SANDIA REPORT
SAND2020-5569
Printed Click to enter a date

Sandia National Laboratories

Uncertainty Analysis Framework for the Hospital Resource Supply Model for Covid-19

Walt E. Beyeler, Christopher Frazier, Daniel Krofcheck, Laura P. Swiler, Teresa Portone, Katherine Klise

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

ABSTRACT

In March and April of 2020 there was widespread concern about availability of medical resources required to treat Covid-19 patients who become seriously ill. A simulation model of supply management was developed to aid understanding of how to best manage available supplies and channel new production. Forecasted demands for critical therapeutic resources have tremendous uncertainty, largely due to uncertainties about the number and timing of patient arrivals. It is therefore essential to evaluate any process for managing supplies in view of this uncertainty. To support such evaluations, we developed a modeling framework that would allow an integrated assessment in the context of uncertainty quantification. At the time of writing there has been no need to execute this framework because adaptations of the medical system have been able to respond effectively to the outbreak. This report documents the framework and its implemented components should need later arise for its application.

ACKNOWLEDGEMENTS

The authors acknowledge the support of the Office of Science for funding this work, and our colleagues Patrick Finley and Sean DeRosa for their contributions and support.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

This page left blank

1. INTRODUCTION

In March and April of 2020 there was widespread concern about availability of medical resources required to treat Covid-19 patients who become seriously ill. A simulation model of the human and technical resources required to treat hospitalized patients was developed to help assess the prospect of shortfall. This Treatment Model was embedded in an uncertainty quantification (UQ) framework to provide information about the risk of resources demands exceeding available capacity, and the uncertainties driving any unfavorable outcomes (Swiler et al., 2020).

Reallocation among existing stockpiles and management of new resource source might be used to mitigate resource shortfall risks, should they be likely. A simulation model of supply management was therefore developed to assess such mitigations (Frazier et al., 2020). Analysis of various outbreak projection scenarios using the Treatment Model showed that forecasted demands for critical therapeutic resources have tremendous uncertainty, largely due to uncertainties about the number and timing of patient arrivals. It is therefore essential to incorporate uncertainty analysis in the evaluate process for managing supplies. We developed a modeling framework that would allow an integrated assessment of the Treatment Model and models of supply sources and routing managers in the context of uncertainty quantification.

Because adaptations of the medical system have been able to respond effectively to the outbreak as of the time of writing, there has been no need to apply the framework on the real system. This report documents the framework, its implemented components, and describes an example application so that it can be rapidly deployed in the event that need arises. The following Chapter (2) describes the framework and its components and illustrates its operation. Chapter 3 discusses specific implementations of Routers based on optimization formulations.

2. FRAMEWORK COMPONENTS AND INTERACTIONS

2.1. Overview

Hospital resource demands arise from patient inflows. Our analysis focuses on those resources needed to treat patients diagnosed with Covid-19 and requiring hospitalization. Other hospital patient flows and Covid-19 patients being cared for elsewhere are not considered.

Figure 1 illustrates the role of the Treater model object in translating an impinging patient arrival stream into hospital resource requirements. Model details are provided elsewhere (Swiler et al. 2020, Beyeler et al. 2020). The interface flows in Figure 1 show how Treaters interact with their environment as resource consumers:

- Treaters are configured at initialization to describe the kinds of processes they implement, the resource requirements associated with each process, and the likelihood of patients with particular characteristics being treated via distinctive sets of processes (trajectories).
- Patients arrive at irregular times during the simulation. Patients may be described as simply as by their number or may include demographic characteristics that can influence their treatment trajectory probabilities. Internally, the Treater tracks the progress of patients through their individual trajectories and derives the consequent resource requirements.
- At any time, the Treater can be interrogated to determine resource use information.

