Multi-criteria Course Mode Selection and Classroom Assignment Under Sudden Space Scarcity

Mehran Navabi-Shirazi, Mohamed El Tonbari,
Natashia Boland, Dima Nazzal, Lauren N. Steimle
H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA
{navabi, mtonbari, natashia.boland, dima.nazzal, steimle}@gatech.edu

Problem Definition: While physical (or 'social') distancing is an important public health intervention during airborne pandemics, physical distancing dramatically reduces the effective capacity of classrooms. During the COVID-19 pandemic, this presented a unique problem to campus planners who hoped to deliver a meaningful amount of in-person instruction in a way that respected physical distancing. This process involved 1) assigning a mode to each offered class as either remote, residential (in-person) or hybrid, and 2) reassigning classrooms under severely reduced capacities to the non-remote classes. These decisions need to be made quickly and under several constraints and competing priorities such as restrictions on changes to the timetable of classes, trade-offs between classroom density and educational benefits of in-person vs. online instruction, and administrative preferences for course modes and classrooms reassignments.

Methodology and Results: We solve a flexible integer program and use hierarchical optimization to handle the multiple criteria according to priorities. We apply our methods using actual Georgia Institute of Technology (GT) student registration data, COVID-19 adjusted classroom and lab capacities, and departmental course mode delivery preferences. We generate optimal classroom assignments for all GT classes at the Atlantacampus, and quantify the trade-offs among the competing priorities. When classroom capacities decreased to 20 - 25% of their normal seating capacities, optimization afforded students 15.5% more in-person contact hours compared to no room re-assignments (NRR). Among sections with an in-person preference, our model satisfies 87% of mode preferences while only 47% are satisfied under NRR. Additionally, in a scenario in which all classes are preferred to be delivered in-person, our model can satisfy 90% of mode preferences compared to 37% under NRR.

Managerial Implications: Multi-objective optimization is well-suited for classroom assignment problems that campus planners usually manage sequentially and manually. Our models are computationally-efficient and flexible with the ability to handle multiple objectives with different priorities, build a new class-classrooms assignment or optimize an existing one, and can apply under normal or sudden capacity scarcity constraints.

Key words: classroom assignment, scheduling, hierarchical optimization, multi-objective, higher education

1. Introduction

Despite the global COVID-19 pandemic that (as of April 2022) has killed more than 900,000 people in the United States and 6.2 million people worldwide (Dong et al. 2020), colleges are under enormous pressure to operate with administrators citing concerns of declining future enrollment and inequitable impact on underrepresented students (Doug 2020). Due to these concerns, most US colleges spent Summer 2020 grappling with decisions on how to open their campuses safely in the Fall 2020 semester. Decisions from these administrators collectively affect 19.9 million students who attend colleges and universities in the United States alone (National Center for Education Statistics 2016).

Due to the pandemic, administrators face critical decisions about whether classes will be offered in-person, remotely using online platforms, or in a hybrid format that delivers some course content in-person and some content online. One fundamental trade-off in these decisions is the balance between the quality of education and the health of students, staff, and faculty. In the absence of risks to public health, the latest research on education supports classroom environments that facilitate discussions to enhance critical thinking and communication skills (Freeman et al. 2014). However, the benefits of having students engaging with their instructor and peers in the classroom are in direct conflict with the research on the COVID-19 pandemic, which has shown that transmission of the coronavirus is highest when people are sitting indoors for a long period and talking (de Oliveira et al. 2021). Reducing in-person instruction time may lead to reduced risk of transmission in classrooms, but may also decrease the quality of discussion and interaction among classmates.

Further, the COVID-19 pandemic created logistical challenges when deciding how to deliver instruction to students due to the sudden scarcity in two critical resources: space and instructional staff. Most college campuses suffer from classroom space constraints even during non-pandemic times. The Center for Disease Control and Prevention (CDC)'s recommended 6-feet physical (or 'social') distancing (Savage et al. 2020) reduces the classroom capacity down to 25 or even 20 percent of its original seating capacity (Lederman 2020). Moreover, some instructional staff were in high-risk groups and were exempted from teaching in-person courses. This sudden reduction in these resources creates a severe mismatch between the supply and demand for classrooms, and limits the number of students that can receive in-person instruction at any given time.

In this article, we propose a flexible hierarchical optimization framework to help colleges balance competing objectives when re-assigning classes to classrooms due to sudden reduction in classroom capacity. The formulations we propose were motivated by our experience working with the Coronavirus Campus Recovery Task Force at the Georgia Institute of Technology (GT).

1.1. Motivation for this study: The impact of COVID-19 at GT

This study was motivated by our experience at our own university. In this section, we describe the context for our analytical modeling. Throughout our modeling process, we worked closely with the Registrar and took steps to address contextual complicating factors such as working within GT's legacy computer systems and ensuring implementable solutions (Gorman 2021).

In June 2020, GT announced its plans for Fall 2020 which allowed students to return to campus for coursework (Rogers 2020). The re-opening plan specified that courses would be offered in four distinct modes:

- *Residential Spread*: In-person, with physical distancing.
- *Hybrid Split:* Lectures in-person, with physical distancing; students attend on a rotating basis.
- *Hybrid Touch Point*: Lectures in an online format; bring students to the classroom several times during the semester for meaningful in-person experiences.
- *Remote*: Fully remote delivery.

The academic departments, with guidance from the Task Force, determined the preferred mode of each course, and the Registrar was left to reallocate classroom space among the non-remote classes, under severely reduced capacities. Three factors determine the modality for a course, as follows. 1) Classroom space: any course that was assigned in-person had to ensure that the size of its assigned classroom is sufficient to accommodate the projected enrollment of the course. This presented a challenge as the largest classroom – after adjustment for physical distancing – was around 140 seats, but 95 classes out of 2249 (4.2 %) had enrollments greater than 140 students. 2) Instructor health accommodation: some faculty members and teaching assistants had risk factors as defined by the CDC and will only deliver classes in Remote mode. 3) Pedagogical considerations: some courses (e.g., chemistry labs and hands-on capstone design projects) lose substantial pedagogical value if delivered remotely, while other courses might be more suitable for online delivery.

GT Scheduling Process: GT has a two-phase registration system. A Fall semester schedule of classes is created in January then released in the Spring of that year. Current students register during Phase I, which occurs in March through early May. Students have an opportunity to make changes to their schedules during Phase II, which spans the two weeks before the Fall semester starts in August. Incoming new students register on a predetermined schedule during the summer semester after orientation events (See https://registrar.gatech.edu/calendarfor details of GT's academic calendar).

Classrooms: GT has classrooms that are controlled by the Registrar (hereafter, *Registrar-controlled* classrooms), and classrooms that are controlled locally by the departments (hereafter, *department-controlled* classrooms). The Registrar's office assigns classes to the Registrar-controlled classrooms based on projected class size, while taking into account departmental requests for specialized classrooms. Large classrooms are a scarce resource even during non-pandemic times, and therefore the Registrar has a process in place for departments to bid for these large classrooms, months ahead of the release of a schedule to the students. Most departments have – under their control – a set of department-controlled classrooms of various sizes that are located in the department's building and serve as specialized classrooms for that department's needs, such as seminars, design studios, special topic classes, and meetings.

COVID-19 Implications for Classroom Assignments: When the US went into lockdown in Spring 2020, current GT students were actively registering for the Fall 2020 semester. When GT started re-planning in May and June 2020 for the Fall 2020 semester, thousands of current GT students had already selected their fall classes, and classrooms had been assigned to classes based on normal seating limits. The need for physical distancing to maintain a 6-foot radius around each seat reduced classroom seating capacities by 70-80%. Rather than recreate a schedule to reflect physical distancing constraints, GT opted to keep the schedule as-is and did not cancel any student registrations to minimize the stress on students inflicted by having to re-register. That decision left the academic departments and Registrar with two levers to accommodate physical distancing within classrooms: 1) change course section delivery mode based on its enrollment, instructor health accommodations, or pedagogical factors. For example, a department may want to ensure that courses requiring significant in-person interaction are taught in Residential Spread, while a course section for which the instructor has an accommodation to teach online should be taught in Remote mode. 2) Re-assign classrooms among the non-remotely delivered sections, but limit the number of reassignments to reduce disruptions to suitability of a classroom to a specific course, or reduce the logistical burden associated with reshuffling classes. We distinguish a "course" from a "course section". A course (e.g., ISYE 3030 Basic Statistical Methods) might have multiple sections delivered by different instructors on different times and days. Each of this course's sections can have its own delivery mode preference, either because the instructor of a particular section has a health accommodation, or the department has determined to offer a variety of modes to its students.

Two of the authors served on the GT Coronavirus Campus Recovery Taskforce and saw the value of modeling and analysis of potential solutions to this problem. Optimization models are well-suited for this problem due to its scale and complexity, especially with all the changing epidemiological facts and COVID-19 guidelines. Having a model that offered quick solutions was critical to re-evaluate potential modes and assignments. This paper discusses the methodology, results, and insights from classroom assignments and pandemic planning at GT.

1.2. Contributions

In this article, we present a new modeling approach for simultaneously assigning course modes and assigning sections to classrooms when classroom capacities are suddenly reduced. The advantages of our methodological framework are summarized as follows.

- **Practical:** Our methodological approach works within a university's existing timetable. Although reworking a university's timetable may lead to better solutions, the time and cost required to change a timetable on short notice is often prohibitive. Our practical approach makes the most of the existing timetable without the need to overhaul it.
- **Multi-criteria:** Our hierarchical optimization approach considers different objectives with different and changing priorities that were identified in the process of working with the Registrar's office: stability, course mode preferences, and in-person contact hours. These criteria were identified to address the needs of the GT Coronavirus Task Force and Registrar.
- **Flexible:** Our approach is flexible in that it always provides a feasible solution. If the preferred course modes cannot be satisfied due to space constraints, the output would recommend a different course mode while optimizing the objective in question. Additionally, the model can adapt to requirements on the minimum number of times that a student attends a course in order for it to provide a meaningful amount of "in-person" experience for hybrid courses.
- **Scalable:** We use the methodological framework above to study course mode assignments and classroom assignments at GT in Fall 2020. The data cover all undergraduate and graduate courses to be offered in Fall 2020, excluding courses that are not taught on

the main campus, asynchronous distance learning, and MBA programs. The final model included 2,249 course sections and 549 classrooms, and considered 204 time slots.

Application of our model to the GT Fall 2020 class schedule disruption due the COVID-19 pandemic led to the following managerial insights.

