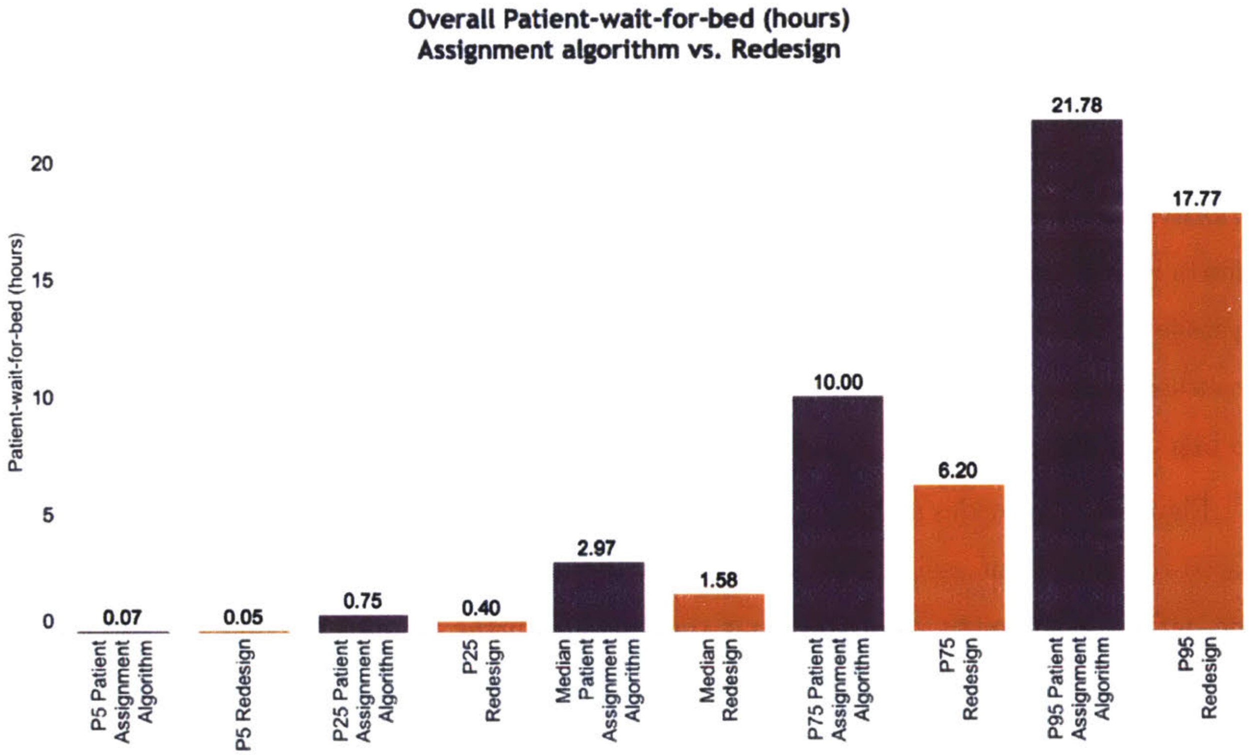

Figure 5-10: Comparison of patient-wait-for-bed quantiles, assignment algorithm vs. redesign

Data sources: Simulation results.

Digging deeper into the effects of the redesign on different patient groups reveals that it provides the most benefit for Level 1 patients and for Oncology patients. Figure 5-11 shows that while Level 1 patients experience a decrease in average patient-wait-for-bed of 40% from 8.09 to 4.86 hours, Level 2 patient-wait-for bed decreased by only 11% from 4.73 to 4.23 hours. Similarly, Oncology patient-wait-for-bed decreased 52% from 6.91 to 3.29 hours while non-Oncology patient wait improved only 29% from 6.64 to 4.73 hours. These effects are well explained by the structure of the redesign.

The redesign adds Level 1 bed capacity by increasing the number of beds on Bigelow 9, covering all beds on house staff units with regionalized teams, and changing Ellison 12 from Level 2 only to Level 1. At the same time, Level 2 dedicated capacity is decreased by the aforementioned changes and the patient cap is lowered on the Green team in the redesign scenario (see Section 4.4.2 and Appendix B for details of team alignment in the redesign). The result is that Level 1 patients experience a dramatic decrease in wait time while Level 2 patients see only a modest one.

The redesign also creates the dedicated sub-unit of 16 beds for Oncology patients on Ellison 16. In an effort to fully utilize this capacity, when these beds are available and an Oncology patient requests a bed she is assigned right away, regardless of her triage level or priority. This results in a shorter average wait time for the Oncology patients. However, these efforts result in an average utilization of only 71% for the Oncology team. In contrast, the average utilization for all regionalized beds is 88%. This suggests that there is an opportunity to either reduce the dedicated capacity for Oncology patients or make the team more flexible to take on other patients.

Figure 5-12 provides a view of patient-wait-for-bed by patient origin for the redesign compared to the patient assignment algorithm without the redesign. The redesign incorporates the patient assignment algorithm, but the changes in the amount of capacity, the balance between Level 1 and Level 2 beds, and the prioritization of Oncology patients have the effect of "leveling the playing field" for patients from different origins. The main reason for this effect appears to be the differences in proportions of patients that experience the most benefit from the redesign (Level 1 and Oncology patients) by origin. In descending order the origins have the following percentages of Level 1 and/or Oncology patients: ICU 89.4%, hospital transfer 88.7%, front door 64.6%, ED 55.7%, floor 50.1%, PACU 42.9%. As a result of the large wait time decreases for Level 1 and Oncology patients, ICU and hospital transfer patients see their average patient-wait-for-bed decreased by over 55% (although they still wait longer than any other patient population). Notably, patients from the floor and front door now wait shorter on average than ED patients. For front door patients, this is likely due to the larger proportion of Level 1 and Oncology patients. There is no clear reason for this change for floor patients, but the lower mean is likely due to the lack of outliers in such a small population (365 patients per replication).

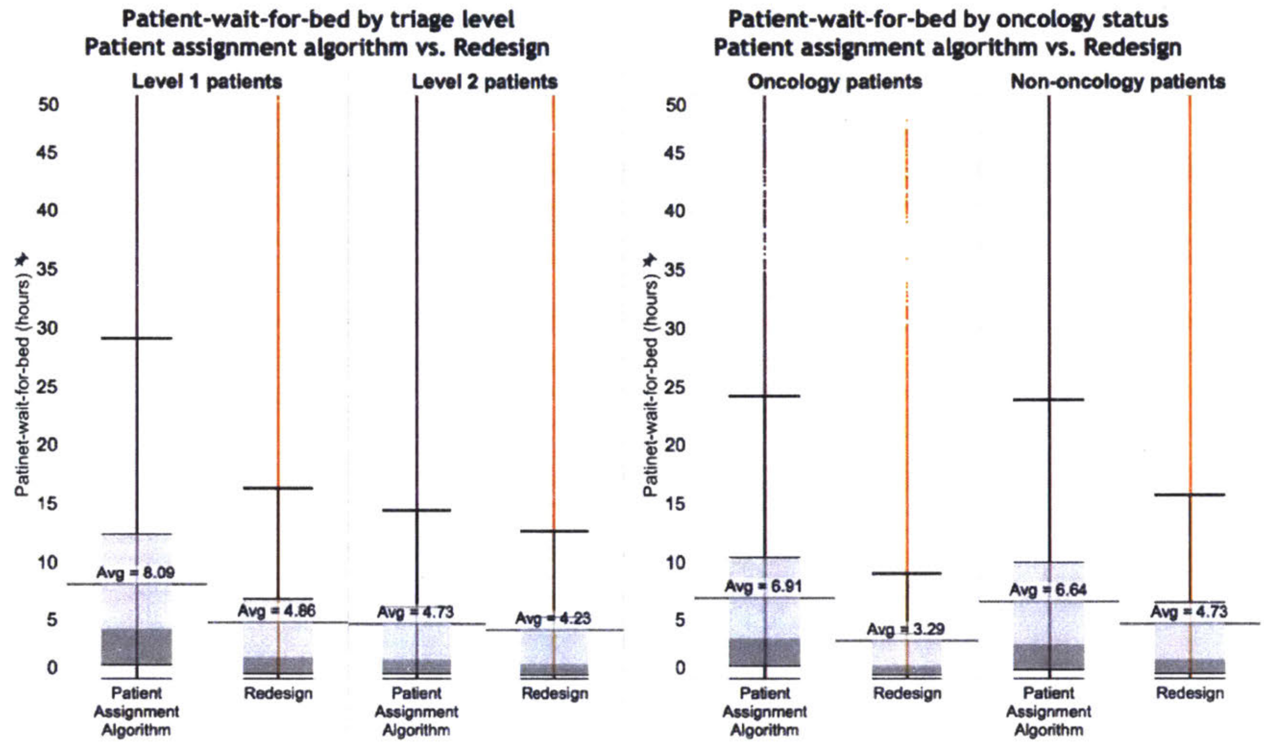

Figure 5-11: Comparison of patient-wait-for-bed times by patient characteristics, assignment algorithm vs. redesign

Note: Number of observations across 20 runs is as follows: Level 1 = 193,200, Level 2 = 141,200, Oncology = 32,220, Non-Oncology = 302,180. Maximum values are truncated on this graph for legibility. Maximum values for the patient assignment algorithm (in hours) are as follows: Level 1 238.13, Level 2 165.48, Oncology 175.47, Non-Oncology 238.13. Maximum values for the redesign (in hours) are as follows: Level 1 115.42, Level 2 164.93, Oncology 57.05, Non-Oncology 164.93. Data sources: Simulation results.