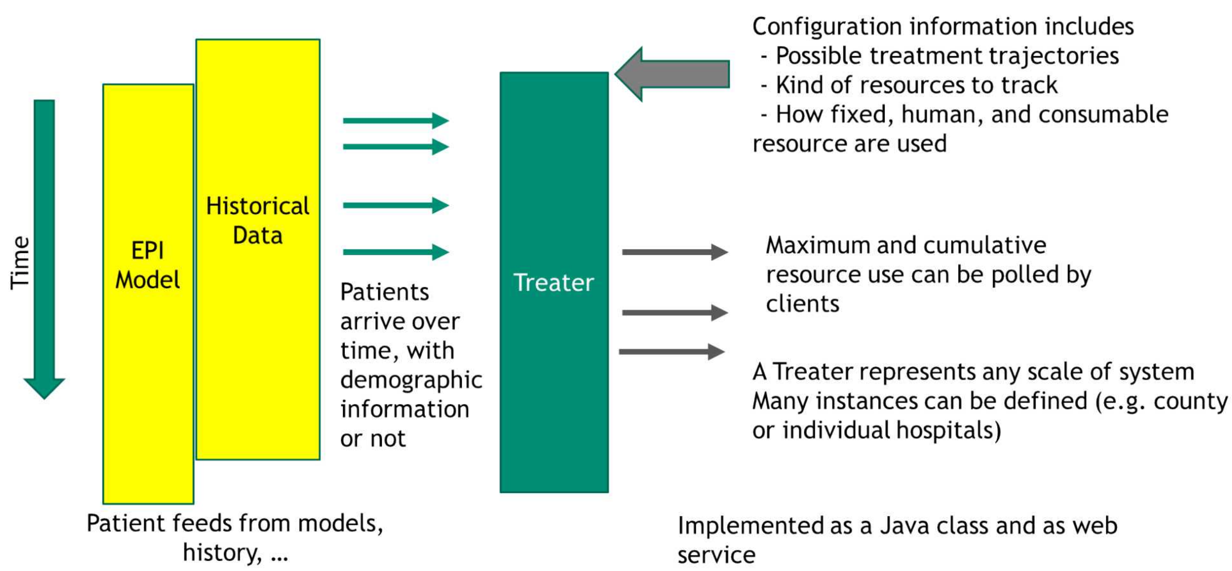

Figure 1. Depiction of Some Treater Interface Flows

Treaters accumulate the resource requirements implied by patient flows and the practices described in the configuration scripts. They do not model responses to resource constraints. In application to the Covid-19 response assessment, Treaters signal the potential need for additional resources to support nominal levels of patient care. Directing those resources to places that need them involves integrating the demands derived by the Treaters with possible sources of supply. Doing this well poses a complicated logistical and forecasting problem. It is therefore important to be able to assess different algorithms or heuristics that might manage resource flows.

The Hospital Resource Supply model adds those elements and defines a framework for their effective interaction. Figure 2 shows the top-level objects considered in the framework, along with illustrative implementations:

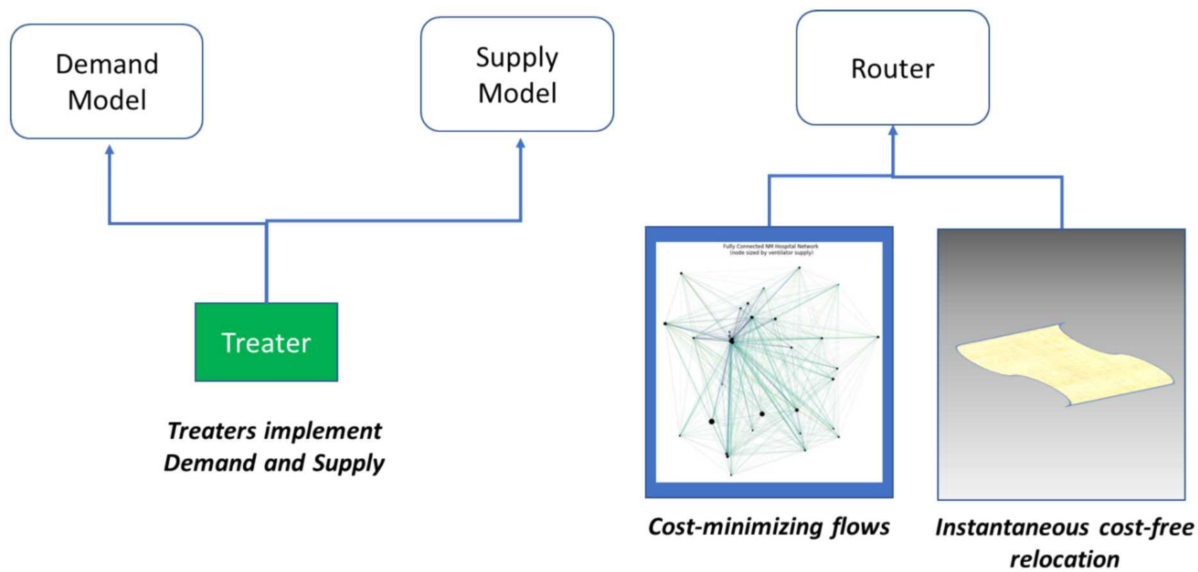

Figure 2. Top-level Objects in the Supply Model Framework with Example Instances

Demand Models represent consumers of resources. The Demand Models anticipated in application are Treaters, however the framework design is open to other possible resource destinations (e.g. foreign aid). Supply Models are potential sources of resources. Although resource inventories are not considered in deriving requirements and simulating treatment processes (as discussed above in Section 1), hospitals do maintain inventories and can share them within and across hospital systems. This aspect of hospital behavior is captured in the Supply Model. Supply Models can also be warehouses with static inventory and producers, whose inventory increases according to an anticipated production schedule.

Routers mediate information and material flows among Demand Models and Supply Models, as detailed below. Various constraints and assumptions can be adopted in order to understand what might be needed to meet aggregate resource constraints. Logistical constraints can be explicitly considered and managed through an optimization approach, for example, as detailed in Section 3 below. In order to simply test the ability of any routing process to reconcile demand and supply, it is useful to posit an ideal routing process that instantaneously affects routing of any surplus resource to points of excess demand.