- **Using optimization to reassign sections to classrooms after sudden capacity reductions leads to significant increase in satisfied course mode preferences. In our data set, optimizing reassignments led to 215% increase in Residential Spread sections compared to keeping classes in their original rooms.** Furthermore, the resulting assignment results in 82% of the maximum amount of in-person contact hours (i.e., if there was no reduction in room capacity for Hybrid and Residential Spread preferred classes) compared to 71% when classes are not reassigned to classrooms.
- **The introduction of a Hybrid Touch Point mode can increase the number of mode preferences satisfied and amount of in-person contact hours. However, these gains are negligible relative to assignments done without considering Hybrid Touch Point mode.** The introduction of a Hybrid Touch Point mode allows the Registrar some flexibility because courses with the preference "Hybrid" can be satisfied either through a Hybrid Split or a Hybrid Touch Point mode assignment. However, this added flexibility does not substantially increase the amount of contact hours provided nor does it substantially increase the amount of mode preferences that can be satisfied. Therefore, if the Registrar would like to ignore the consideration of a Hybrid Touch Point mode and just keep Hybrid Split mode, we would not expect to see a large drop in performance.
- **Centralized scheduling can lead to substantial gains in number of in-person contact hours.** Compared to decentralized scheduling in which departments do not share their own classrooms, centralized scheduling can result in an increase of 10.9% in the in-person contact hours available to students through Residential Spread or Hybrid classes.
- **Capacity reductions of more than 50% lead to rapid decline in contact hours.** If the adjusted capacity remains over 50% of the full capacity, 100% of all non-Remote preference classes can be satisfied and more than 90% of the maximum amount of in-person contact hours for these non-Remote classes can be delivered. However, when the adjusted capacity drops to 25%, 65% of the maximum contact hours can be delivered, if classrooms can be reassigned and centralized scheduling is enforced.

- **Decreasing the amount of in-person instruction delivered can significantly decrease the distance by which classes are relocated.** Although maximizing in-person instruction hours is a strong priority, departments have also conveyed a strong desire to minimize relocation of sections. Our work shows that we can significantly reduce the relocation distance by slight sacrifices to other objectives. For example, in our experiments, we can reduce distance relocation by two orders of magnitude at a less than 1% sacrifice in mode preferences and about 10% sacrifice in contact hours.

1.3. Organization of the paper

The remainder of this article is organized as follows. In Section 2, we provide background on course scheduling, room assignment, and timetabling for university classes, and discuss the related literature. In Section 3, we propose optimization models for assigning course modes and reassigning classrooms due to reduced capacity, and in Section 4, we present our hierarchical optimization approach. In Section 5, we demonstrate our formulations using a case study of GT's Fall 2020 schedule, and in Section 6, we describe the contributions of our optimization modeling on the decision-making process at GT. Finally, in Section 7, we summarize our most important findings, discuss the limitations of our work, and present opportunities for future research.

2. Literature Review

The larger problem of university course scheduling has been thoroughly studied in a variety of settings. University course scheduling consists of two key sub-challenges. Planners must determine (1) when courses should be taught (timetabling), and (2) where each course should be taught (room assignment). In the literature, course timetabling and room assignment are either tackled simultaneously or independently.

Lach and Lübbecke (2008) state that solving both problems simultaneously in one large IP is likely to be computationally impractical for larger universities, and instead decompose the problem into two stages, by first scheduling the courses and then assigning courses to rooms. The authors exploit the structure of their constraints, generating facets of a partial transversal polytope to obtain tight cuts. Their decomposition method solves the timetabling and room assignment problems simultaneously and exactly. The authors only consider one objective, however, defined by the professors' preferences for each feasible pair of a course and timeslot. Several (meta)heuristics have been proposed in the literature to address both problems simultaneously. Ueda et al. (2000) consider a two-phase genetic algorithm, where two populations are used, one for scheduling classes and one for allocating them to rooms. Kostuch (2004) use a simulated annealing procedure to handle timetabling and room allocations. The authors also take student enrollments into account, so that no two courses in which a student is enrolled can be scheduled at the same time. This is known as the post-enrollment stage of the problem.

There are two main reasons for approaching timetabling and room assignment independently. First, the administrative process for timetabling and room assignment are often managed by different administrative departments and performed disjointly. Secondly, even in cases where it is of administrative interest to solve both problems simultaneously, it is helpful to solve these problems independently for tractability . Although considering the two aspects of course scheduling independently greatly reduces complexity, tractability is still a fundamental challenge. For a fixed course schedule, Carter and Tovey (1992) state the conditions under which the room assignment problem is polynomially solvable or is NP-hard, where the former holds if the time windows for sections are non-overlapping. In the latter, more practical case, integer programs (Waterer 1995, Phillips et al. 2015) and heuristics (Mulvey 1982, Gosselin and Truchon 1986, Glassey and Mizrach 1986) have been proposed.

While most studies focus on a single objective, more recent papers have considered multiple objectives. Given a ranking of the objectives, Barnhart et al. (2021) propose a hierarchical optimization approach to tackle the problems of timetabling and room assignment together, as well as term planning. The authors consider whether their objective functions would benefit from a shift to a three semester calendar. The paper addresses the problem of university course scheduling under constraints posed by COVID-19 and its physical distancing requirements. Phillips et al. (2015) propose a flexible integer program to analyze the room assignment problem, and consider hierarchical objective functions as well. The objectives considered include event hours, seated student hours, seat utilization, room preference, course room stability, and sparse seat robustness. Some of these objectives are similar to objectives considered in our study, but are calculated differently. For instance, their objective of "seated student hours" is conceptually the same as our "contact hours", but since we consider the possibility of courses being taught in different modes, our objective functions adapt correspondingly.

Our work considers course scheduling in the context of sudden reductions in room capacity due to a pandemic. While several studies have explored the risk associated with moving classes online during the COVID-19 pandemic (Lopman et al. 2021, Gressman and Peck 2020, Weeden and Cornwell 2020), there has been limited consideration of whether the proposed strategies are feasible given the sudden capacity drops associated with the CDC's recommended 6-feet physical distancing (Savage et al. 2020). Others have considered the impact of sudden capacity reductions due to a pandemic on bus scheduling on university campuses, another important consideration for campus operations, using optimization and simulation techniques (Chen et al. 2020). To our knowledge, Barnhart et al. (2021) is the only paper that examines course scheduling in the context of sudden scarcity in room capacity due to a pandemic and our work adds to this area of research. While Barnhart et al. (2021) adopt a more granular approach and focus on tackling the larger problem of timetabling and room assignment simultaneously to a single department, we focus on scalability and the ability to handle sections across all departments in a university, where we are given an existing timetable and solve the room assignment problem. The goal is to be able to quickly solve the room assignment problem and give planners the ability to run what-if scenarios. Moreover, we expand on the room assignment problem and implicitly determine the mode a section is taught in (e.g., Residential Spread vs Hybrid Split) through our optimization model, which differs from past studies that treat the mode of a section as an input parameter.

3. Classroom and Course Mode Assignment Formulations

We now formally introduce the classroom and course mode assignment problem (CCMAP). We present integer programming formulations to help campus administrators simultaneously assign classrooms and course modes to satisfy various administrative preferences. As mentioned previously, we assume that all course sections are delivered in their previously assigned time slots due to existing registrations and ease of administrative implementation.

3.1. Preliminaries: CCMAP under reduced capacity

If a section is delivered entirely remotely, there is no need to assign a room for that section. If a section is delivered partially in-person (i.e., in Hybrid Split, Hybrid Touch Point, or Residential Spread mode), then the section needs to be assigned a classroom. Thus, given a set of sections and enrollments, CCMAP is defined as assigning sections to both a room and delivery mode, where we optimize over one or multiple objectives determined by administrative preferences.

If the section $x \in \mathcal{X}$ is assigned to a room $r \in \mathcal{R}$, the delivery mode of x is determined by a relationship between the assigned classroom's capacity n_r and the section's enrollment p_x. The mode also depends on other specifications, such as the number of times the section x is scheduled to meet each week (denoted m_x), the number of weeks in the semester (W), and the minimum number of "touch points" (S) required for Hybrid Touch Point courses. A touch point is defined as the number of times a student must be able to attend a class in person over the semester for the section to not be deemed fully remote, with $1 \leq S \leq W$. The user may select any feasible value of S that is most appropriate for their administrative needs. The assigned "mode" for a section x taught in room r is determined by the following logic.

- $p_x \leq n_r$: "Residential Spread", meaning that the section is held in-person with all students attending every class.
- $n_r < p_x \leq m_x n_r$, $m_x = 2, \ldots, 5$: "Hybrid Split", meaning that the section is held in person and students attend at least one class per week. When $m_x = 2$, students attend exactly one class per week. When $m_x = 3$, if $p_x \leq 2n_r$ each student attends 2 classes per week ; otherwise each student attends 1 class per week. Similarly, if $(k-1)n_r < p_x \leq kn_r$ for $k = 1, \ldots, m_x$ then each student attends $m_x - k + 1$ classes per week.
- $m_x n_r < p_x \leq W m_x n_r / S$: "Hybrid Touch Point", meaning that much of the class delivery takes place online, but there is scheduled face-to-face time between students and instructors, so that students can touch base with the instructor at least S times during the semester.

An illustration of CCMAP is provided in Figure 1, where we give an example of a feasible solution. We now consider factors that might make one feasible solution to CCMAP more desirable than another in the eyes of campus administration, and important constraints on reassignments that ensure they are implementable.

Each section $x \in \mathcal{X}$ has a mode preference, which is typically based on the instructor's or department's preference. Additional considerations that must be accounted for when assigning class sections to rooms are restrictions on rooms for some classes due to the need for specialized equipment and constraints that prohibit two courses from being assigned to the same room at the same time. We now discuss the details of these restrictions.