5.3.3 Swaps within Pods

The intervention known as 'swaps within pods' was found to have a very minor effect on overall patient-wait-for-bed. The mean wait time for all patients drops from 4.59 hours with the redesign to 4.58 hours when swaps within pods are introduced, which was found to not be statistically significant. The quantiles for wait times stay virtually identical. More detailed analysis of the impact of this intervention by patients' isolation status shows that the effect on the mean patient-wait-for-bed time is not statistically significant for any subgroup. It appears that the reason for these minimal impacts is the relatively small number of swaps within pods that occur. In the 20 replications of the simulation with swaps within pods implemented, only 9921 such swaps occur. This averages out to 496 swaps per replication or about 1.36 such swaps on each day of each replication. The small number of swaps between pods is due in part to the simulation preferring swaps within the unit to those across units

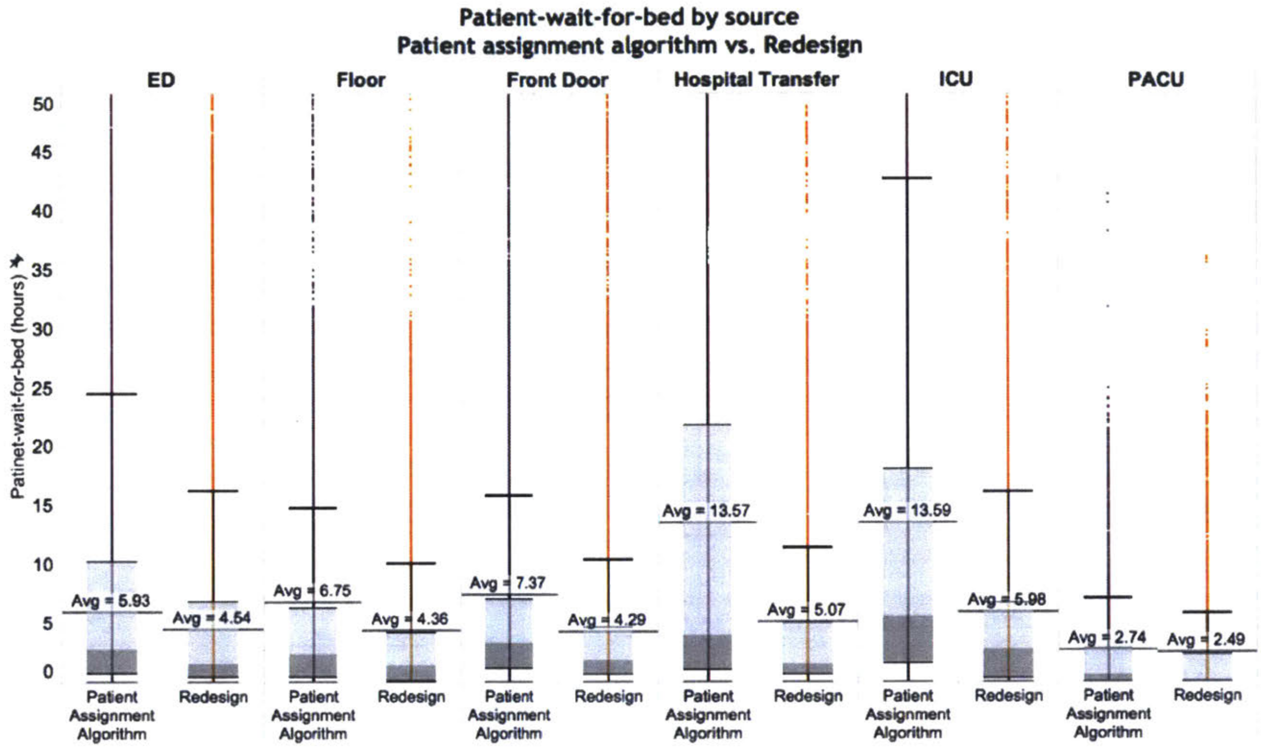

Figure 5-12: Comparison of patient-wait-for-bed times by origin, assignment algorithm vs. redesign

Note: Number of observations across 20 runs is as follows: ED = 266,960, Floor = 7300, Front Door = 30,040, Hospital transfer = 6,880, ICU = 20,140, PACU = 3080. Maximum values are truncated on this graph for legibility. Maximum values for the assignment algorithm scenario (in hours) are as follows: ED 94.27, Floor 165.48, Front Door 148.48, Hospital transfer 219.63, ICU 238.13, PACU 41.43. Maximum values for the redesign (in hours) are as follows: ED 83.42, Floor 164.93, Front door 148.58, Hospital transfer 93.78, ICU 115.42, PACU 36.12. Data sources: Simulation results.

within the pod. This preference is implemented because swaps across units would require nursing team handoffs in the current nursing structure. Although the effect of swaps within pods was quite small, based on the experimental design, they are incorporated in the early team assignment intervention.

5.3.4 Early Team Assignment

There were 21,709 patients selected for early assignment across the 20 replications of the simulation. This averages to 1085.45 patients per replication and 4.36 patients per non-holiday weekday (these are the days that early assignment is performed). Since we designed the intervention to choose up to six early assignment patients each day, this means that there are days that six appropriate patients cannot be found[33]. For those patients that are

[33]The requirements for patients to be selected for early team assignment are that they are non-Oncology, clean, Level 1, originating in the ED, and cannot be assigned to a bed immediately upon their request. We look for patients for early assignment between 8 am and 4 pm.

selected, mean patient-wait-for-team decreases from 9.93 to 7.28 hours. This represents a transfer of approximately 2,876 hours of patient care from ED physicians to DOM physicians in the course of a year, helping to ease workload on the ED and advance patients' care at an inpatient level.

The patient-wait-for-team results for those patients that were selected in the early assignment intervention are summarized in Figure 5-13. Patients selected for early assignment generally waited for their team longer than other ED patients when early assignment was not in effect. Figure 5-14 shows the patient-wait-for-team for all ED patients across quantiles. By comparing Figure 5-13 and Figure 5-14, one can see that the median patient-wait-for-team time for all ED patients without early assignment was 3.27 hours, but for those who are selected for early assignment it was 9.93 hours without the intervention. With the early assignment intervention, the selected patients' wait for teams is reduced with the median falling to 7.45 hours. Patients selected for early team assignment experience even larger decreases of approximately 25% in patient-wait-for bed at the higher quantiles, 3.63 hours at the 75th percentile and 4.89 hours at the 95th percentile.