2.2. Details and Illustration

The Integration Framework manages the flows of information and resources among a collection of otherwise independent objects. Supply inventories and demand amounts are calculated by individual (possibly independent) models. In general, flows related to each kind of resource are managed separately, with one or more independent routers used to manage flows of each resource. Resources can include durable goods, consumables, healthcare workers and patients. Routers apply strategies to determine movements of resources. The Framework allows resource projections used by routers to differ from the ground truth defined for the system, meaning the overall evaluation can also assess allocation problems that might arise from errors in the router's projection algorithm.

Figure 3 shows a simple illustrative application involving ventilators and masks. Note that Treaters demand multiple types of resources, and so the Demand:Masks and Demand:Ventilators interactions may connect to the same underlying Treater object.

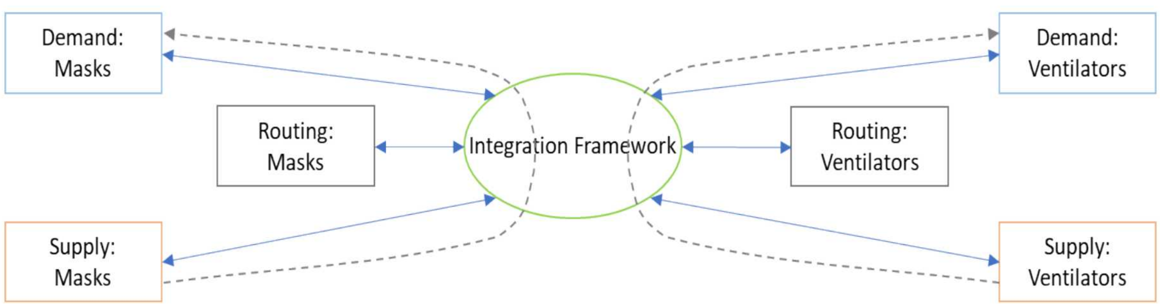

Figure 3 – Illustration of Information and Material Flows Managed

The Framework's operation can be clearly illustrated through an example sequence of interactions. The following figures depict the flow of information and ventilators mediated by a Router that builds internal forecasts of future demand and supply as part of its optimization process.