First, each section $x \in \mathcal{X}$ has a restricted set of rooms $\mathcal{R}_x$ which it can be assigned to. A chemistry lab section, for instance, may only be assigned to rooms with the appropriate

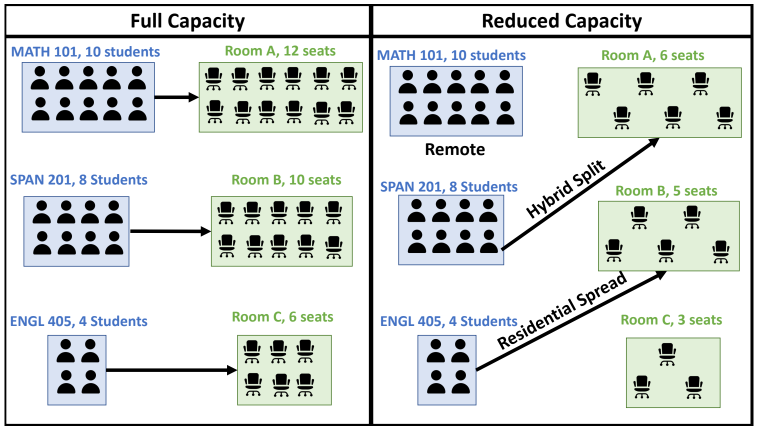

Figure 1 **An illustration of CCMAP for 3 class sections (shown in blue) and 3 rooms (shown in green) when capacity is reduced to 50%. MATH 101, SPAN 201, and ENGL 405 have 10, 8, and 4 students, respectively. Without capacity reductions, MATH 101 is delivered in Room A, SPAN 201 is delivered in Room B, and ENGL 405 is delivered in Room C. After capacity reductions, MATH 101 is assigned Remote mode and thus has no assigned room. SPAN 201 has 8 students and so can fit into Room A with its reduced capacity of 6 seats when delivered in Hybrid Split mode (students alternate attendance). ENGL 405 has 4 students and is assigned to Room B in Residential Spread mode; all ENGL 405 students are able fit in the room at the same time even with Room B's reduced capacity..**

lab equipment. The definition of $\mathcal{R}_x$ is also influenced by room capacity. For a given section x, any room r with $n_r < Sp_x/(Wm_x)$ does not meet the minimum threshold for in-person learning and so is excluded from $\mathcal{R}_x$. Conversely, the set of sections that may be taught in room r for $r \in \mathcal{R}$ is denoted $\mathcal{X}_r$, and these sets are based on the same restrictions. Later, we will consider the effects of *centralized* planning compared to *decentralized* planning, which affects the feasible set of classrooms that are allowed for a given section.

Second, we consider constraints that ensure two sections are not assigned to the same room in timeslots that overlap. Our model makes a few mild assumptions regarding the structure of these timeslots. Each timeslot should consist of a pattern of days, a single start time, and a single end time. Concretely, 11:00-11:50am MWF (Monday, Wednesday, Friday) and 9:30-10:45am TR (Tuesday, Thursday) are each valid timeslots. However, 11:00-11:50am MW 10:00-10:50am F, is not a valid timeslot, as the start time and end time vary based on the day of the week. A course with such a schedule will be modeled as two separate sections, i.e. two different elements in the set $\mathcal{X}$. The first section would have the predetermined timeslot 11:00-11:50am MW, and the second section would have the timeslot F 10:00-10:50am. This does mean there may be a different room assignment for Monday and Wednesday than for Friday. We argue that such timeslots are relatively uncommon for university schedules. When they do occur, it is not unreasonable to ask faculty and students to use different rooms on different days of the week.

Recall there may be multiple sections of the same course, and each section is modeled independently. We do not require that sections of the same course be taught in the same room, in the same mode, nor do we impose any other form of interdependence between sections of the same course. This modeling choice was informed by discussions with the GT Task Force. On the one hand, if different sections of the same class were taught in the same delivery mode, that would give some consistency across sections; on the other hand, if they were taught in different formats, this would allow for students to register for the mode they preferred and at this point in time, students had mixed opinions about in-person vs. hybrid vs. remote instruction (Steimle et al. 2022). Because there seemed to be benefits to both outcomes, the Task Force was not concerned about modeling different sections of the same class independently.

Lastly, we mention that the set $\mathcal{X}$, and the CCMAP in general, only consider sections that would typically be taught in Residential Spread. In other words, sections for courses that are exclusively taught online are not considered in CCMAP, as no further decisions need to be made for such sections.

Table 1 Summary of notation

Set/Parameter	Description
$\mathcal{X}$	A set of class sections, where there might be multiple sections of the same class
$\mathcal{R}$	A set of classrooms
$\mathcal{ST}$	A set of day-time pairs during the week that are a *start time* of some meeting of some class, e.g. if a class has timeslot 11:00-11:50am MWF then the three meeting start times of that class, (M, 11:00am), (W, 11:00am) and (F, 11:00am) are all elements of $\mathcal{ST}$
$\mathcal{R}_x$	The set of rooms suitable for section x to be held in, for each $x \in \mathcal{X}$ (e.g., if the class needs certain lab equipment, it must be held in a room with that equipment)
$\mathcal{X}_r$	The set of sections that may be taught in room r for each $r \in \mathcal{R}$
$\mathcal{X}_{d,t}$	The set of sections that have a meeting on day d that starts before time t on that day and that makes the classroom it is assigned to unavailable for other classes until time t or later, for $(d,t) \in \mathcal{ST}$
$\mathcal{X}_{rdt} = \mathcal{X} \cap \mathcal{X}_r \cap \mathcal{X}_{d,t}$	the set of non-remote sections which can be assigned to room $r \in \mathcal{X}_r$, and which meet on day and time $(d,t) \in \mathcal{ST}$
p_x	the number of students enrolled in section $x \in \mathcal{X}$
n_r	the maximum number of students that can be seated in classroom $r \in \mathcal{R}$

With these considerations in mind, we now define the feasible region of the CCMAP. The required parameters and sets are summarized in Table 1. We define the decision variable X

to be a binary vector, where $X_{x,r}$ is one if and only if section x is assigned to room r, and P to be the set of feasible assignments defined by the following set of constraints:

$$\sum_{x\in\mathcal{X}_{rdt}} X_{xr} \leq 1, \quad \forall r \in \mathcal{R}, \quad \forall (d,t) \in \mathcal{ST} \tag{1}$$

$$\sum_{r\in\mathcal{R}_x} X_{xr} \leq 1, \quad \forall x \in \mathcal{X} \tag{2}$$

$$X_{xr} \in \{0,1\}, \quad \forall x \in \mathcal{X}, \quad \forall \mathcal{R}_x$$

Constraint (1) ensures we assign at most one section to any room $r \in \mathcal{X}_r$ for each meeting day and time $(d,t) \in \mathcal{ST}$; constraint (2) ensures all non-remote sections are assigned to at most one room. Note that in constraint (2), the inequality provides the flexibility to have originally non-remote sections be taught remotely.

In the subsections that follow, we discuss different objectives that can be used to optimize the decisions in the CCMAP. We discuss the motivation for considering each objective, and we formulate a unique IP for each objective. When necessary, we introduce new sets and parameters.

3.2. Delivery Mode Maximization

The first objective we discuss is delivery mode preference maximization. A key administrative concern is to ensure that as many sections as possible are taught in their preferred mode. The mode preferences discussed are influenced primarily by the instructor health accommodations and pedagogical concerns mentioned in Section 1.

In order to properly model the desired objective, we introduce the categorical parameter, $l_{x,r}$, which will denote the delivery mode for each section x, if it is assigned to room r following the logic described in Section 3.1:

$$l_{x,r} = \begin{cases} \text{Residential Spread} & \text{if } p_x \leq n_r, \\ \text{Hybrid Split} & \text{if } (k-1)n_r < p_x \leq kn_r, k \in \{2,\ldots,m_x\} \\ \text{Hybrid Touch Point} & \text{if } m_x n_r < p_x \leq W m_x n_r / S. \end{cases}$$

Note that $l_{x,r}$ is not defined if $p_x > W m_x n_r / S$ because this implies room r is not large enough to ensure that each student in section x could have S touch points over the course of the semester. In this case assigning x to r assumes that the section will be run in Remote mode and $r \notin \mathcal{R}_x$. Note that we do not impose the requirement that all sections of the same course be assigned the same mode. This was intentional because for some courses, the academic units aimed to offer several mode options to students.

The choice of room assigned (or not assigned) dictates the highest-contact mode that is *possible* for a section to be delivered in. In practice at GT, the unit responsible for each course provides a preference for one of Residential Spread, Hybrid, or Remote for the course section, as discussed in Section 1. We define sets to capture the unit preferences below. Currently, we assume that if a unit specifies Hybrid, then they will be satisfied with either Hybrid Split or Hybrid Touch Point mode. However, units may express a more specific hybrid preference. For example, they may provide a lower bound on the number of meetings each student could have with the instructor during the semester for a given course (the parameter S would then become S_x, and be section-dependent).

- L_x: the set of delivery modes section x may be taught in, as determined by its unit or by university administration. We call this the set of modes *preferred* for the section. L_x must be a subset of the partially ordered set $L = \{\text{Residential Spread}, \text{Hybrid Split}, \text{Hybrid Touch Point}, \text{Remote}\}$. The partial ordering of L follows Residential Spread $\geq$ Hybrid Split $\geq$ Hybrid Touch Point. L_x obeys upward closure on L. For example, if Hybrid Split $\in L_x$, then Residential Spread $\in L_x$. Intuitively, we are assuming that there would be no objection to teaching a section in a mode with increased contact hours. This does not hold for sections with preference for remote instruction, as there may be general objections to teaching such sections in person.
- $\hat{\mathcal{R}}_x = \{r \in \mathcal{R}_x \,:\, l_{x,r} \in L_x\}$: the set of rooms that match the room type required for section x, (classroom, chemistry lab, etc.) and in which the class may be delivered in a preferred mode.

We now present a formulation that maximizes the number of sections delivered in a preferred mode. For a given section, $x \in \mathcal{X}$, if x is assigned to a room enabling it be delivered in preferred mode, then the expression $\sum_{r\in\hat{\mathcal{R}}_x} X_{x,r} = 1$, while if x is not assigned to a room and Remote $\in L_x$ then x may also be delivered in preferred mode. Thus, the number of sections that are delivered in a preferred mode is

$$O_p(X) = \sum_{x\in\mathcal{X}} \sum_{r\in\hat{\mathcal{R}}_x} X_{x,r} + \sum_{x\in\mathcal{X}:\text{Remote}\in L_x} (1 - \sum_{r\in\mathcal{R}_x} X_{x,r})$$

and the corresponding maximization problem is

$$\underset{X \in P}{\text{maximize}} \quad O_p(X). \tag{3}$$

3.3. Contact Hours Maximization

Another objective of interest to campus planners relates to the total number of student-credit hours delivered in-person (either via a Residential Spread course or via the in-person component of a Hybrid course). We refer to these in-person student-credit hours as 'student contact hours'. We first provide some notation for contact hours and then describe the objective in detail.