Early team assignment results in an increase in average overall patient-wait-for-bed time from 4.58 hours to 4.74 hours, a difference of 0.16 hours 95% CI [0.12, 0.19 hours]. A similar average increase is observed for ED patients from 4.54 to 4.74 hours, a difference of 0.20 hours 95% CI [0.17, 0.24 hours]. The selected patients see their mean patient-wait-for-bed increase from 8.86 to 12.73 hours. This occurs because the early assignment patients are now restricted in the units that they can be assigned to and cannot go to the first available bed. Figure 5-15 shows these changes to patient-wait-for-bed for the relevant populations. However, it is important to refer back to Section 4.3.3 and to note that the implementation of early team assignment in the simulation does not capture all of the anticipated benefits of the intervention. Namely, discharges do not occur any earlier in the day. It is possible that if this effect were indeed realized that the negative effect on overall patient-wait-for bed would be neutralized or possibly even transformed into an improvement. It is also worth noting that the average patient-wait-for-bed for patients that are not selected for early assignment actually decreases slightly when the intervention is introduced, falling from 4.29 to 4.18 hours, a difference of 0.11 hours (95% CI [0.07, 0.14 hours]), which shows that the increase

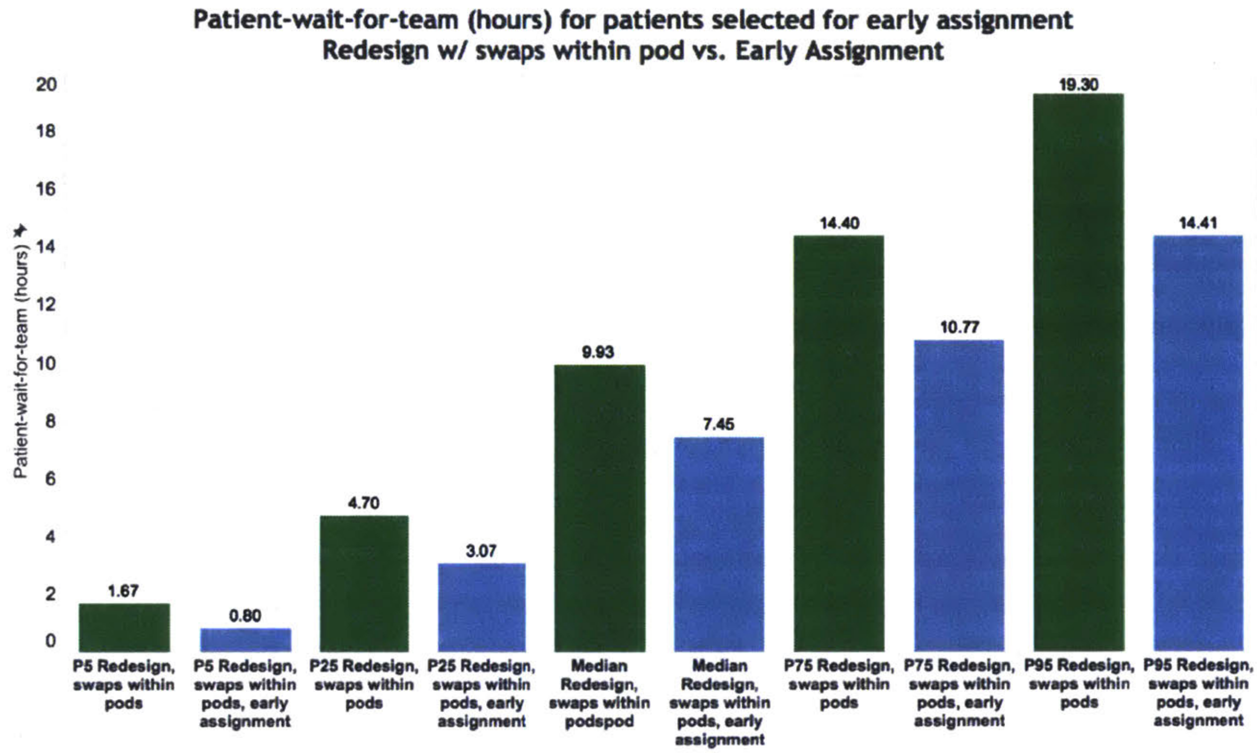

Figure 5-13: Comparison of patient-wait-for-team quantiles for patients selected for early assignment, redesign with swaps within pods vs. redesign with swaps within pods and early assignment

Data sources: Simulation results.

in overall average patient-wait-for-bed is driven by the patients who are selected for early assignment.

5.3.5 Full Flexibility

When full flexibility is introduced (meaning that all three teams in a pod can cover all the beds and that all patients can be swapped between units in a pod), the number of average daily swaps within pods increases from 1.48 to 4.43 (under the swaps within pod intervention the average was 1.36, but it increases to 1.48 with early assignment). This increase is associated with a decrease in average patient-wait-for bed for all patients from 4.74 to 4.51 hours, a difference of 0.23 hours 95% CI [0.19, 0.26 hours]. The average wait time for those patients that are assigned to the house staff units that form the pods decreases from 4.67 hours to 4.43 hours as cohorting is improved. The full distributions of patient-wait-for-bed in the two scenarios are shown in Figure 5-16. Patients from all origins experience

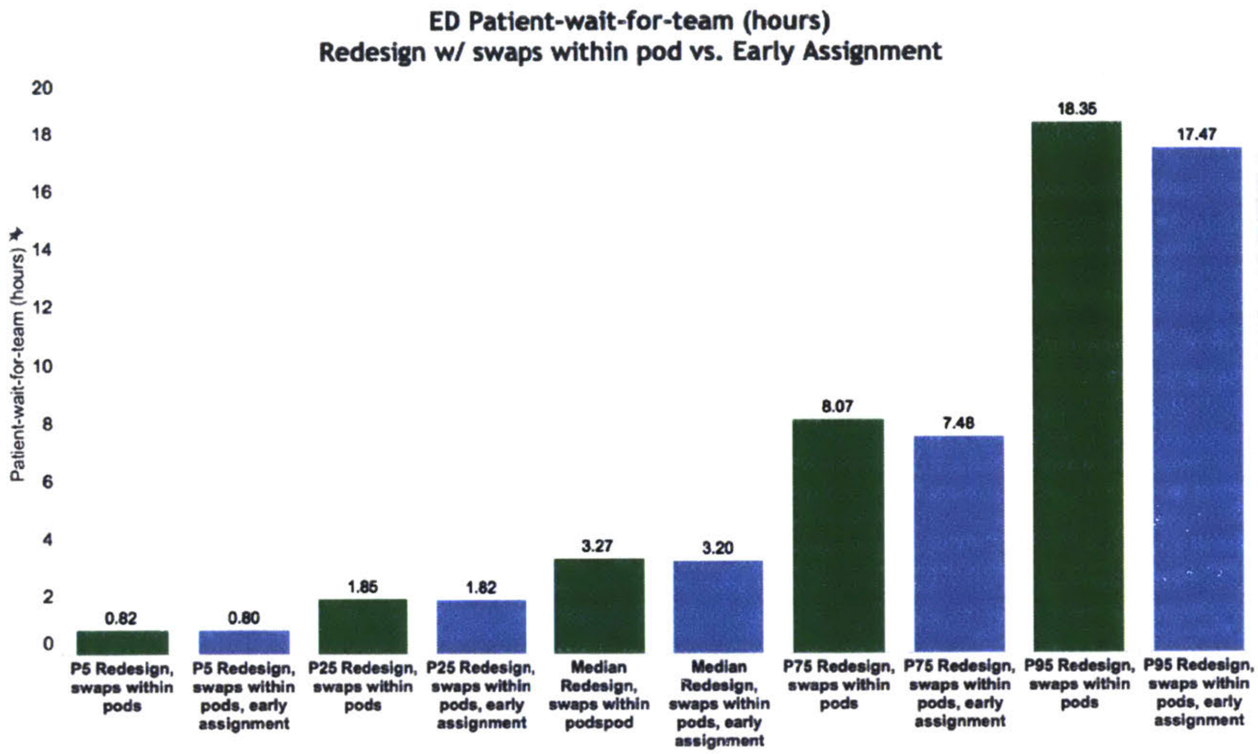

Figure 5-14: Comparison of patient-wait-for-team quantiles for all ED patients, redesign with swaps within pods vs. redesign with swaps within pods and early assignment
Data sources: Simulation results.

similar decreases.

5.3.6 First-come-first-served Prioritization

The implementation of first-come-first-served prioritization, where patient origin is disregarded when making assignment decisions, results in an overall decrease in average patient-wait-for-bed from 4.51 hours to 4.45 hours, a difference of 0.06 hours 95% CI [0.03, 0.10]. Despite this relatively small change in the overall wait, there are large changes based on patient origin as shown in Figure 5-17. When first-come-first-served prioritization is applied, hospital transfer and ICU patients see the largest decreases of 33% (1.6 hours) and 30% (1.73 hours) respectively, ED patients have the longest average patient-wait-for-bed time of all the origins at 4.65 hours, PACU patients continue to have the shortest wait at 2.64 hours.