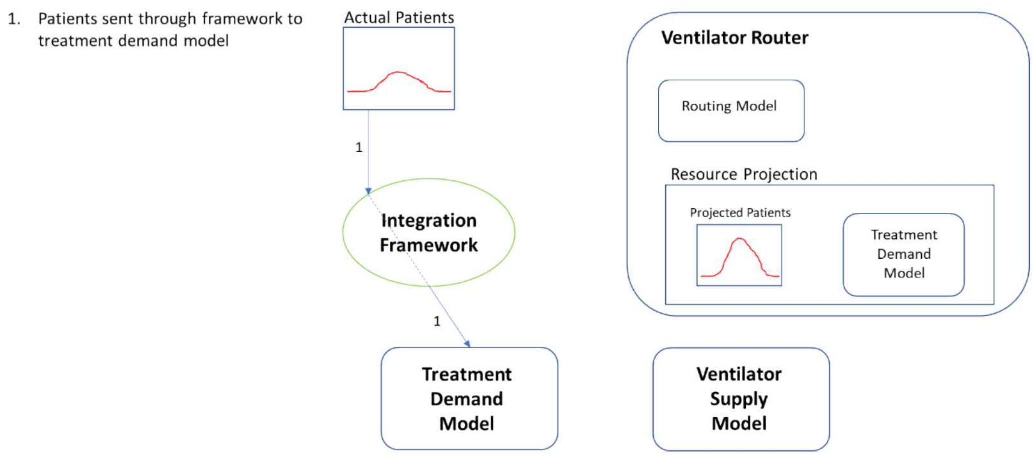

Figure 4 - Framework Illustration, Step 1

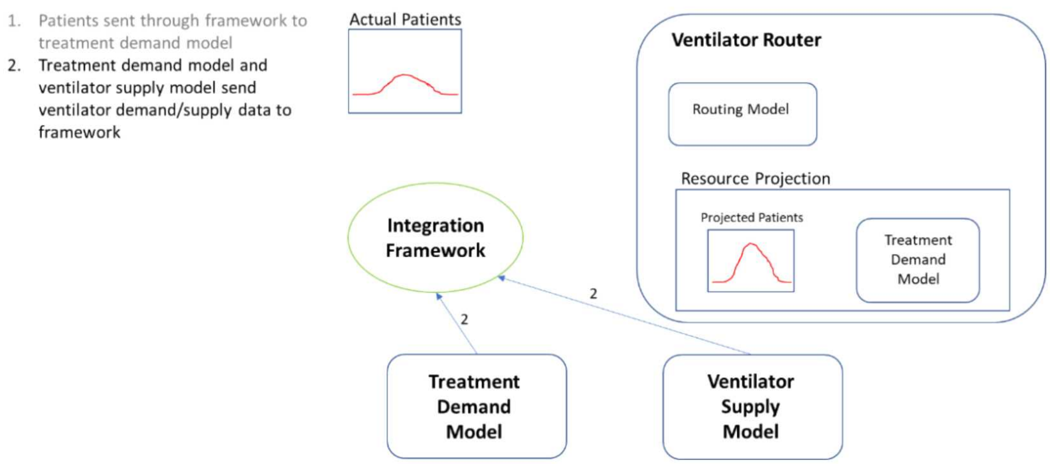

Figure 5 - Framework Illustration, Step 2

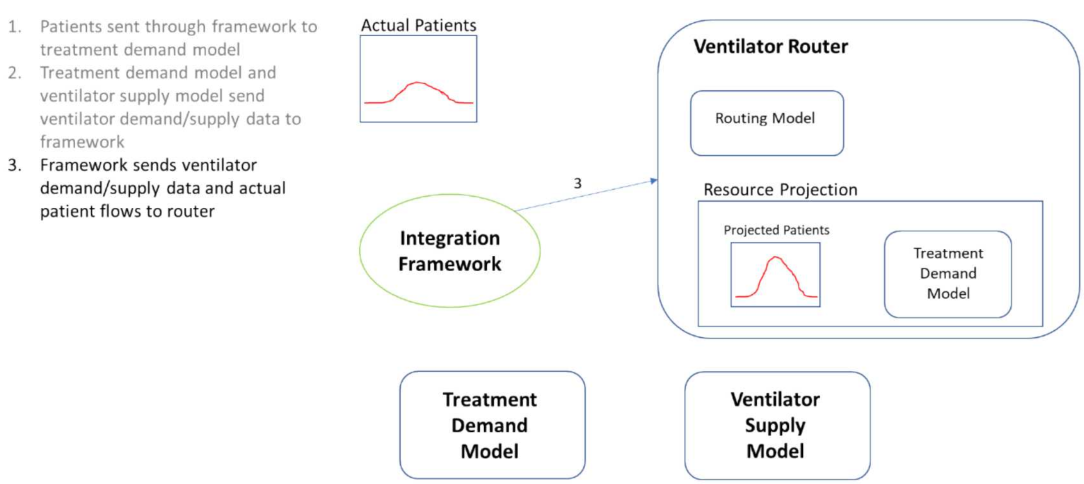

Figure 6 - Framework Illustration, Step 3

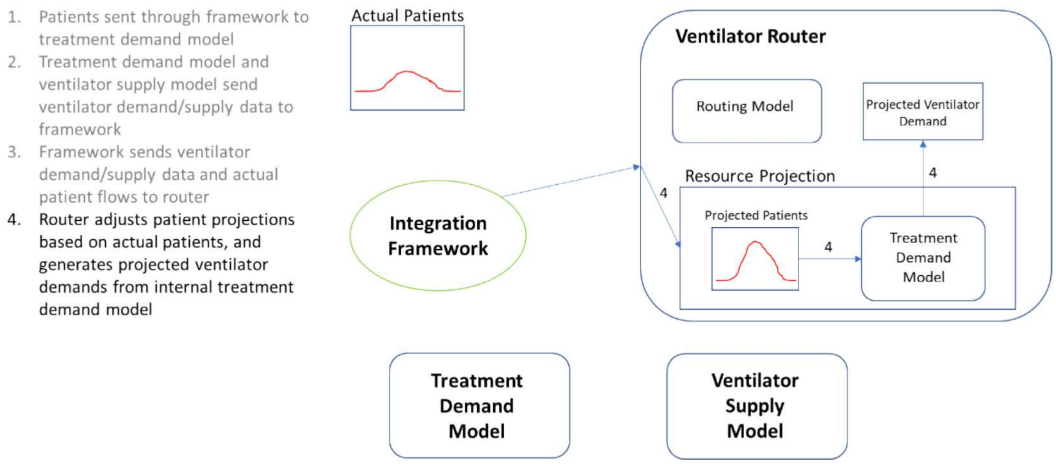

Figure 7 - Framework Illustration, Step 4

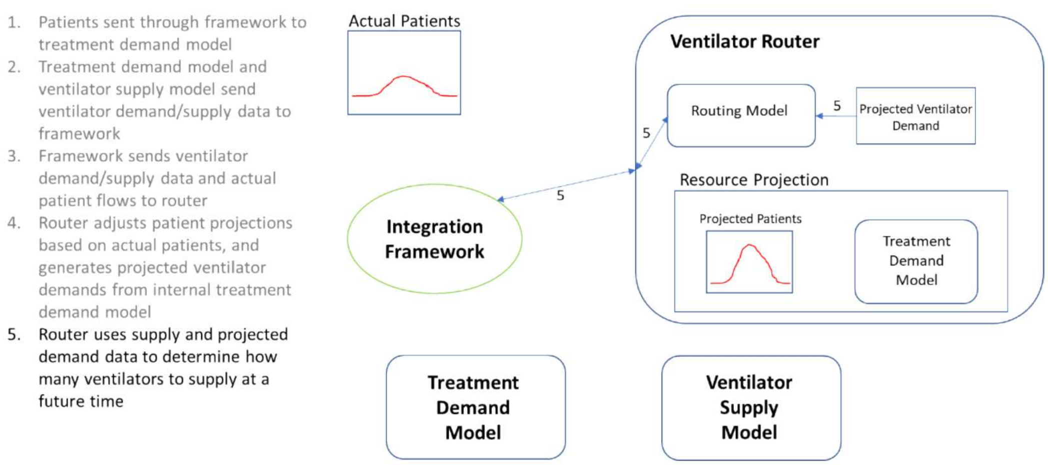

Figure 8 - Framework Illustration, Step 5

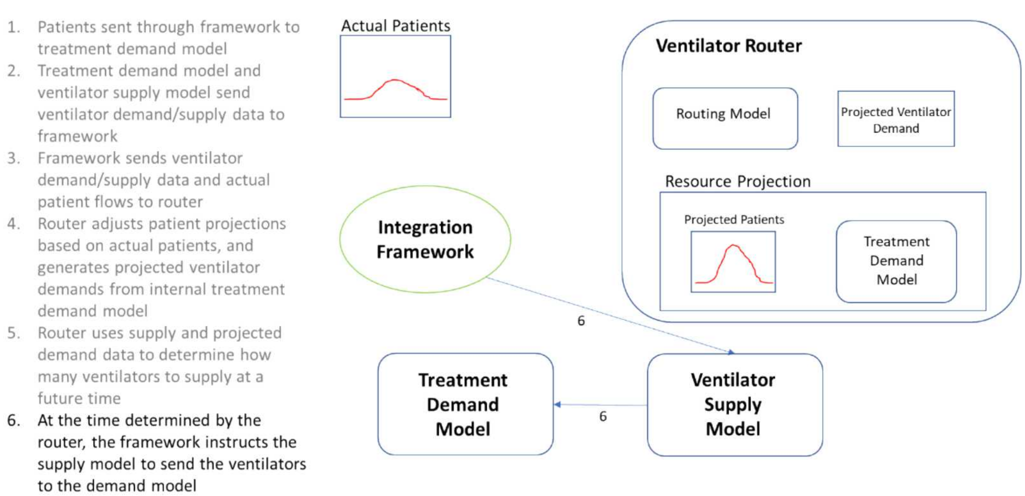

Figure 9 - Framework Illustration, Step 6

The Framework was implemented and exercised in a trial application involving allocation of ventilators among US states based on alternative patient loads in combination with demand projections that matched, overshot, or undershot those projections.

The application focuses on ventilators and uses a network-based optimization model (described below in Section 3) to design inter-state ventilator transfers based on inventories and projected needs. All inventories were held by the Treaters used to represent individual states, and no production was included.

The network optimization model uses distance as a cost measurement, so ventilator moves between closer locations are preferred. The optimization is run every 10 days, and both current state information and forecasted demand over the following 30 days are used to recommend a pattern of ventilator moves. The model will attempt to group its ventilator moves into one day 10 days in the future. When that day comes, ventilators are moved, and the optimization is run again.

The framework allows for the projected patient flows used by routers to differ from the actual patient flows arriving at Treaters in the future. This allows the performance of the router algorithm as a whole, including its forecasts of future requirements, to be assessed so that its performance in the real