For a given section $x \in \mathcal{X}$, let h_x denote the number of contact hours per meeting of the section. So the total number of contact hours delivered face-to-face under normal circumstances over the semester is $h_x m_x W$ and the total number of student-contact hours is $h_x m_x W p_x$. If x is assigned to a given classroom $r \in \mathcal{R}$, then it must be that $p_x \leq W m_x n_r / S$ and the number of student-contact hours delivered face-to-face is dictated by the mode choice. We calculate a new parameter, dependent on the section, room and preferred mode set, to denote the number of contact hours per student and per week, given by

$$f_{x,r} = \begin{cases} h_x m_x, & \text{if } p_x \leq n_r \text{ and Residential Spread} \in L_x, \\ h_x(m_x - 1), & \text{if } p_x \leq n_r \text{ and Residential Spread} \notin L_x, \\ h_x(m_x - k + 1) & \text{if } (k-1)n_r < p_x \leq k n_r, k \in \{2, \dots, m_x\} \\ \lfloor W m_x n_r / p_x \rfloor h_x / W, & \text{otherwise, so } p_x \leq W m_x n_r / S. \end{cases}$$

The second alternative makes the assumption that if a section for which a Hybrid mode is preferred is assigned to a room big enough for Residential Spread, then it will be taught in the highest contact hours possible for Hybrid Split, meaning that each student skips one class per week. The last alternative is deduced by observing that if all students meet the instructor in-person the same number of times during the semester, for the duration of a section meeting, then each student can experience at most $\lfloor W m_x n_r / p_x \rfloor h_x$ contact hours in total over the semester.

Let $O_h(X)$ be the total number of contact hours delivered by an assignment X, where

$$O_h(X) = \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}_x} f_{x,r} p_x X_{x,r}.$$

Then the problem of finding an assignment which maximizes the number of contact hours delivered is

$$\underset{X \in P}{\text{maximize}} \quad O_h(X). \tag{4}$$

One note is that the contact hours objective coefficient factor $f_{x,r}$ defined above assumes that if a section, x, meets m_x times per week, then in Hybrid Split mode students could attend the class only once per week, twice, three times, etc., up to $m_x - 1$ times, i.e., each student skips the same number of classes per week, which is at least 1 and at most $m_x - 1$.

Alternatively, Hybrid Split mode may be implemented using distinct cohorts of students so that students in each cohort never meet students from the other cohorts when attending their class for the section. Let κ_x denote the number of times each student attends class per week, for section x. In the case that $m_x = 2$, Hybrid Split implies that $\kappa_x = 1$ and the class can be split into two distinct cohorts that meet on alternate class days. In the case that $m_x = 3$, Hybrid Split with $\kappa_x = 1$ implies that the class can be split into three distinct cohorts that meet on one of each of the class days. The case that $m_x = 3$, Hybrid Split with $\kappa_x = 2$ prohibits distinct cohorts, e.g. a MWF section has three cohorts that attend MW, WF and MW, but each of these interact with the others in one class meeting per week. Similarly, $m_x = 4$ run in Hybrid Split with $\kappa_x = 1$ or $\kappa_x = 2$ imply 4 or 2 distinct cohorts, while $\kappa_x = 3$ prohibits distinct cohorts. Since the ability to implement distinct cohorts could be significant in reducing disease spread, campus planners have an interest in only permitting these modes. This leads to the following, alternative, definition of the contact hours objective coefficients:

$$f_{x,r} = \begin{cases} h_x m_x, & \text{if } p_x \le n_r \text{ and Residential Spread} \in L_x, \\ h_x & \text{if } m_x \in \{2,3\}, p_x \le m_x n_r, \text{ and } (p_x > n_r \text{ or Residential Spread} \notin L_x), \\ h_x & \text{if } m_x = 4 \text{ and } 2n_r < p_x \le 4n_r, \\ 2h_x & \text{if } m_x = 4, p_x \le 2n_r \text{ and } (p_x > n_r \text{ or Residential Spread} \notin L_x), \\ \lfloor W m_x n_r / p_x \rfloor h_x / W, & \text{otherwise, so } p_x \le W m_x n_r / S. \end{cases}$$

3.4. Plan Stability Objectives

Campus planners, instructors, and students have preconceived expectations of the room assignments: prior to the need to accommodate physical distancing, rooms were already assigned to courses. Given a pre-existing assignment, we consider two types of plan stability objectives: a relocation distance defined as the sum of the travel distances between the buildings in which sections are located in the pre-existing assignment versus a new assignment, and the number of assignment changes.

Pre-existing room assignments are usually in buildings that are preferred by the unit responsible for the course, and assigning a course to a building far away from the pre-existing assignment is likely to create long travel times for students and faculty moving between classes. The second stability objective relates to the amount of labor required in the Registrar's Office to modify room assignments in the system. It also measures the degree of interaction needed with instructors to check that undocumented features of their new rooms meet their requirements for the course. In general, the allocation of sections to classrooms is not solved once, but is iterated over multiple phases as new enrollment information is received, as is the case with GT's two-phase registration system, for example. In the earlier phases, planners might be more willing to make bigger changes to a previous assignment, where labor is not as much of an issue. As we get closer to the start of the semester and enrollments become less volatile, planners might be less flexible, equally penalizing any change in the current assignment. We formally define both objectives.

3.4.1. Relocation Distance Minimization The preference for pre-existing room assignments and short travel distances can be captured in a single objective using the following parameters.

- $\theta_x \in \mathcal{R}$: the pre-existing room assigned to section x, which may be null, for example, if the section is newly created.
- $\mathcal{B}$: the set of buildings in which there are classrooms.
- $d_{b,b'}$: the travel distance from building b to building b', for $b, b' \in \mathcal{B}$. In the case that $b = b'$, this distance may still be positive, to model the likely (average) distance between different rooms within the same building.
- $\mathcal{R}_b \subseteq \mathcal{R}$: the set of rooms in building $b \in \mathcal{B}$.
- $b_r \in \mathcal{B}$: the building in which room $r \in \mathcal{R}$ is located.

We model the cost of assigning section x to room r versus its pre-existing assignment as:

$$\Delta_{x,r} = \begin{cases} 0, & \text{if } r = \theta_x \\ d^2_{b_{\theta_x}, b_r}, & \text{otherwise.} \end{cases}$$

The sum of squared distances between a given assignment X and a pre-existing one is:

$$O_d(X) = \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}_x} \Delta_{x,r} X_{x,r},$$

and minimizing the sum of squared relocation distances is then

$$\underset{X \in P}{\text{minimize}} \quad O_d(X). \tag{5}$$

Note that the value assigned to $d_{b,b}$ is unlikely to greatly affect the result, provided it satisfies

$$d_{b,b} < \delta \min_{b' \in \mathcal{B}, b' \neq b} d_{b,b'}$$

for $0 < \delta << 1$. For example, a δ value of 0.25, 0.1, or 0.01 may yield reasonable solutions. Note also that the travel distance $d_{b,b'}$ between two buildings b and b' could be any measure of distance that is appropriate for the campus, such as walking time or Manhattan distance. For our study, we approximate the travel distance parameters, $d_{b,b'}$, as the haversine distance between pairs of longitude and latitude coordinates of the buildings.

3.4.2. Assignment Change Minimization In this stability objective, we are interested in the number of changes to a pre-existing assignment. Note that this quantity is related to the relocation distance, as minimizing the latter indirectly lowers the former if $d_{b,b} > 0$ for all buildings $b \in \mathcal{B}$.

To minimize the number of changes to a pre-existing assignment, we equivalently maximize $O_a(X)$, given by

$$O_a(X) = \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}_x} X_{x\theta_x}.$$

and the associated maximization problem is

$$\underset{X \in P}{\text{maximize}} \quad O_a(X). \tag{6}$$

4. Multi-Objective Optimization

Given the many objectives of interest, campus planners will want to ideally take multiple, if not all, objectives into account when deciding on an assignment. We thus turn to a multi-objective optimization framework to handle the different objectives. While the literature is rich with multi-objective optimization algorithms (see Przybylski and Gandibleux (2017), Pardalos et al. (2017), Antunes et al. (2016)), we provide a classical hierarchical method for its simplicity and interpretability.

4.1. Hierarchical Optimization

Let $z_1(X), \ldots, z_K(X)$ be the objectives of interest from $O_p(X), O_h(X), -O_d(X)$ and $O_a(X)$. The problem of simultaneously optimizing these objectives can be written as

$$\underset{X \in P}{\text{maximize}} \quad \{z_1(X), z_2(X), \ldots, z_K(X)\}.$$

Note that we use the negative of the relocation distance objective as it is originally a minimization problem (see (5)). In a hierarchical optimization approach, we assume we are given a ranking of the objectives by the campus planners. We then optimize each objective sequentially in the order of the ranking while constraining the previous objectives in the ranking to be within a specified tolerance of their optimal value resulting from the previous solves. For example, if objectives $z_1(X), \ldots, z_K(X)$ are in order of importance, we first maximize $z_1(X)$ to obtain an optimal value z_1^*. We then maximize $z_2(X)$ with the added constraint $z_1(X) \geq (1-\alpha_1)z_1^*$, where α_1 is a tolerance parameter between 0 and 1 which we vary until we are satisfied with the value of the second objective. We thus obtain α_1^* and z_2^*. When optimizing objective ℓ in the ranking, we solve

$$\begin{aligned} \underset{X \in P}{\text{maximize}} \quad & z_\ell(X) \\ \text{subject to} \quad & z_k(X) \geq (1-\alpha_{\ell-1})z_{\ell-1}^*, \\ & z_k(X) \geq (1-\alpha_k^*)z_k^* \qquad , \; k = 1, \ldots, \ell-2 \end{aligned}$$

where $0 \leq \alpha_k \leq 1$ for $k = 1, \ldots, \ell-1$. Tolerances $\{\alpha_k^*\}_{k=1}^{\ell-2}$ are fixed from the previous solves in the hierarchy, and we only vary $\alpha_{\ell-1}$ to obtain $\alpha_{\ell-1}^*$ and z_ℓ^*. Note that one can also increase any of the "fixed" values $\{\alpha_k^*\}_{k=1}^{\ell-2}$ if desired; decreasing them might lead to an infeasible solution if other tolerances are not relaxed accordingly.

Such a hierarchical procedure is interpretable and is simple to implement. In fact, this approach can be done automatically using a commercial solver like Gurobi (Gurobi Optimization, LLC 2021).