5.3.7 Eliminate Staffing Closures

The final intervention implemented is the elimination of staffing closures. This was found to have a minimal impact on patient-wait-for-bed, decreasing the overall mean from 4.45 to

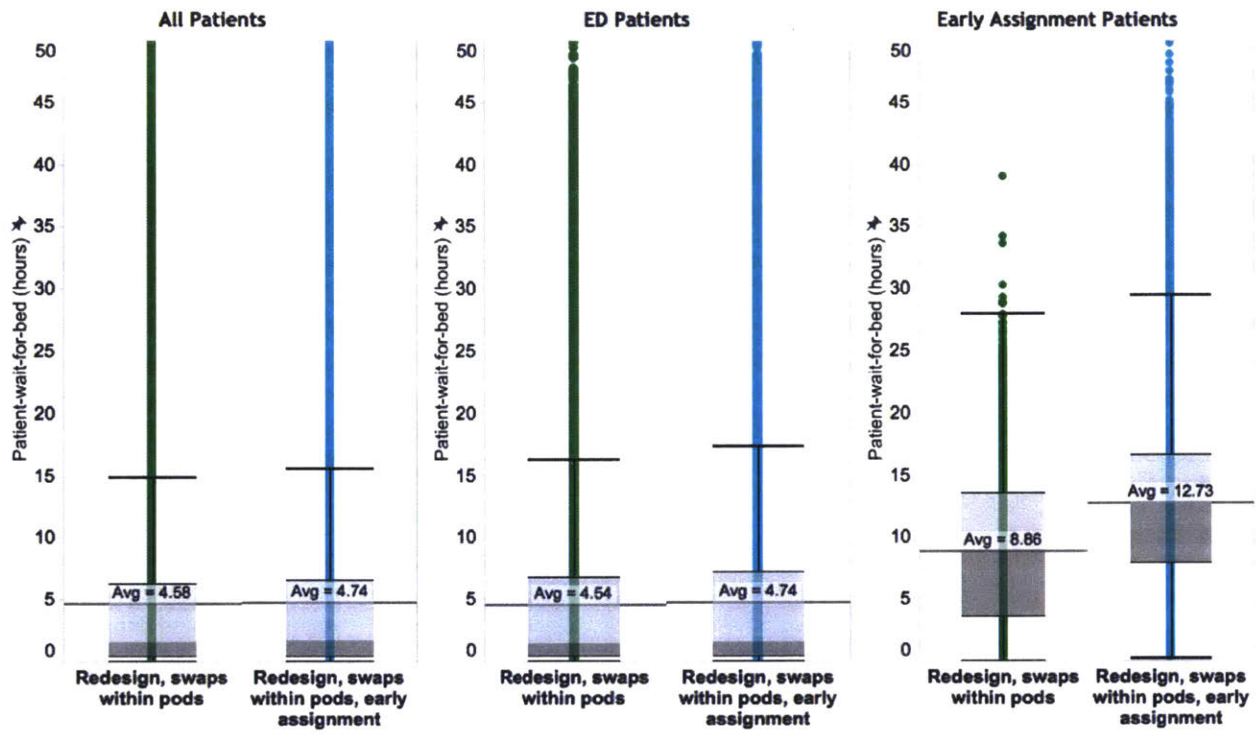

Figure 5-15: Comparison of patient-wait-for-team times for relevant populations, redesign with swaps within pods vs. redesign with swaps within pods and early assignment

Note: Number of observations across 20 runs is as follows: All Patients = 334,400, ED = 266,960, Early Assignment Patients = 21,709. Maximum values are truncated on this graph for legibility. Maximum values for the redesign with swaps within pods (in hours) are as follows: All patients 165.1, ED 91.47, early assignment patients 39.05. Maximum values for the redesign with swaps within pods and early team assignment (in hours) are as follows: All patients 165.4, ED 92.08, early assignment patients 80.05. Data sources: Simulation results.

4.41 hours, a difference of 0.04 hours 95% CI [0.01, 0.07 hours]. The full distributions of patient-wait-for-bed in the two scenarios are shown in Figure 5-18. Patients from all origins experience similar decreases.

5.4 Summary of Intervention Results

The major effects of the interventions are summarized as follows. All effects reported are incremental to the previous intervention.

- **Patient assignment algorithm:** Decrease overall mean patient-wait-for-bed approximately 9% from the base scenario (from 7.36 to 6.67 hours). Large reductions in mean wait time for hospital transfer and ICU patients (from 20.17 to 13.57 and 24.09 to 13.59 hours respectively) and modest increases for ED patients (from 5.67 to 5.93).
- **DOM redesign:** Decrease overall mean patient-wait-for-bed approximately 31% (from

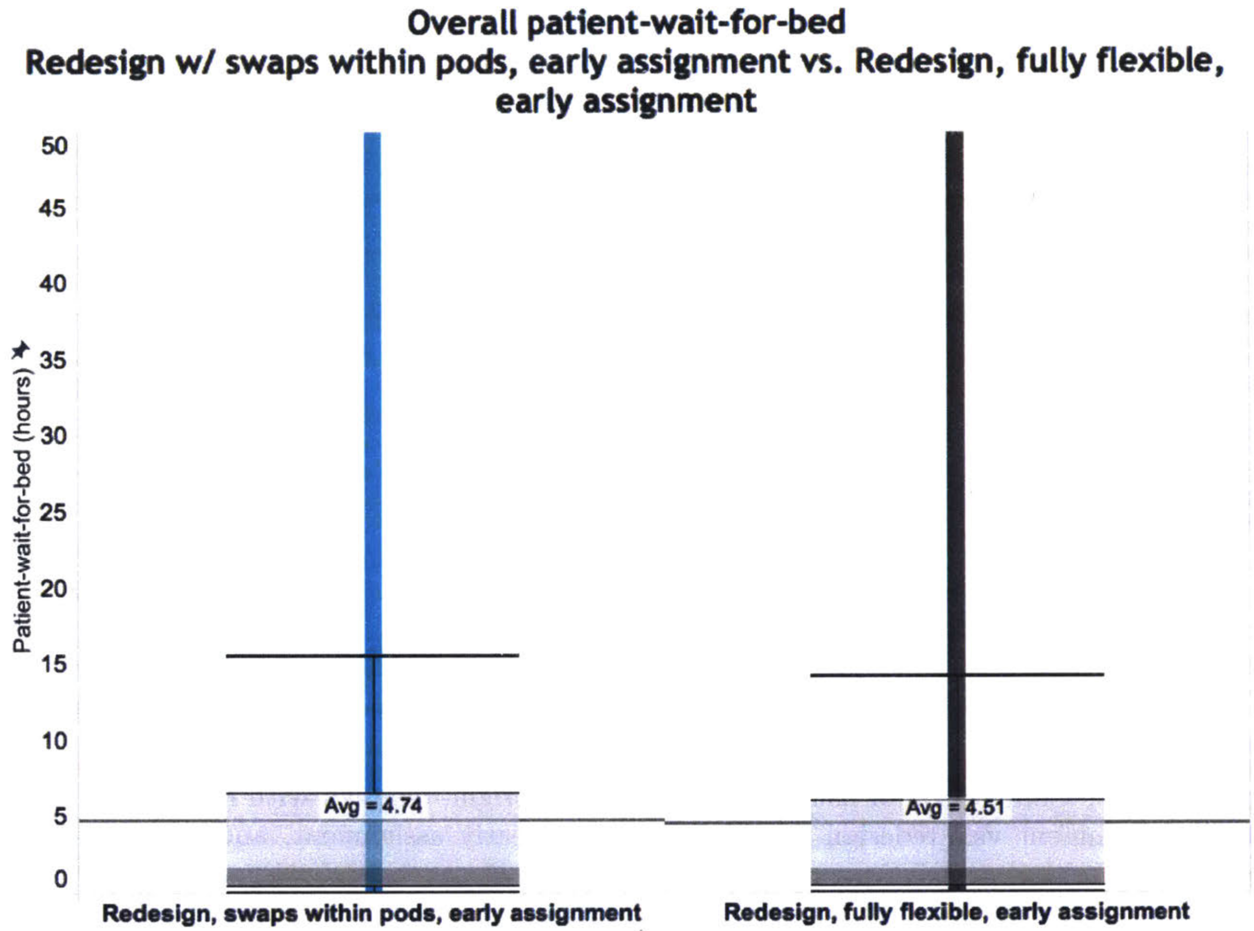

Figure 5-16: Comparison of patient-wait-for-bed times for all patients, redesign with swaps within pods and early assignment vs. redesign with full flexibility and early assignment

Note: Number of observations across 20 runs is as follows: All Patients = 334,400. Maximum values are truncated on this graph for legibility. Maximum values for the redesign with swaps within pods and early team assignment (in hours) are as follows: All patients 165.4. Maximum values for the redesign with full flexibility and early team assignment (in hours) are as follows: All patients 165.5. Data sources: Simulation results.