world can be more accurately gauged. For example, the projections used to make routing decisions might use patient flows from a "low" epi model, while the actual patient flows are generated using a "high" model. The projection process can be recalibrated based on observed patient flows and prior projections about those flows.

The application considered the continental United States (48 states + Washington D.C.), and used patient arrival rates derived from the upper and mean projections published by IHME on April 7, 2020 to define both the actual patient arrival streams as well as the projections used in the optimization. The two patient arrival streams and their use as both actual and projected conditions created four experimental cases. These four experiments, and the associated deficits experienced over the simulation period, are summarized in Table 1 below.

Patient Streams		**Total Supply**	**Maximum Demand**	**Largest Deficit**	**Total Supply-Peak Demand**
Actual	**Projected**				
Mean	Mean	28472	15975	0	12497
Mean	Upper	28472	15975	-1136	12497
Upper	Mean	28472	45168	-22796	-16696
Upper	Upper	28472	45168	-16723	-16696

Table 1. Ventilator Deficits Observed for Different Combinations of Actual and Projected Cases

The total supply is how many ventilators are in the system. The maximum demand is the maximum ventilator demand over all of the days in the simulation. The largest deficit is the worst total ventilator shortage, summed over all states experiencing a shortage, over all of the days. The final column provides a theoretical best-case deficit resulting from instantaneous cost-free routing. The optimal routing solution deviates from this in the experiment with shortfalls and perfect foreknowledge (Upper/Upper) because relocations are constrained to occur once every 10 days, and due to internal stochasticity.