4.2. Obtaining Non-Dominated Solutions

Hierarchical optimization is not guaranteed to produce non-dominated solutions if the relaxation parameters are strictly positive; there may be solutions optimizing the final objective that give better values for the relaxed objectives. For example, consider using the hierarchy of objectives where we first maximize the number of sections delivered in a preferred mode $O_p(X)$, then maximize the number of contact hours $O_h(X)$, and finally minimize the relocation distance objective $O_d(X)$. Denote the optimal values of the first two optimization problems in the hierarchy by z_p^* and z_h^*, respectively. The last optimization problem where we minimize $O_d(X)$ with added constraints $O_p(X) \geq (1-\alpha_p)z_p^*$ and $O_h(X) \geq (1-\alpha_h)z_h^*$ has optimal value z_d^* and optimal solution X^*. There may be another solution $\hat{X}$ where $O_d(\hat{X}) = O_d(X^*)$ and either $O_p(\hat{X}) > O_p(X^*)$, $O_h(\hat{X}) > O_h(X^*)$ or both. In other words, there may exist a solution $\hat{X}$ which dominates X^*, having equal total relocation distance, but being at least as good in all relaxed objectives and better in at least one.

Final solutions used in practice should be non-dominated. We thus first use hierarchical optimization, then optimize a weighted combination of the objectives, where each objective is constrained to be at least as good as their value in the hierarchical optimization. To obtain a non-dominated solution, we may solve an IP of the following general form:

$$\begin{aligned} \underset{X \in P}{\text{maximize}} \quad & \sum_{k=1}^{K} w_k z_k(X) \\ \text{subject to} \quad & z_k(X) \geq z_k^*,\ k = 1, \ldots, K \end{aligned} \tag{7}$$

where w_k are positive parameters and z_k^* is the value of objective k resulting from the hierarchical optimization. Parameters w_k can be chosen based on the relative scale of the objectives, for example.

5. Application to GT's Fall 2020 Schedule

In this section, we present our modeling approach applied to the Fall 2020 schedule at GT. We first describe the data used in our study, and then present how the approach and models can be used to inform trade-offs for an entire university in the baseline scenario. All models were implemented in Python 3.8.3 using the commercial solver Gurobi 9.1.0 (Gurobi Optimization, LLC 2021), and all mixed-integer programs were solved using default settings unless specified otherwise. Experiments were run on a 3.6 GHz Linux machine with 16 GB RAM, 6 cores and 2 threads per core.

5.1. Data preparation and model implementation

We began working with the registrar in June 2020. At this point in time, with the exception of incoming freshmen and first-year graduate students, the majority of GT students had already registered for Fall 2020 classes during Phase I registration in Spring 2020 reserving more than 57,000 seats.

The Registrar provided us with the following data: (1) course sections and associated time slots, faculty, and teaching mode preferences; and (2) a list of all Registrar-controlled and department-controlled classrooms, with original seating capacities. Our team had, in a separate project to support physically-distanced classroom layouts, redesigned all the Registrar-controlled classrooms. Therefore, we used the adjusted seating capacities based on those layouts and interpolated those numbers for the department-controlled classrooms, which we did not redesign precisely.

We excluded courses that are not taught on Atlanta's maincampus, courses that do not need a classroom, such as thesis and undergraduate research. We combined cross-listed sections with their originating department section, given that these sections meet at the same time and place, and with the same instructor. Table 2 provides descriptive statistics of the course sections and classrooms in the model. Notably, there are 173 sections with a projected enrollment of over 100 students and only one classroom that can accommodate over 100 students simultaneously with 6-ft physical distancing. The average projected enrollment is 44.3 students, while the average adjusted-capacity of classrooms is 14 seats. Courses with mode preferences of Hybrid could be taught in either Hybrid Split or Hybrid Touchpoint in order to have their mode preference satisfied. Moreover, these courses did not specify the minimum number of touch points, which increases the space of feasible solutions. However, many of those solutions would dictate 1 or 2 touch points in the semester, which is an inferior experience for students who register in a Hybrid course in hopes of more in-person interactions and then receive an almost entirely remote experience. Therefore, the power of our approach is its ability to identify reassignments that not only satisfy mode preferences - which would not be difficult to achieve with such a large proportion of Remote sections and such a small proportion of Residential Spread sections - but primarily in maximizing the number of in-person contact hours for the students.

To investigate the robustness of our managerial insights, we also investigated an alternative data set corresponding to a hypothetical scenario in which Spring 2020 courses were also subject to sudden scarcity (as opposed to a complete shutdown of campus). We describe the data set and the resulting analysis the findings in the E-companion. We found that our managerial insights are consistent in the other data set.

5.2. Results

We present the results of our study in this section. We first describe a baseline scenario presented to us by the Registrar, then discuss the implications of decentralized planning and requirements on number of touchpoints, and the value of hierarchical optimization and the trade-offs between the different objectives.

5.2.1. Baseline scenario: Satisfied preferences, contact hours, and plan stability Given the ranking of objectives from the GT campus planners, we solved the hierarchical optimization problem as explained in Section 4.1, where we first maximized the mode preference satisfaction, then maximized contact hours, and finally minimized the relocation distance. Specifically, we set a tolerance of 0.01 and 0.1 for the number of mode preferences and contact hours objectives, respectively. Since campus planners prioritized satisfying mode preferences

Table 2 Characteristics of course sections and classrooms in the data set

Characteristic	**Value**
Number of Course Sections	2249
Mean Current Enrollment	25.4
Median Current Enrollment	14
Mean Projected Enrollment	44.3
Median Projected Enrollment	30
Number of Unique Time Slots	204
Number of Buildings	37
Projected enrollment	**Sections**
150+ students	88
100 - 149 students	85
50 - 99 students	550
25 - 49 students	669
Mode preference	**Sections**
Residential Spread mode preference	118
Hybrid (Split or Touchpoint) mode preference	1291
Remote mode preference	840
Adjusted capacity	**Rooms**
150+ seats	0
100-149 seats	1
75-99 seats	4
50-74 seats	16
25-49 seats	51
20-24 seats	61
15-19 seats	75
10-14 seats	108
5-9 seats	144
2-4 seats	117
Total with 2+ seats	577
Control	**Rooms**
Registrar-controlled	152
Department-controlled	441

the most, we selected a tight tolerance for this objective. In other words, in order to improve the number of contact hours achieved, we are only willing to decrease the number of mode preferences satisfied by 1%, and in order to improve the relocation distance, we are willing to decrease the contact hours, by 10%.

We also note that if a section has a mode preference for Remote but is taught with an in-person component, this would improve the contact hours objective (4) and potentially provide improved quality of education, at the expense of the mode preferences objective (3). However at GT, Remote mode preferences were often dictated by instructors' health accommodations. Therefore, courses with Remote mode preference were automatically assigned to Remote mode.

Table 3 demonstrates properties of the solution determined by our hierarchical optimization model and compares the optimization model solution to a scenario in which no rooms are reassigned (no room re-assignment, denoted by NRR). Under NRR, each section would keep its originally assigned room, unless that section's preferred mode is Remote, in which case it is removed from the classroom. In some cases, this led to a section with Residential Spread mode being assigned a hybrid mode because the classroom cannot fit the enrollment while maintaining physical distancing. In other cases, the section indeed was forced to Remote mode because the classroom was too small to hold a Hybrid format of this class. The registrar had initially considered adhering to NRR, in order to minimize the administrative burden of re-assigning rooms.

Our model is able to achieve a total of 86,529 in-person contact hours across all students. If there was no need for physical distancing and each section with a non-Remote mode preference could be taught in Residential Spread, there would be 105,030 total contact hours. Thus, our solution satisfied 82% of the maximum possible contact hours compared to 71% from the NRR strategy. We note our solution kept 400 sections, or 17.3% of sections, in their originally assigned room. The solution also kept 491 sections, or 38.6% of all sections, in different rooms within the same building.

Our model was able to satisfy 99.2% of mode preferences compared to 95.7% in the NRR strategy. NRR is likely able to achieve a high satisfaction of mode preferences because nearly 40% of sections were preferred to be in Remote mode. However, among sections with a Residential Spread mode preference, NRR only achieves 47% mode satisfaction, which was one of the Registrar's key concerns with the NRR policy. Meanwhile, our optimization model achieves 87% mode satisfaction among these sections, indicating its ability to significantly improve accommodation for sections with preference for Residential Spread. In the E-Companion, we consider a scenario in which all sections are preferred to be delivered in Residential Spread. The benefit of the optimization model is clearly seen in this scenario since NRR is only able to satisfy 37% of mode preferences compared to 90% using the optimization model (see E-Companion §EC.1.3.2).

5.2.2. Implications of decentralized planning In the previous analysis, we assumed the Registrar is able to assign any section to any classroom on campus. In reality, each classroom at GT is classified as either a Registrar-controlled or department-controlled classroom. Department-controlled classrooms are managed by individual departments, rather than campus planners, and the individual departments choose which sections are assigned to decentralized classrooms. Meanwhile, Registrar-controlled classrooms are managed by campus

Table 3 Comparison of 2 strategies for course mode selection and room re-assignment: No room re-assignment (NRR) in which each section keeps its original pre-pandemic room assignment & our model which simultaneously assigns modes and re-assigns classrooms

Metric	Strategy	
	NRR	Our model
Courses delivered in Remote mode	155	843
Courses delivered in Hybrid mode	1771	377
Courses delivered in Residential Spread mode	323	1019
Contact hours, % of maximum	71	82
Mode preferences satisfied, %	95.7	99.2
Residential Spread mode preferences satisfied, %,	47	87
Sections with initial room assignment, %	100	36
Sections with initial building assignment, %	100	14

planners, rather than individual departments. Any section may be taught in a Registrar-controlled classroom. We considered the impact of *centralized planning*, in which the Registrar is able to assign sections to both Registrar-controlled classrooms and department-controlled classrooms, in comparison to *decentralized planning* in which the Registrar is only able to assign sections to Registrar-controlled classrooms.

To replicate the effect of decentralized planning in our optimization models, we assume that if a section was previously assigned to a classroom in a department-controlled building, that section must stay in its preassigned classroom. Other sections may be assigned to any room. This implies that if during a particular day and time, a department-controlled classroom does not have a preassigned section, this classroom is available to any other section. This assumption will give us an upper bound on the value of centralization because we are assuming that, in the decentralized planning scenario, departments are not optimally allocating their classroom space in line with the hierarchy of preferences over relocation, modes, etc. Our modeling framework is flexible enough to relax this assumption by modifying the set of allowable classrooms for each section to be restricted to the set of classrooms that its corresponding department has control over. An even stronger interpretation of decentralized planning might prohibit any section from being scheduled in a department-controlled classroom, even when it is empty. We found such an assumption to be overly restrictive and unrealistic.