6.67 to 4.59 hours). 'Level the playing field' by source due to large decreases for Level 1 and Oncology patients.

- **Swaps within pods:** No significant effect on overall mean patient-wait-for-bed. Only 1.36 swaps between different units take place each day on average.

- **Early team assignment:** Increase overall mean patient-wait-for-bed approximately 3% (from 4.58 to 4.74 hours). Decrease mean patient-wait-for-team for patients selected approximately 27% (from 9.93 to 7.28 hours). Transfer approximately 2,876 hours of workload from the ED physicians to DOM physicians in a year.

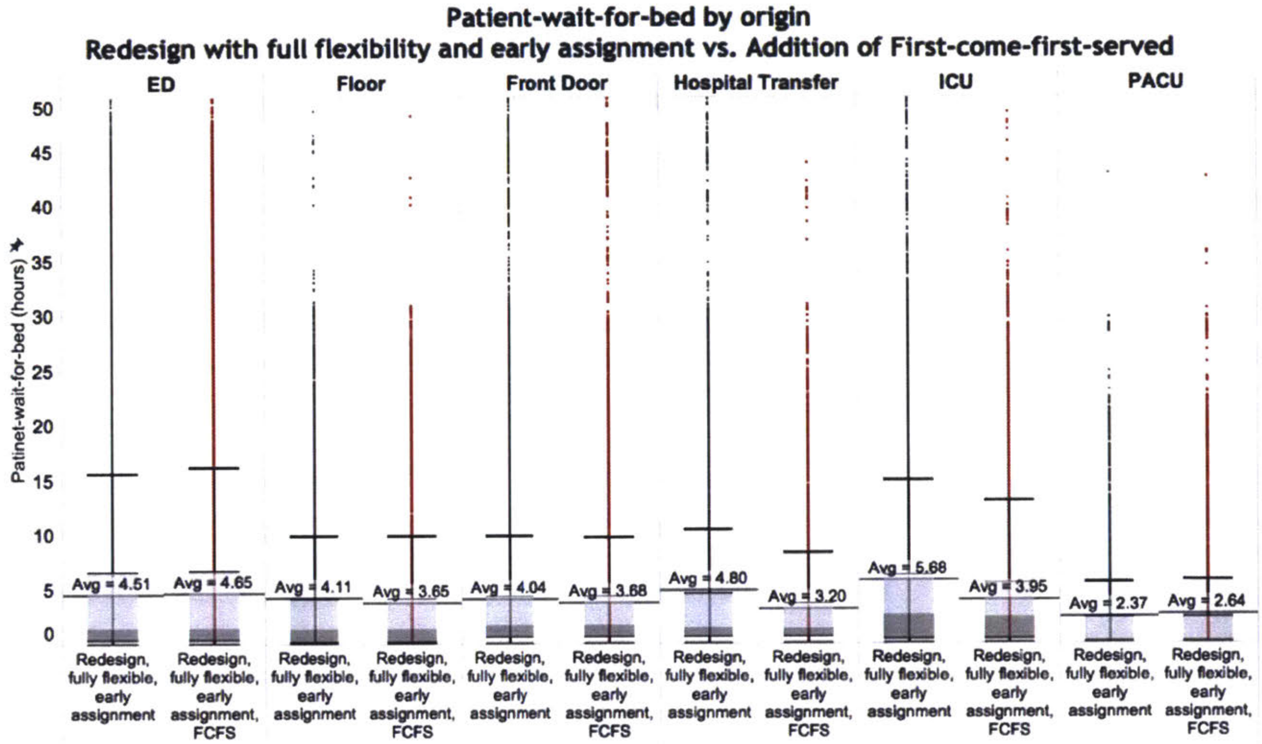

Figure 5-17: Comparison of patient-wait-for-bed by origin, redesign with full flexibility and early assignment vs. redesign with full flexibility, early assignment, and first-come-first-served

Note: Number of observations across 20 runs is as follows: ED = 266,960, Floor = 7300, Front Door = 30,040, Hospital transfer = 6,880, ICU = 20,140, PACU = 3080. Maximum values are truncated on this graph for legibility. Maximum values for the scenario without FCFS (in hours) are as follows: ED 91.47, Floor 165.52, Front Door 148.45, Hospital transfer 90.13, ICU 114.98, PACU 42.95. Maximum values for the FCFS scenario (in hours) are as follows: ED 108.02, Floor 57.28, Front Door 64.58, Hospital transfer 43.90, ICU 56.35, PACU 42.58. Data sources: Simulation results

- **Full flexibility:** Decrease overall mean patient-wait-for-bed approximately 5% (from 4.74 to 4.51 hours). Increase number of swaps between different units in pod to 4.43 per day on average.

- **First-come-first-served prioritization:** Decrease overall mean patient-wait-for-bed approximately 1% (from 4.51 to 4.45 hours). Patient-wait-for-bed decreases by more than 30% for hospital transfer and ICU patients, ED patients now have the longest mean wait.

- **Eliminate staffing closures:** Decrease overall mean patient-wait-for-bed slightly (from 4.45 to 4.41 hours).

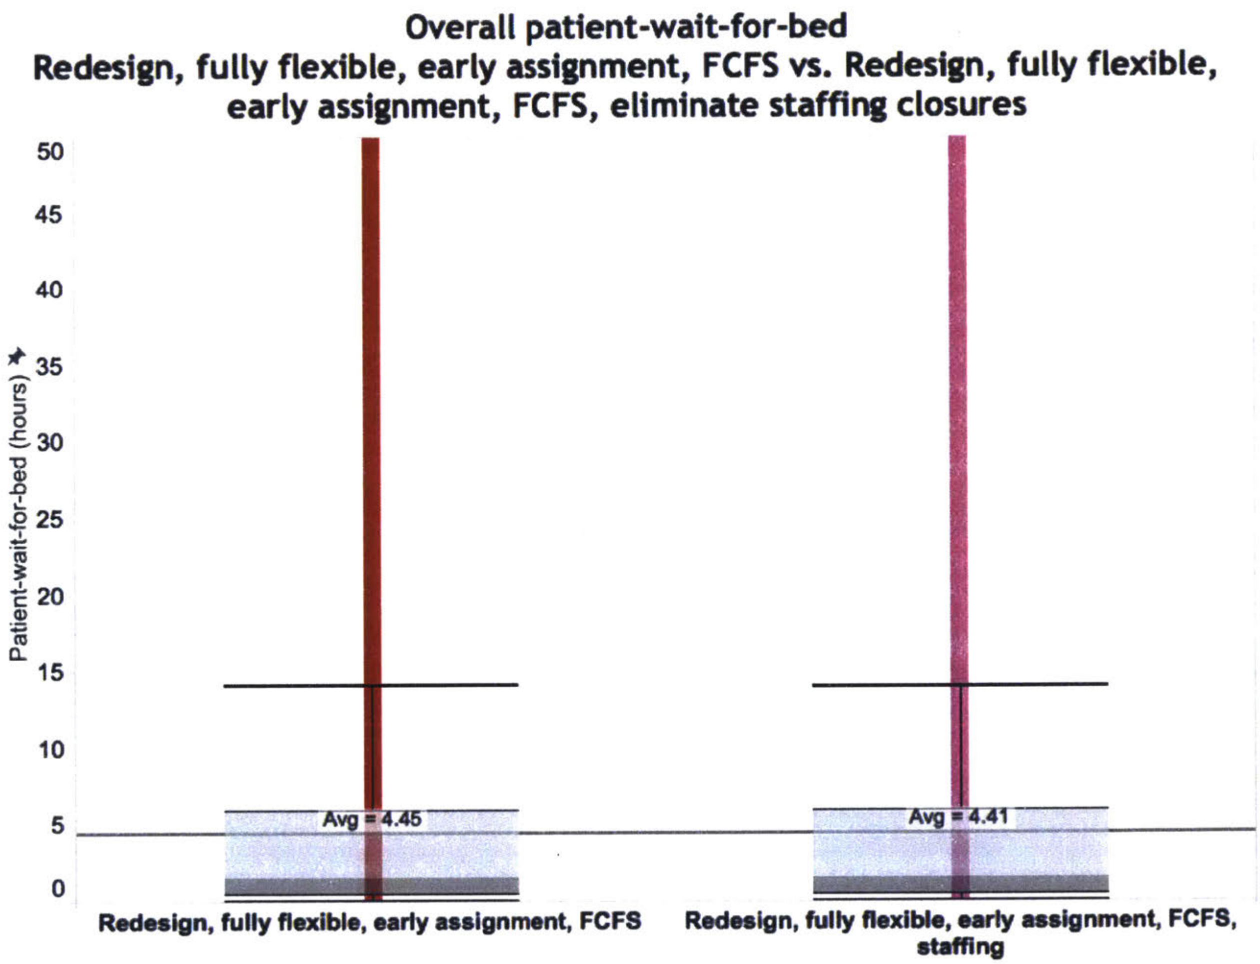

Figure 5-18: Comparison of patient-wait-for-bed times for all patients, redesign with full flexibility, early assignment, and first-come-first-served vs. redesign with full flexibility, early assignment, first-come-first-served and no staffing closures