The results also suggest, for this experiment, that planning for a "worse" projection may be preferable, for even though there is a deficit if the Mean is actualized, it is far smaller than if the Mean projection is used with the Upper actualization. Whether this holds more generally could be evaluated with a systematic experiment.

Figures 10-12 show ventilator supply and demand over time for three experimental cases. These illustrate how the demand and routing decisions play out over time. Figure 10 shows the demand (orange line), the deficit (unmet demand) (grey line), and the unused supply (ventilators that aren't being used) (yellow line) for the Mean projection, Upper actual scenario:

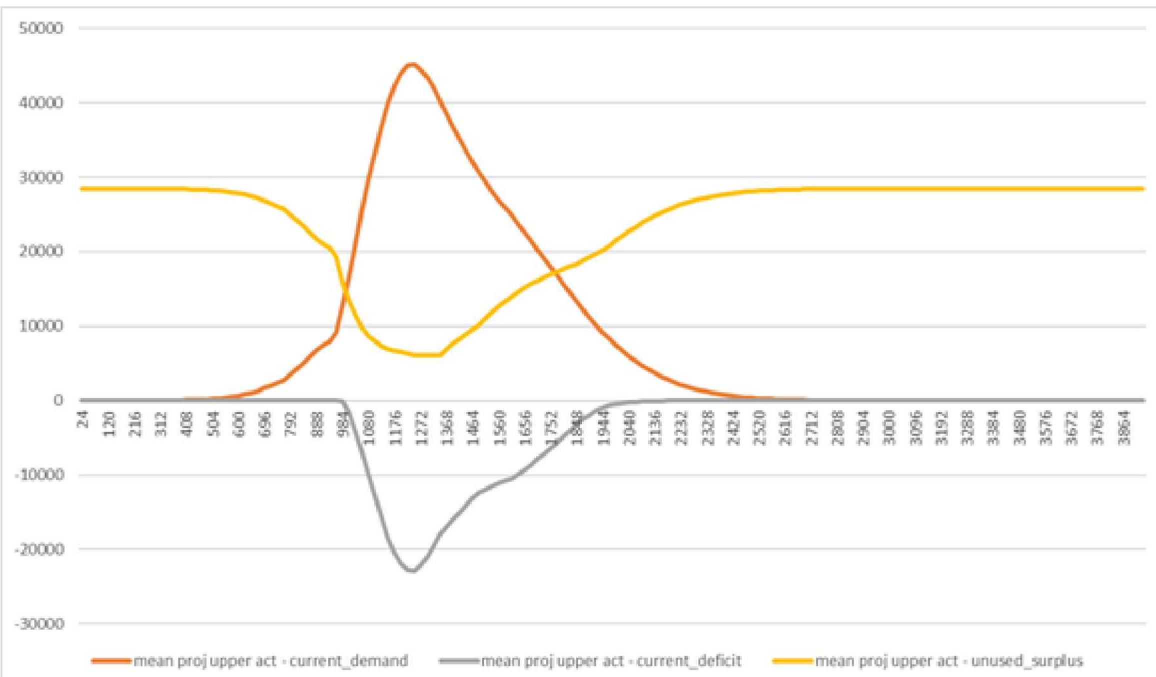

Figure 10 - Demand, Deficit (unmet demand), and Unused Supply for the Mean projection, Upper actual experiment

Figure 11 shows the same variables for the Upper projection, Upper actual experiment, and Figure 12 shows them for the Upper projection, Mean actual experiment.