We run the model under two different settings: decentralized planning and centralized planning, as described above. We expected the effect of decentralized planning to depend on the level of physical distancing enforced.

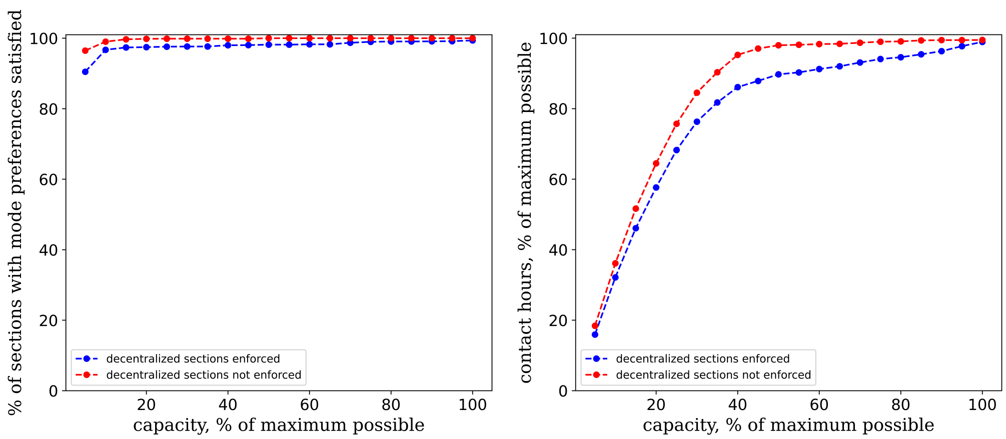

(a) Mode Preferences Satisfied vs. Classroom Capacity

(b) Contact Hours vs. Classroom Capacity

Figure 2 **Each point in (a) represents the optimal value of the optimization model (3), with the singular objective function of mode preference maximization. The effective capacity of the rooms is changed when generating various points. Similarly, each point in (b) represents the optimal value of the optimization model** (4)**, with the singular objective function of contact hour maximization. . In both plots, the red dots represent centralized planning, whereas the blue dots represent decentralized planning.**

We individually optimize over contact hours and number of mode preferences satisfied. For simplicity, we only consider a single objective at a time. We also optimize over a variety of effective capacity levels. Specifically in each model run, we assume each classroom has an effective capacity that is equal to $v\%$ of the total room capacity, where $v \in \{5, 10, 15, ..., 100\}$, $v = 100\%$ implies full capacity, i.e., no physical distancing.

Figures 2a and 2b show that compared to decentralized planning, centralized planning can result in a noticeable increase in contact hours and modest increase in mode preference satisfaction. Specifically, at 25% capacity, we see a 10.9% relative increase in contact hours and 2.3% relative increase in mode preference satisfaction. The small increase in mode preferences satisfaction is due to the large proportion of sections with Remote or Hybrid mode preference in this case study, 37% and 57%, respectively. That, combined with the fact that department-controlled classrooms are generally small to medium in size, diminished the benefit of adding more classrooms to the available central inventory.

5.2.3. Sensitivity analysis on the number of touchpoints We now conduct a sensitivity analysis over the parameter S, the minimum number of times a student can attend a class over the semester for the class to not be deemed fully remote. Lower values of S may result

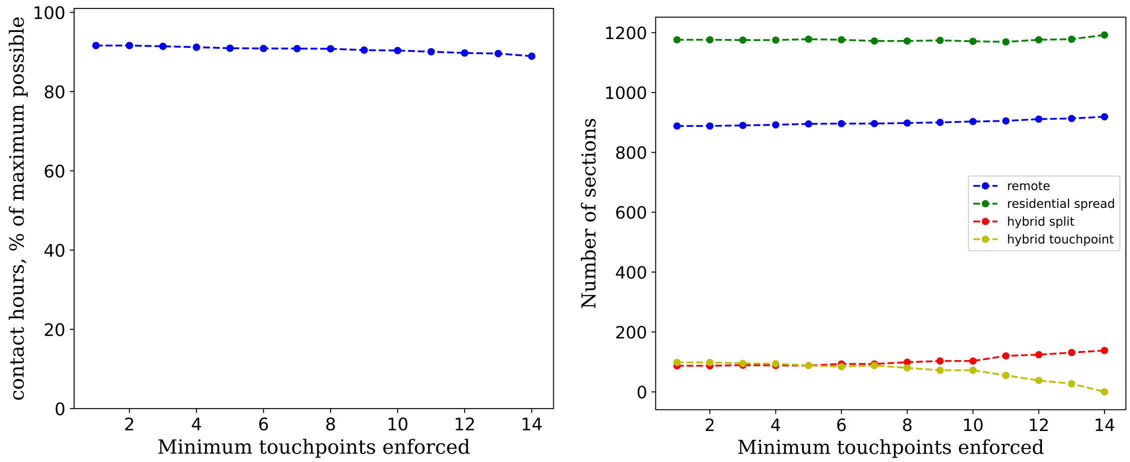

(a) Contact Hours vs Minimum Number of Touchpoints

(b) Mode Preferences Satisfied vs Minimum Number of Touchpoints

Figure 3 **Each point in plot (a) corresponds to the optimal value of the of the optimization model** (4)**, with the singular objective function of contact hour maximization. A different value of** S **is used to generate each point, where** $S \in \{1, 2, \ldots, 14\}$**. Plot (b) corresponds to the output of the same optimization model** (4)**, and the y-axis represents the number of sections taught in each delivery mode.**

in unconventional scenarios in which students rarely attend in-person. Consider the case of $S = 1$ where students may attend a class in person only once throughout the semester. In general, higher values of S may result in higher educational value by providing more consistent in-person meetings.

At the same time, a lower value of S relaxes our optimization problems, and can therefore only improve the value of our objectives. We analyze the relationship between the value of S and the number of contact hours, i.e. $O_h(X)$, by solving (4) for varying values of of S. Recall that there are 14 weeks of instruction in a semester at GT. We test values of $S \in \{1, 2, \ldots, 14\}$. Setting $S = 14$ effectively means that the student can attend class in person at least once per week, meaning no sections can be taught in Hybrid Touch Point (i.e., they would be in Hybrid Split mode).

Figure 3a shows that the minimum number of touch points has limited effect on maximum contact hours achieved by optimization model (4). Meanwhile, Figure 3b illustrates that the number of sections taught in Hybrid Touch Point decreases as the value of S increases, with no sections taught in Hybrid Touch Point at $S = 14$. In other words, as S increases, the optimization model is able to move the sections that would be taught in Hybrid Touch Point into some combination of Residential Spread, Hybrid Split, and Remote delivery modes, without significantly reducing the total number of contact hours. Thus, if campus planners are focused on maximizing the contact hours, they may consider removing the option of teaching sections in Hybrid Touch Point, as the logistical inconvenience it causes may not be worth the marginal increase in contact hours it provides.

5.2.4. Value of hierarchical optimization In this section, we perform sensitivity analysis on the tolerance values to better understand the trade-offs between the different objectives. For simplicity, the minimum number of touch points is set to one, and we assume all classrooms are centralized, i.e. all classrooms are managed by the Registrar. We change the relative MIP optimality gap from 10^{-4} to 10^{-3} for the experiments in this section.

To analyze the trade-offs between the objectives, we approximate the Pareto frontier by solving

$$
\begin{aligned}
\underset{X}{\text{minimize}} \quad & O_d(X) && \text{(8a)}\\
\text{subject to} \quad & O_p(X) \geq (1-\beta_p) z_p^*, && \text{(8b)}\\
& O_h(X) \geq (1-\beta_h) z_h^*, && \text{(8c)}\\
& X \in P && \text{(8d)}
\end{aligned}
$$

where z_p^*, and z_h^* are the optimal objective values of (3), and (4), respectively; β_p and β_h are the tolerances associated with the number of sections whose mode preferences are satisfied and contact hours, respectively. For the purpose of this analysis, we choose the relocation distance objective O_d as the stability objective. We solve (8) for a grid of tolerance values for β_p and β_h. There are more sophisticated ways to construct the Pareto frontier for multi-objective integer programs in the literature (see e.g. Zhang and Reimann (2014), Dächert and Klamroth (2015)), but these approaches are out of the scope of this paper, and we consider this approach as a simple way to analyze how the different objectives interact with each other.

We present the full results of the above analysis in the E-companion. To summarize, we observe a significant trade-off between the relocation distance objective and contact hours, where the former increases exponentially with the number of contact hours. Moreover, we note that relaxing the number of mode preferences satisfied has less of an effect on the relocation distance objective, unless we also relax the number of contact hours. It seems that to decrease the relocation distance objective, decreasing contact hours is more effective than decreasing the number of mode preferences satisfied.

6. Contribution to Decision-Making at GT

In this section, we describe the influence of our optimization modeling on decision-making at GT. We emphasize that our optimization model was used as a decision support tool in an *iterative* and *collaborative* decision-making process with the GT Task Force (as opposed to providing GT with '*the* optimal solution' for a particular hierarchy of objectives). For context, the optimization modeling process began in late spring 2020 and continued throughout the summer of 2020. During this time, the COVID-19 knowledge landscape was ever-evolving with vast amounts of uncertainty around whether or not it would be feasible to re-open the university, how the virus spreads, the guidance on physical distancing from the CDC, and other constraints from GT's governing university system. Given this uncertainty, our optimization model was critical to helping the GT COVID-19 Task Force with "what-if" scenarios to help speed decision-making during the summer by adding these considerations into the model and help GT adjust plans as new information became available. In particular, our analysis was useful in helping the COVID-19 Task Force better understand the trade-offs and how placing priority on one aspect of campus planning (e.g., administrative burden and not wanting to re-locate classes to new rooms) had impacts on the other aspects of the system (e.g., how many contact hours could be delivered.) Next, we elaborate on a few specific examples of this process.

First - and early in the decision-making timeline - the model was adapted to account for the administrative burden associated with moving classes. The goal was to demonstrate the benefits of re-assigning classrooms and achieve higher in-person contact hours for students without violating adjusted classroom capacities for physical distancing. The Task Force was assessing the degree to which the reduced classroom capacity could accommodate existing enrollments if all classes remained in their previously assigned classrooms. Our team demonstrated the benefit of using optimization to simultaneously re-assign classrooms and the mode of delivery. However, throughout this discussion, members made it clear that the time required to input these changes into the legacy system would be prohibitive and that there were administrative preferences for keeping classes in similar locations as in the past. With this in mind, we incorporated relocation distance into the model to capture these preferences and helped the Registrar consider the trade-offs between relocation and mode satisfaction.