Note: Number of observations across 20 runs is as follows: All Patients = 334,400. Maximum values are truncated on this graph for legibility. Maximum values for the redesign with full flexibility, early team assignment, and FCFS (in hours) are as follows: All patients 108.0. Maximum values for the redesign with full flexibility, early team assignment, FCFS, and elimination of staffing closures (in hours) are as follows: All patients 107.9. Data sources: Simulation results.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Recommendations and Conclusions

This section provides recommendations for the implementation of interventions based on the results discussed in Section 5. It also discusses opportunities for study that could supply valuable insights on patient flow at the hospital and opportunities for process optimization. Several of these opportunities could be evaluated using the simulation developed in this thesis. It is important to remember that all of the results presented in the previous chapter and discussed here are obtained using 2015 data. Prior to implementing any interventions it is recommended that the hospital consider whether there are changes in any underlying variables (e.g. LOS, Level 1 vs. Level 2 mix) that could affect the results.

6.1 Recommendations

This thesis provides evidence that the implementation of the DOM redesign was a productive step towards reducing patient-wait-for-bed and estimates the improvement in wait time that could be realized if the redesign were implemented in combination with a patient assignment algorithm. One further recommendation is that the DOM look at the extent to which the Oncology sub-unit is being utilized and potentially reduce the capacity or add flexibility for this team to cover other patients.

A key recommendation of this thesis is that Admitting and the DOM take a closer look at their patient assignment practices. This recommendation is based on simulation results that indicate that overall patient-wait-for-bed for 2015 could have been reduced by 9% if such an algorithm had been in use, and by 38% under the combination of the redesign and the patient assignment algorithm. The proposed patient assignment algorithm still prioritizes patients by origin, with ED patients having the highest priority. This is important because it supports the hospital's goal of minimizing ED congestion. For this reason, we do not recommend that the hospital adopt a completely first-come-first-served prioritization strategy. The incremental impact of such a strategy on overall patient-wait-for-bed time is shown to be very small compared to what is obtained by implementing the assignment algorithm that does consider patient origin. In our judgment, this improvement is not worth the increased wait for ED patients.

The implementation of a patient assignment algorithm needs to be spearheaded by the Admitting department. Bed managers would need to be trained on the proper prioritization of patients based on the algorithm. Although the algorithm presented here is straightforward, it is likely that it will be challenging to instill the mindset that a bed should never go unassigned when there is an appropriate patient waiting, regardless of their origin. We see this as the biggest hurdle to implementation of the patient assignment algorithm.

The results of the simulation show that the implementation of early assignment can help the ED patients who historically wait the longest to get inpatient-level care faster and decrease ED physician workload while causing minimal increases to overall average patient-wait-for-bed. The implementation of this intervention in the simulation does not include the expected benefit of earlier discharges due to incentive re-alignment, which we hypothesize may negate the observed small increase is overall average patient-wait-for-bed. We believe that it is worth pursuing this intervention through an extended pilot program. A short pilot was completed in Autumn 2016 with promising results. By completing a longer pilot, the DOM would be able to better assess the impact of early assignment on its educational goals and potentially refine the criteria for selecting patients.

We recommend that the DOM consider adding more flexibility to the redesign. Our results show that full flexibility of physician teams in pods and the ability to swap all patients within the pods could produce overall average patient-wait-for-bed savings of approximately 5%. To decide whether to implement this change, the DOM team should consider all the costs of implementation. These changes would require all the pod-based house staff physician teams to see patients on two different units; it should be carefully considered whether this would cause any degradation to patient care. Additionally, implementing swaps across units within pods means that the nursing team structure would either need to change to be pod-based or that patients would need to change nurses when they move from one unit to another. Given the modest improvements to wait times, the team needs to carefully consider the implications of either change before implementation. If the decision is made to transition nursing to a pod-based structure, there may be an added benefit of reducing staffing closures through the pooling of resources. While a reduction in staffing closures would certainly be beneficial to the daily operations of the hospital, based on the low historical volume of closures, it should not be expected to have a large effect on patient-wait-for-bed.

6.2 Opportunities for Further Study

This thesis has raised several questions that encourage further study. The simulation of patient flow developed for this thesis provides a framework that can be used to evaluate other interventions that are developed as these questions are investigated.

One important insight gleaned from this simulation is the importance of the algorithm used to assign patients to beds. This thesis uses a straightforward algorithm informed by our understanding of current practices and priorities. There is an opportunity to develop a more sophisticated algorithm based on optimization that could then be tested in this modeling environment. Based on the reductions in wait time seen from the algorithm developed for this thesis, implementing a more sophisticated algorithm could be very beneficial to the hospital.

In addition to the proposed pilot of early team assignment, we believe that there could be more work done to model the effect of the anticipated earlier discharges using the simulation framework. Such an effort might incorporate data collected during the pilot to better estimate the anticipated impact of the early team assignment intervention (with earlier discharges taken into account) and simulate the full implementation of early team assignment.

This thesis built on previous work in several ways, one of which was the study of delays leaving the ICU and their effect on hospital LOS. We did not attempt to study the relationship between assignment delays and hospital LOS for other patient populations. Estimating these effects and modeling them in the simulation could provide more accurate results and recommendations on other interventions.

The current de-prioritization of patients transferring out of the ICU and the consistent use of non-regionalized beds for General Medicine patients suggest a fundamental imbalance in the allocation of beds at the hospital. This simulation provides a framework that if extended to the greater patient population of MGH could help to assess alternative proposals on how the capacity should be allocated to services and levels of care.

Appendix A

List of General Medicine Teams from PEPL

Team descriptions from PEPL used to identify general medicine patients: Medicine Medicine V, Medicine Medicine J, Medicine Medicine IV, Medicine Med Team AHS - Z, Medicine Med Team 2 INT3, Medicine Med Team 2 INT2, Medicine Med Team 1 INT4, Medicine Med Team 1 INT3, Medicine Med Team 1 INT2, Medicine Med Team 1 INT1, Medicine Med Team 1 DF, Medicine Bigelow X, Medicine Bigelow E, Medicine Bigelow D, Medicine Bigelow C, Medicine Bigelow B, Medicine Bigelow A

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

Team Coverage and Caps by Scenario

Team coverage of units by scenario.