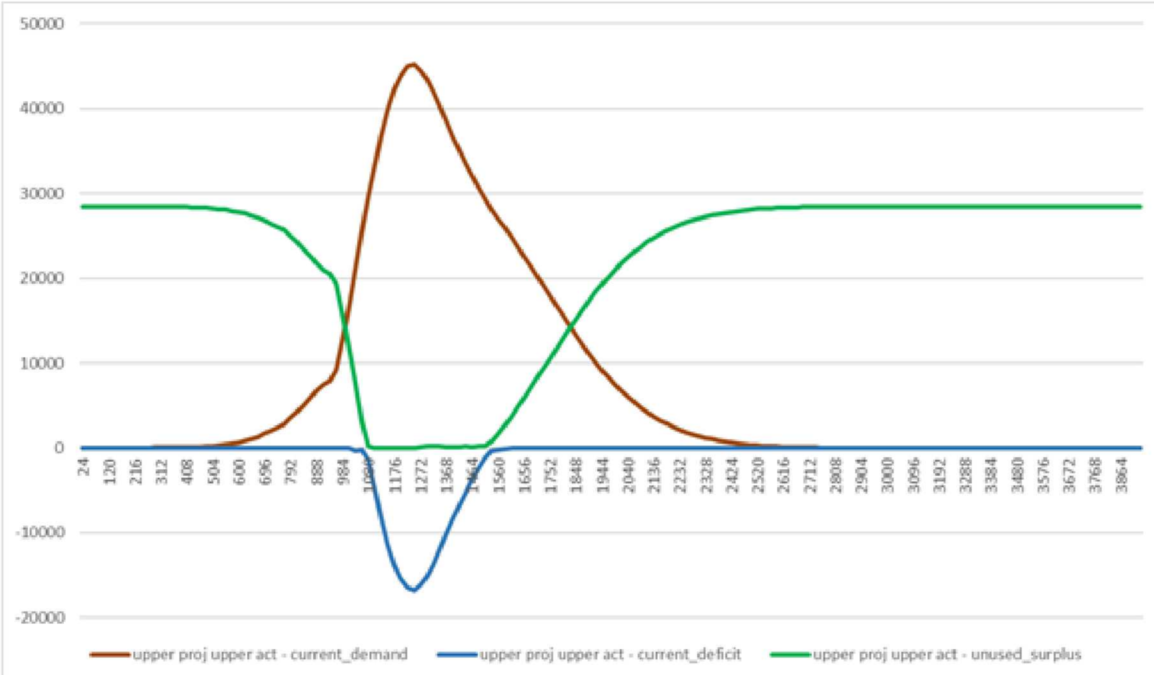

Figure 11 - Demand, Deficit (unmet demand), and Unused Supply for the Upper projection, Upper actual experiment

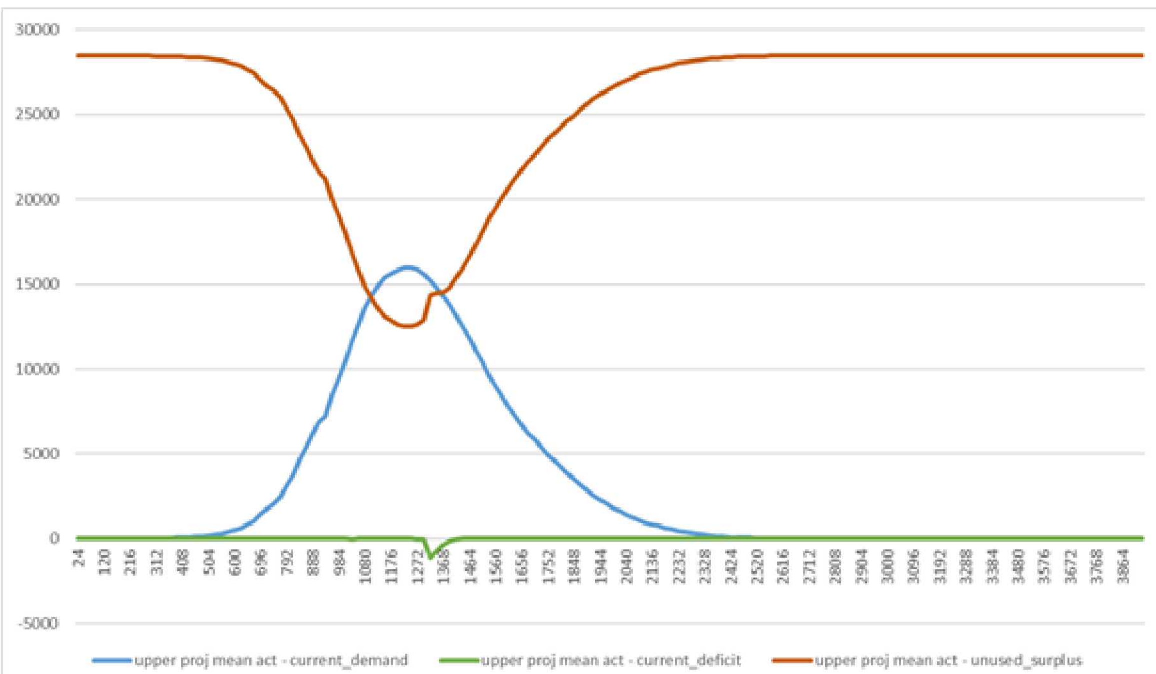

Figure 12 - Demand, Deficit (unmet demand), and Unused Supply for the Upper projection, Mean actual experiment

3. ROUTING ALGORITHMS

The framework is designed to allow interoperability among different implementations of Demands, Supplies, and Routers. While Treaters are likely to be the only kind of Demand objects, many kinds of Supply and Router instances can be anticipated. Supplies for example might come from strategic stockpiles, characterized by an initial inventory and criteria for release, while manufacturers' inventories can accumulate over time, and may depend on access to upstream suppliers. Routers can also be expected to have different subtypes depending on their scope of operation, the kind of information available to them, and the processes they consider in modeling resource relocation.

The bounding case of instantaneous cost-free transport is useful for screening purposes because it distinguishes cases in which routing of existing (and anticipated) resources could conceivably meet projected demand from cases in which more supply or less demand is (also) required. Besides this trivial implementation, two optimization algorithms have been developed to help bridge gaps between demand and supply.