Second, the model was used to consider the impacts of various physical distancing guidelines. During the summer of 2020, there was limited understanding of how physical distancing in classrooms would impact capacity and in turn, the feasibility of delivering in-person classes. Moreover, different universities were considering different physical distancing requirements (e.g., 3-feet vs. 6-feet) and administrators were interested in the impacts if these requirements were to change before the fall semester. Our team used the model to show the impact of different physical distance thresholds on the amount of in-person class that could be delivered, which helped the Registrar evaluate the feasibility of different mode preferences.

Third, the model was used to get a sense of the benefit of moving large classes online. At the beginning of summer 2020, different schools were proposing to move larger classes (e.g., those with 50 or more students) remotely and keeping other classes delivered in-person. Our model was used to evaluate different thresholds for what constituted a 'large' class and the corresponding amount of in-person sections that could be delivered for various thresholds. In the end, GT decided to not use a set threshold for their mode decisions. Indeed, due to our analysis, the Registrar sought to focus on the delivery of Residential Spread classes first, and then determine which course sections would be placed in Remote format.

In summary, our optimization model was used to inform the decision-making process at GT in a collaborative and iterative way as new information and considerations came to light throughout the summer of 2020.

7. Conclusions

This paper presents a hierarchical optimization framework for course mode selection and room reassignment to mitigate the impacts of sudden space reductions, such as those that occur during an emerging epidemic. Our modeling framework is flexible in that the model will always return feasible solutions which was important during the iterative process of identifying solutions that met multiple criteria.

Application of our model to data from GT led to insights for campus administrators during the COVID-19 pandemic. First, we found that optimization after sudden capacity reductions can substantially increase the amount of in-person instruction that can be delivered. Second, we show that when room capacities remain above 50% of their original capacity, our model is able to meet all mode preferences and keep in-person instruction to within 50% of pre-pandemic levels. However, when capacity drops below 50%, in-person contact hours starts to drop rapidly with every percent loss in seating capacity. This may motivate administrators to take other measures to reduce the physical distancing requirements (e.g., mask-wearing). Third, we show the value of centralized planning. Although normal operations typically operate in a decentralized fashion, centralized planning could be especially useful during times of sudden capacity reductions. We suspect that centralized planning would be even more beneficial if administrators desired more Residential Spread classes.

Although our study was motivated by sudden capacity reductions, our modeling approach could also be applied to sudden changes in course enrollments. For instance, the proposed approach can be used to optimize classroom reassignments during the GT active phase 2 registration period, even in the absence of a pandemic causing sudden capacity losses. During phase 2, students add and drop courses, and departments change enrollment limits or add course sections to the timetable. This dynamic environment results in a daily need to reassign classes to accommodate changes in demand and supply. The plan stability objective is well-suited to handle this problem efficiently and with minimal disruption to the existing assignments.

There are several limitations to our work which motivate future research. First, we do not consider revisions to the time table of courses in combination with their room assignments and mode. This modeling choice was decided after discussions with the Registrar's office which stated that the administrative burden associated with re-coding the time table would be prohibitive given the time frame to prepare for Fall 2020. In addition, creating a unified framework for simultaneously designing the time table with course mode and room assignments for the entire university would introduce new computational challenges. Second, while our model had accurate estimates of reduced capacities due to physical distancing from the GT Facilities Management team, our model relied on projected enrollment which does not consider the effects of student attendance and enrollment behavior. Course modes have to be announced before students could enroll in their courses, allowing students to switch course sections based on which course mode they prefer. Future work could incorporate these behavioral aspects to better assign course modes and allocate classrooms. Third, our models are formulated at the class-level and do not consider the impact on individual students. For instance, our model maximizes contact hours but in-person contact hours are not necessarily evenly distributed across all students; future work may consider the impact of these decisions at the student-level. The course mode selections will influence how many students have at least some in-person class and are expected to return to campus. These students will then in-turn interact with other students and people in the broader community outside of classes which could pose health risks. Therefore, these scheduling decisions should be made in conjunction with public health experts.

In summary, this paper presents a hierarchical optimization approach to help campus planners simultaneously specify course modes and re-assign classrooms during times of sudden capacity reduction and could be applied to sudden enrollment changes as well. While we demonstrate the benefits of this approach using a case study for GT during the COVID-19 pandemic, the methodology could also apply to other institutions of primary and secondary education and other emerging epidemics more broadly.

Acknowledgments

This work was supported by a grant from the Thos and Clair Muller Research Endowment Fund of the Georgia Institute of Technology's H. Milton Stewart School of Industrial and Systems Engineering, and a grant from the Georgia Institute of Technology Executive Vice President for Research COVID-19 Rapid Response Seed Grant Program. The authors would like to thank Sabrina Miller, Jimmie Hardin, Reta Pikowsky, Paul Kohn, Michelle Reinhart, and Nisha Botchwey for providing data and feedback throughout the project.

References

Antunes CH, Alves MJ, Clímaco J (2016) *Multiobjective linear and integer programming* (Springer).

Barnhart C, Bertsimas D, Delarue A, Yan J (2021) A unified model for course scheduling under sudden scarcity: Applications to pandemic planning. *Manufacturing Service Operations Management* .

Carter MW, Tovey CA (1992) When is the classroom assignment problem hard? *Operations Research* 40(1-supplement-1):S28–S39.

Chen G, Fei X, Jia H, Yu X, Shen S (2020) An optimization-and-simulation framework for redesigning university campus bus system with social distancing. *arXiv preprint arXiv:2010.10630* .

Dächert K, Klamroth K (2015) A linear bound on the number of scalarizations needed to solve discrete tricriteria optimization problems. *Journal of Global Optimization* 61(4):643–676.

de Oliveira PM, Mesquita LC, Gkantonas S, Giusti A, Mastorakos E (2021) Evolution of spray and aerosol from respiratory releases: theoretical estimates for insight on viral transmission. *Proceedings of the Royal Society A* 477(2245):20200584.

Dong E, Du H, Gardner L (2020) An interactive web-based dashboard to track COVID-19 in real time. *The Lancet infectious diseases* 20(5):533–534.

Doug L (2020) Low-income students top presidents' COVID-19 worry list. *Inside Higher Ed* .

Freeman S, Eddy SL, McDonough M, Smith MK, Okoroafor N, Jordt H, Wenderoth MP (2014) Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences* 111(23):8410–8415.

Glassey CR, Mizrach M (1986) A decision support system for assigning classes to rooms. *Interfaces* 16(5):92–100.

Gorman MF (2021) Contextual complications in analytical modeling: When the problem is not the problem. *INFORMS Journal on Applied Analytics* .

Gosselin K, Truchon M (1986) Allocation of classrooms by linear programming. *Journal of the Operational Research Society* 37(6):561–569.

Gressman PT, Peck JR (2020) Simulating COVID-19 in a university environment. *Mathematical Biosciences* 328:108436, ISSN 18793134, URL `http://dx.doi.org/10.1016/j.mbs.2020.108436`.

Gurobi Optimization, LLC (2021) Gurobi Optimizer Reference Manual. URL `https://www.gurobi.com`.

Kostuch P (2004) The university course timetabling problem with a three-phase approach. *International Conference on the Practice and Theory of Automated Timetabling*, 109–125 (Springer).

Lach G, Lübbecke ME (2008) Optimal university course timetables and the partial transversal polytope. *International Workshop on Experimental and Efficient Algorithms*, 235–248 (Springer).

Lederman D (2020) Trying to make sense of a fluid fall. *Inside Higher Ed* .

Lopman B, Liu CY, Le Guillou A, Handel A, Lash TL, Isakov AP, Jenness SM (2021) A modeling study to inform screening and testing interventions for the control of SARS-CoV-2 on university campuses. *Scientific Reports* 11(1):1–11.

Mulvey JM (1982) A classroom/time assignment model. *European Journal of Operational Research* 9(1):64–70.

National Center for Education Statistics (2016) The NCES fast facts tool provides quick answers to many education questions (national center for education statistics) .

Pardalos PM, Žilinskas A, Žilinskas J (2017) *Non-convex multi-objective optimization* (Springer).

Phillips AE, Waterer H, Ehrgott M, Ryan DM (2015) Integer programming methods for large-scale practical classroom assignment problems. *Computers & Operations Research* 53:42–53.

Przybylski A, Gandibleux X (2017) Multi-objective branch and bound. *European Journal of Operational Research* 260(3):856–872.

Rogers V (2020) Campus leaders share plan to return to campus. *Georgia Tech News Center* URL `https://news.gatech.edu/2020/06/17/campus-leaders-share-plan-return-campus`.

Savage MP, Fischman DL, Mamas MA (2020) Social intervention by the numbers: Evidence behind the specific public health guidelines in the covid-19 pandemic. *Population Health Management* .

Steimle L, Sun Y, Johnson L, Besedeš T, Mokhtarian P, Nazzal D (2022) Students' preferences for returning to colleges and universities during the COVID-19 pandemic: A discrete choice experiment. *Socio-economic Planning Sciences* .

Ueda H, Ouchi D, Takahashi K, Miyahara T (2000) A co-evolving timeslot/room assignment genetic algorithm technique for university timetabling. *International Conference on the Practice and Theory of Automated Timetabling*, 48–63 (Springer).

Waterer H (1995) A zero-one integer programming model for room assignment at the University of Auckland. *Proceedings of the 1995 ORSNZ Conference*.

Weeden KA, Cornwell B (2020) The small-world network of college classes: Implications for epidemic spread on a university campus. *Sociological science* 7:222–241.

Zhang W, Reimann M (2014) A simple augmented -constraint method for multi-objective mathematical integer programming problems. *European Journal of Operational Research* 234(1):15–24.

E-Companion
Multi-criteria Course Mode Selection and Classroom Assignment Under Sudden Space Scarcity

EC.1. Analysis of an Alternative Data Set

To test the robustness of our managerial insights, we considered another data set besides the Fall 2020 semester, which motivated this work. This section describes the alternative data set and the corresponding results.

EC.1.1. Description of Alternative Data Set

The alternative data set considers a hypothetical scenario in which the Registrar attempted to accommodate physically-distanced learning during the Spring 2020 semester, as opposed to delivering all classes virtually. Because this is a hypothetical scenario, we used the following data sources to construct our data set for this Hypothetical Spring 2020 scenario:

- Timetable, course sections, and original room assignments: These inputs were extracted from the Spring 2020 schedule before all courses were moved to online format.
- Mode preferences: We consider two scenarios:
  - — Scenario 1: Mode preferences were based on the observed modes from the Spring 2021 semester. We note that the observed modes are a function of the mode preference and the classroom capacities available on-campus.
  - — Scenario 2: Mode preference were set to Residential Spread for all sections.