Scenario	Unit	Lvl1Team_1	Lvl1Team_2	Lvl1Team_3	Lvl2Team_1	Lvl2Team_2	Lvl2Team_3	Lvl2Team_4
Base	White 8	Big A			Green	Big A		
	Bigelow 9	Red 1			Red 1	Green		
	White 9	Big B	Red 2		Green	Big B	Red 2	
	White 10	Big C			Big C			
	Bigelow 11	Big E			Green	Big E		
	White 11	Big D			Green	Big D		
	Ellison 12				Blue			
	Ellison 16	Team 1			Team 1			
	Ellison 19	Yellow			Yellow			
	Phillips 20	Yellow			Yellow			
	Non-regionalized				Green	Blue	Yellow	Red
Redesign	White 8	Big A	Flex X		Big A	Flex X		
	Bigelow 9	Big B	Flex X		Big B	Flex X		
	White 9	Big C	Flex Y		Big C	Flex Y		
	White 10	Big D	Flex Y		Big D	Flex Y		
	Bigelow 11	Big E	Flex Z		Big E	Flex Z		
	White 11	Big F	Flex Z		Big F	Flex Z		
	Ellison 12	Blue			Blue			
	Ellison 16	Red			Red			
	Ellison 19	Yellow			Yellow			
	Phillips 20	Yellow			Yellow			
	Non-regionalized				Green	Red	Blue	Yellow
Fully Flexible	White 8	Big A	Flex X	Big B	Big A	Flex X	Big B	
	Bigelow 9	Big B	Flex X	Big A	Big B	Flex X	Big A	
	White 9	Big C	Flex Y	Big D	Big C	Flex Y	Big D	
	White 10	Big D	Flex Y	Big C	Big D	Flex Y	Big C	
	Bigelow 11	Big E	Flex Z	Big F	Big E	Flex Z	Big F	
	White 11	Big F	Flex Z	Big E	Big F	Flex Z	Big E	
	Ellison 12	Blue			Blue			
	Ellison 16	Red			Red			
	Ellison 19	Yellow			Yellow			
	Phillips 20	Yellow			Yellow			
	Non-regionalized				Green	Red	Blue	Yellow

Team patient caps by scenario.

Team	Base Scenario	Redesign	Fully Flexible
Bigelow A	24	16	16
Bigelow B	20	16	16
Bigelow C	20	16	16
Bigelow D	20	16	16
Bigelow E	24	17	17
Bigelow F	N/A	17	17
Flex X	N/A	12	12
Flex Y	N/A	13	13
Flex Z	N/A	15	15
Team 1	36	N/A	N/A
Oncology	N/A	16	17
Red 1	8	N/A	N/A
Red 2	2	N/A	N/A
Red	N/A	20	20
Blue	36	36	36
Yellow	34	34	34
Green	82	31	31

Appendix C

Glossary

Acuity: Intensity of clinical care a patient requires.

Admission: Event when a patient is admitted to the hospital and reclassified as an inpatient.

Capacity Disaster: The state in which Code Help has been activated for two hours or more and there are still ED Boarders present.

Code Help: A state-mandated policy requiring hospitals to move all admitted patients out of the ED within a 30-minute period after the ED's maximum occupancy is reached or exceeded. Maximum occupancy is influenced by the number of patients present and their acuity/clinical complexity. A prolonged Code Help requires the hospital to make a report to government officials and might require the ED to divert patients to other hospitals.

Discharge: In the context of this project, event where a patient's stay in a General Medicine bed ends.

Hospital Medicine Group (HMG): Physican teams composed of board-certified Internal Medicine physicians, known as Hospitalists. Hospitals cover patients in a regionalized and non-regionalized fashion.

House Staff: Physician teams that are primarily composed of residents who are supervised by Attending physicians. House staff only cover patients in a regionalized fashion.

Level 1: A patient whose acuity high enough that he/she needs to be cared for in a regionalized setting.

Level 2: A patient who can be cared for in a regionalized or non-regionalized setting.

Non-regionalized bed: A bed that is not covered by a physician team that is physically based on the unit.

Patient stay: The time that a patient is in a General Medicine unit, whether it is regionalized or non-regionalized. See Section 4.1.2 for details.

Regionalized bed: A bed that is covered by a physician team that is physically based on the unit. A bed on a unit that has a local physician team could be non-regionalized if that team is at their capacity and the bed is covered by a non-local team.