Klise and Bynum (2020) have developed an approach for evaluating alternative locations for facility expansion based on minimizing aggregate patient access cost. It has been applied to analyze alternative sites for adding hospital capacity in New Mexico (Figure 13 below). This formulation might also interface with the framework in two ways: as a Router of patients to Treaters based on cost minimization, and as a Router of new bed capacity to new Treater locations.

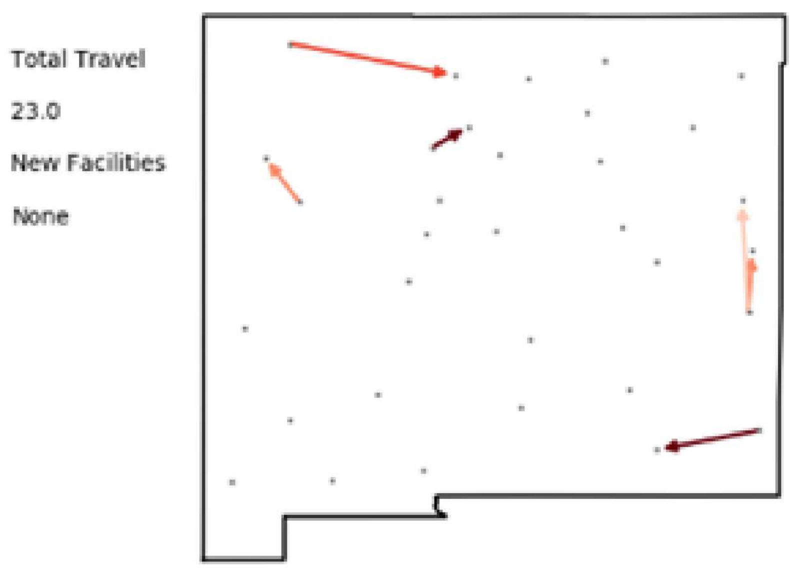

Figure 13 - Possible Hospital Expansion Locations to Minimize Patient Travel under a Hypothetical Surge Scenario

Frazier et al. (2020) have implemented a general resource-routing algorithm for minimizing unmet demand by relocation over network characterized by edge costs. The algorithm uses both current and forecasted demands to plan relocations and can constrain relocation moves to occur only at certain intervals in order to capture practical logistical considerations.

4. SUMMARY

The urgent need to provide critical care to hospitalized Covid-19 patients has placed extraordinary strain on the US healthcare system. As an aid to managing that strain, Sandia has developed a modeling capability that integrates resource use, production, and allocation. This capability has been designed to help identify and alleviate prospective resource shortfalls. By exercising this system model over a range of possible conditions and contingencies, decision-makers can characterize the risk of shortfalls as a function of time and location and find mitigations that represent tolerable trade-offs among desired outcomes.

Comment [DS]: Update when available

REFERENCES

1. Swiler,L. P., T. Portone, W. E. Beyeler. "Uncertainty Analysis of Hospital Resource Demand Model for Covid-19," SAND2020-49004900, Sandia National Laboratories, Albuquerque, NM.
2. Beyeler, W. E., C. Frazier, D. Krofcheck, L. Swiler, T. Portone (2020). "Treatment Model Design and Use." SAND2020-4944 PE. Sandia National Laboratories, Albuquerque, NM.
3. IHME COVID-19 health service utilization forecasting team & Murray, C. J. (2020). "Forecasting the impact of the first wave of the COVID-19 pandemic on hospital demand and deaths for the USA and European Economic Area countries." *medRxiv*. doi: https://doi.org/10.1101/2020.04.21.20074732
4. Klise, K. and M. Bynum (2020). "Facility Location Optimization Model for COVID-19 Resources". Unpublished white paper
5. Frazier, C., D. Krofcheck, D., J. Gearhart, W. Beyeler, Integrated Resource Supply-Demand Routing Model for the COVID-19 Crisis, SAND2020-5569

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Walter E. Beyeler	6617	webeyel@sandia.gov
Justin Garretson	6617	jrgarr@sandia.gov
Sean DeRosa	8721	sderosa@sandia.gov
Patrick D. Finley	8721	pdfinle@sandia.gov
Daniel J. Krofcheck	8721	djkrofc@sandia.gov
Chad E. Davis	8721	cedavis@sandia.gov
Christopher R. Frazier	8721	crfraz@sandia.gov
Dean A. Jones	8720	dajones@sandia.gov
Laura P. Swiler	1463	lpswile@sandia.gov
Daniel Z. Turner	1463	dzturne@sandia.gov
John T. Feddema	1460	jtfedde@sandia.gov
S. Scott Collis	1400	sscoll@sandia.gov
M. Daniel Rintoul	1462	mdrinto@sandia.gov
Teresa Portone	1463	tporton@sandia.gov
Jaideep Ray	8759	jairay@sandia.gov
Technical Library	01977	sanddocs@sandia.gov

This page left blank

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.