Appendix D

Basic Patient Flow in Simulation

Chart appears on next page due to size

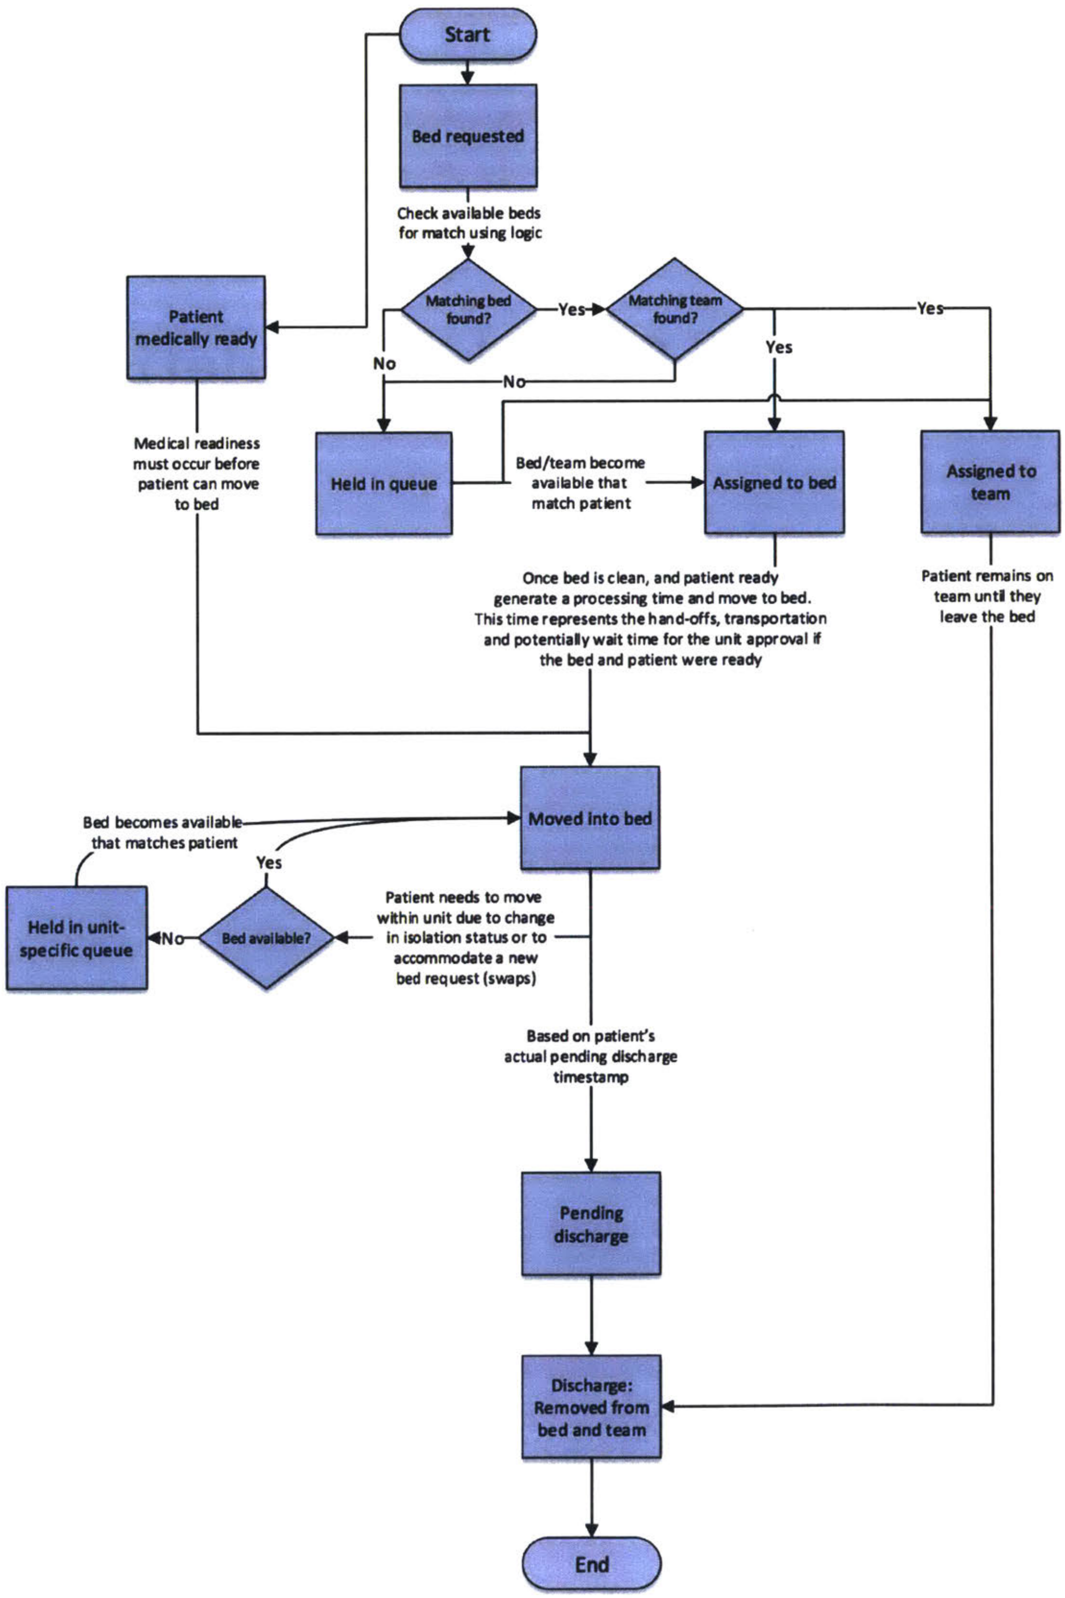

Start
Bed requested
Check available beds for match using logic
Matching bed found?
Yes
Matching team found?
Yes
Yes
No
No
Patient medically ready
Medical readiness must occur before patient can move to bed
Held in queue
Bed/team become available that match patient
Assigned to bed
Assigned to team
Once bed is clean, and patient ready generate a processing time and move to bed. This time represents the hand-offs, transportation and potentially wait time for the unit approval if the bed and patient were ready
Patient remains on team until they leave the bed
Moved into bed
Bed becomes available that matches patient
Yes
Held in unit-specific queue
No
Bed available?
Patient needs to move within unit due to change in isolation status or to accommodate a new bed request (swaps)
Based on patient's actual pending discharge timestamp
Pending discharge
Discharge: Removed from bed and team
End

Appendix E

Procedure for Correcting Isolation Status

This appendix describes the procedure used to adjust patients' isolation statuses when they were found to be in conflict with the patients' historical room placement. The procedure for determining the patients' isolation status initially is described in Section 4.1.2.

For two patients X and Y that shared a semi-private room and had conflicting isolation statuses (known as an 'incorrect cohort') lasting more than four hours we made the following changes:

- If the incorrect cohort began when patient X changed isolation status (X and Y were already in the room together and had previously matching isolation statuses) and ended when patient X or Y left the room:
  - If the incorrect cohort ended when patient X was discharged, remove the last isolation status from patient X.
  - If the incorrect cohort ended when patient X moved rooms and subsequently became non-cohortable, change patient x's conflicting isolation status to 'non-cohortable' (to match subsequent non-cohortable status).
  - If the incorrect cohort ends by patient Y moving to another room, delay the onset of the incorrect isolation status of X to cause a maximum of four hours of the

incorrect cohort.

- If the incorrect cohort occurred immediately when patient X entered the room and stopped when X or Y changed isolation status:
  - If patient X changes isolation status to match Y, make this isolation status change occur when patient X entered the room.
  - If patient Y changes isolation status to match X, make this isolation status change occur when patient X entered the room.
- If the incorrect cohort started and stopped while patients X and Y were in the room together when X changed to and from a conflicting isolation status:
  - Remove the isolation status that causes the conflict.
- If the incorrect cohort occurs when two non-cohortable patients are staying together in a semi-private room:
  - Change both isolation statuses to 'Influenza' to allow cohorting to occur.
- If the incorrect cohort occurs when patient X enters the room and stops when patient X or Y leaves the room:
  - Here, we check to see if we are confident in the labeling of the isolation statuses for each patient. We check to see if patient X had roommates before and after and award a +1 for each roommate isolation status that matches the isolation status of X during the period of the incorrect cohort and a -1 for each isolation status that matches the isolation status of Y during the incorrect cohort. A +1 is also given if X was non-cohortable during the incorrect cohort and did not have a roommate before or after. These points are awarded similarly for Y. Thus, X and Y receive a maximum score of +2 each and a minimum score of -2 each.
  - If patient X is non-cohortable during the incorrect cohort, we change X's isolation status to match Y according to the confidence levels just calculated. So, if X received a +1 for his/her time after the incorrect cohort (thus supporting the

label as non-cohortable), we only change the non-cohortable status up to the end of the incorrect cohort and apply the same principle for the time before the incorrect cohort. We repeat a similar process if Y is non-cohortable.

- If X and Y are not non-cohortable, we change the isolation status of the patient with the lower confidence score to match the roommate, again following the confidence levels previously calculated.
- If X and Y receive the same confidence level (a tie), we change the confidence level of the patient with the more 'severe' isolation status, where severity is ranked as follows: clean < VRE < MRSA < MRSA & VRE < Influenza < non-cohortable. Again, in making these changes, we follow the confidence levels previously calculated (to avoid changing isolation statuses unnecessarily that would cause more conflict).

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] Avery Comarow and Ben Harder. 2016-17 best hospitals honor roll and overview. http://health.usnews.com/health-care/best-hospitals/articles/best-hospitals-honor-roll-and-overview. Accessed: 2016-12-21.

[2] Massachusetts General Hospital (MGH). Massachusetts General Hospital - Hospital Overview. http://www.massgeneral.org/about/overview.aspx. Accessed: 2016-12-21.

[3] Massachusetts General Hospital (MGH). Massachusetts General Hospital - Partners Healthcare. http://www.massgeneral.org/partners.aspx. Accessed: 2016-12-21.

[4] Broadcast MGH. From the desktop of Peter L. Slavin, MD - Mass Generally Speaking. Email Newsletter, February 2016.

[5] Andrew Johnston. Quantifying the impact of care team discontinuities on medically unnecessary delays in inpatient flow. Master's thesis, Massachusetts Institute of Technology, 2016.

[6] Sara Dolcetti. Analyzing the impact of delays for patient transfers from the ICU to general care units. Master's thesis, Massachusetts Institute of Technology, 2015.

[7] Linda V. Green. The vital role of operations analysis in improving healthcare delivery. *Manufacturing & Service Operations Management*, 14:488–494, 2012.

[8] Benjamin C. Sun, Renee Y. Hsia, Robert E. Weiss, David Zingmond, Li-Jung Liang, Weijuan Han, Heather McCreath, and Steven M. Asch. Effect of emergency department crowding on outcomes of admitted patients. *Annals of Emergency Medicine*, 61:605–611, 2013.

[9] Drew B. Richardson. The access-block effect: relationship between delay to reaching an inpatient bed and inpatient length of stay. *Medical Journal of Australia*, 177:492–495, 2002.

[10] Farah Mustafa, Peadar Gilligan, Deborah Obu, Patrick O'Kelly, Eimear O'Hea, Catherine Lloyd, Sherif Kelada, Attracta Heffernan, and Patricia Houlihan. Delayed discharges and boarders: a 2-year study of the relationship between patients experiencing delayed discharges from an acute hospital and boarding of admitted patients in a crowded ED. *Emergency Medicine Joural: EMJ*, 33:636–640, 2016.

[11] Emilie S. Powell, Rahul K. Khare, Arjun K. Venkatesh, Ben D. Van Roo, James G. Adams, and Gilles Reinhardt. Administration of emergency medicine: The relationship between inpatient discharge timing and emergency department boarding. *Journal of Emergency Medicine*, 42:186–196, 2012.

[12] R. Ben Bachouch, Alain Guineta, and Sonia Hajri-Gaboujb. An integer linear model for hospital bed planning. *International Journal of Production Economics*, 140:833–843, 2012.

[13] Bex George Thomas, Srinivas Bollapragada, Kunter Akbay, David Toledano, Peter Katlic, Onur Dulgeroglu, and Dan Yang. Automated bed assignments in a complex and dynamic hospital environment. *Interfaces*, 43:435–448, 2013.

[14] Steven Thompson, Manuel Nunez, Robert Garfinkel, and Matthew D. Dean. Efficient short-term allocation and reallocation of patients to floors of a hospital during demand surges. *Operations Research*, 57:261–273, 2009.

[15] Alexandre Mazier, Xiaolan Xie, and Marianne Sarazin. Real-time patients assignment: A method for improving emergency department flow. *2010 IEEE Workshop on Health Care Management (WHCM)*, 2010.

[16] Benjamin A. Christensen. Improving ICU patient flow through discrete-event simulation. Master's thesis, Massachusetts Institute of Technology, 2012.

[17] Jonas Hiltrop. Modeling neuroscience patient flow and inpatient bed management. Master's thesis, Massachusetts Institute of Technology, 2014.

[18] Sean T. McNichols. Reducing intraday patient wait times through just-in-time bed assignment. Master's thesis, Massachusetts Institute of Technology, 2015.

[19] Douglas C. Montgomery. *Applied statistics and probability for engineers.* John Wiley and Sons, Inc, Reading, Massachusetts, sixth edition, 10 January 2014.