In this section, we present our modeling approach applied to the Spring 2020 schedule at GT. We first describe the data used in our study, and then present how the approach and models can be used to inform trade-offs for an entire university in the baseline scenario. All models were implemented in Python 3.8.3 using the commercial solver Gurobi 9.1.0 (Gurobi Optimization, LLC 2021), and all mixed-integer programs were solved using default settings unless specified otherwise. Experiments were run on a 3.6 GHz Linux machine with 16 GB RAM, 6 cores and 2 threads per core.

EC.1.2. Data preparation and model implementation

Most of the assumptions regarding the alternative data set are consistent with the assumptions made for the original data set described in Section 5. To summarize, we use registrar provided data on course sections and associated time slots. We excluded courses that are

Table EC.1 Characteristics of course sections and classrooms in the alternative data set

Characteristic	**Value**
Number of Course Sections	2052
Mean Actual Enrollment	33.8
Median Actual Enrollment	23
Number of Unique Time Slots	193
Number of Buildings	37
Actual enrollment	**Sections**
150+ students	71
100 - 149 students	51
50 - 99 students	228
25 - 49 students	927
Mode preferences, Scenario 1	**Sections**
Residential Spread mode preference	465
Hybrid mode preference	1586
Remote mode preference	0
Mode preference, Scenario 2	**Sections**
Residential Spread mode preference	2052
Hybrid mode preference	0
Remote mode preference	0
Adjusted capacity	**Rooms**
150+ seats	0
100-149 seats	1
75-99 seats	4
50-74 seats	16
25-49 seats	51
20-24 seats	61
15-19 seats	75
10-14 seats	108
5-9 seats	144
2-4 seats	117
Total with 2+ seats	577

not taught on Atlanta's main campus and courses that do not need a classroom, such as thesis and undergraduate research. We combined cross-listed sections with their originating department section, given that these sections meet at the same time and place, and with the same instructor. Regarding the available classroom and their physically distanced capacities, we use the same data set as in Section 5. Unlike Section 5 however, we have access to actual enrollment data for all sections taught during Spring 2020. We use this enrollment data as our input. Table EC.1 provides descriptive statistics of the course sections and classrooms used in this case study.

The most important difference between Table EC.1 and Table 2 from the original case study regards the mode preferences. For Scenario 1 under the new case study, recall we use the actual modes assigned during Spring 2021 as our mode preferences input, and all sections had some in person component during Spring 2021. For Scenario 2, recall we are setting all mode preferences to Residential Spread. As a result, both scenarios had no sections with mode preference for Remote. In contrast, the original case study included 840 sections with preference for Remote.

Table EC.2 Comparison of the No Room Re-assignment (NRR) & our model's strategies for course mode selection and room re-assignment in preference Scenario 1. No room re-assignment (NRR) & our model. NRR keeps each section in its original pre-pandemic room assignment while our model which simultaneously assigns modes and re-assigns classrooms. Results are generated using the alternate data set, and mode preferences are are drawn from the Spring 2021 room assignments.

Metric	Strategy	
	NRR	Our model
Courses delivered in Remote mode	0	5
Courses delivered in Hybrid mode	1297	427
Courses delivered in Residential Spread mode	761	1608
Contact hours, % of maximum	48	70
Mode preferences satisfied, %	95.9	99.6
Sections with initial room assignment, %	100	39
Sections with initial building assignment, %	100	15

EC.1.3. Results for Alternative Data Set

To allow for a clear comparison, we solved the same sequence of objective functions as in Section 5.2.1. Specifically, we first maximized mode preference satisfaction with a tolerance of 0.01, then maximized contact hours with a tolerance of 0.1, and finally minimized the relocation distance.

EC.1.3.1. Scenario 1: Mode Preferences Based on Observed Modes First we consider the scenario where sections' mode preferences are based on the observed modes of the Spring 2021 semester. The results are listed in Table EC.2. They show over 99% all mode preferences can be satisfied with the optimization model. These results are consistent with those in Section 5.2.1.

Additionally, there is an increase in contact hour satisfaction from 48% in NRR to 70% with the optimization model. This result is similar to, and in fact somewhat stronger than, the increase in contact hour satisfaction from 71% to 82% in Section 5.2.1. Note contact hours satisfaction is lower compared to Section 5.2.1, both for the NRR policy and for our optimization model output. This is an expected result. In Section 5.2.1, there were a maximum of 105,030 possible contact hours, if all sections with preference for non-Remote instruction were taught in Residential Spread. Meanwhile there is a maximum of 150,611 contact hours for the new data set, where the increase is driven by the fact that the new data set has no sections with preference for Remote instruction. It is therefore more challenging to reach this new maximum number of contact hours, given the same set of classrooms. Regardless, the optimization model clearly improves the contact hours achieved in this more challenging scenario.

Lastly, the one can confirm that the percentage of sections that remain in their initially assigned buildings, and the percentage of sections that remain in their initially assigned classrooms, are nearly the same as in the original case study.

In summary, both the original case study and this scenario show that mode preference satisfaction and contact hours can be improved using our optimization model, while maintaining a similar level of classroom stability. Therefore the alternate data set demonstrates the generalizability of our results.

EC.1.3.2. Scenario 2: All Residential Spread Mode Preferences Now we consider the scenario where all sections' mode preferences are set to Residential Spread. This is clearly a more challenging scenario, and the results are listed in Table EC.3.

Notably, the NRR policy only achieves 37% mode satisfaction, while the optimization model achieves 90%. This result illustrates the effectiveness of the optimization model in achieving in person instructions for a large majority of sections. Similar to Scenario 1, we are able to achieve 68% contact hour satisfaction. The slight decrease compared to Scenario 1's value of 70% may be attributed to the more challenging mode preference objective preceding the contact hour objective. Meanwhile, the percentage of sections that can be taught in the same building or the same room drop noticeably in this scenario. One explanation may be that it is necessary move sections around campus in order to find a new room assignments that can accommodate for physical distancing while teaching 90% of sections in Residential Spread.

Again, this analysis demonstrates the generalizability of the optimization model in its ability to improve mode preference satisfaction and contact hours compared to the NRR policy. Moreover, this scenario clearly illustrates the ability of the optimization model to dramatically improve the model preference satisfaction.

Table EC.3 Comparison of the No Room Re-assignment (NRR) & our model's strategies for course mode selection and room re-assignment in preference Scenario 2. No room re-assignment (NRR) & our model. NRR keeps each section in its original pre-pandemic room assignment while our model which simultaneously assigns modes and re-assigns classrooms. Results are generated using the alternate data set, and mode preferences are all set to Residential Spread

Metric	Strategy	
	NRR	Our model
Courses delivered in Remote mode	0	7
Courses delivered in Hybrid mode	1297	193
Courses delivered in Residential Spread mode	761	1853
Contact hours, % of maximum	48	68
Mode preferences satisfied, %	37	90.3
Sections with initial room assignment, %	100	29
Sections with initial building assignment, %	100	9

EC.2. Additional Results from the Primary Analysis

We now present additional results from our primary analysis, focusing on the trade-offs between our objectives.

We choose 11 equidistant values between 0 and 0.2 inclusive for β_p and β_h. For each combination of tolerance values, we first solve (8), then optimize a weighted combination of the objectives to ensure we get a non-dominated solution as in (7). We pick weights of 60, 1, and 10^{-3} for $O_p(X), O_h(X)$, and $O_d(X)$, respectively. Weights were roughly picked based on relative scale of the optimal value of optimizing each objective individually. We omit combinations of tolerances which lead to an infeasible model due to the objective function constraints, since setting tolerances that are too tight for multiple objectives can lead to an infeasible model. In Figure EC.1, we plot the objective values resulting from solving (8) over the grid of tolerances, analyzing the trade-offs between the relocation distance objective and contact hours, and between the relocation distance objective and number of mode preferences satisfied.

In Figure EC.1a, we plot the relocation distance objective value on a logarithmic scale as a function of contact hours, O_h, for fixed values of the tolerance on the number of mode preferences satisfied, β_p. We also plot reference lines at $\log O_d \approx 11.9$ and $\log O_d \approx 18.6$, in which cases the average distance a section moves is about 2 meters and 86 meters, respectively. In the former, this roughly means sections stay in the same building on average. Although the average distance a section moves appears low in the latter case, we note that there is a significant jump in standard deviations between objective values $\log O_d \approx 11.9$ and $\log O_d \approx 18.6$, from about 13 meters to 280 meters. Moreover, the maximum relocated

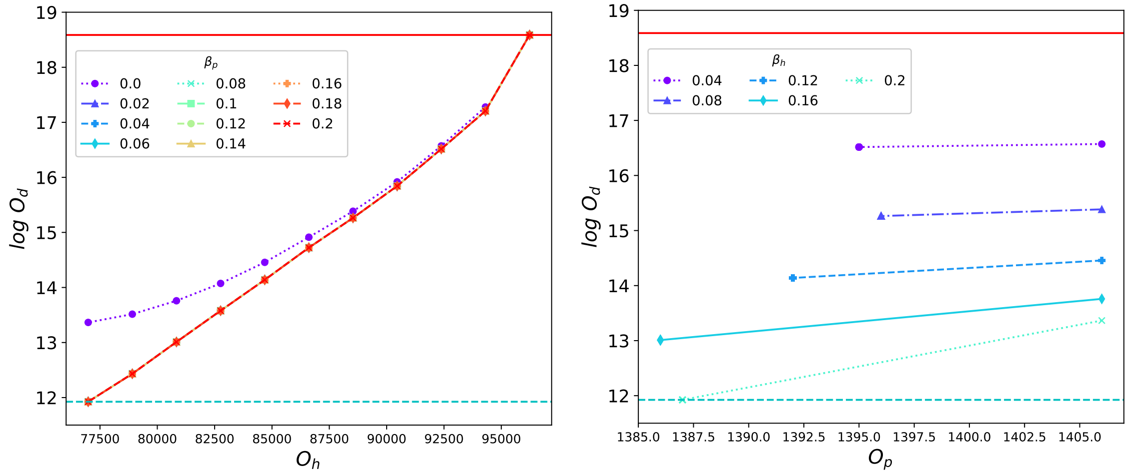

(a) Relocation Distance and Contact Hours Trade-Off

(b) Relocation Distance and Mode Preferences Trade-Off

Figure EC.1 **Objective values resulting from solving** (8) **for various tolerance values. In each figure, we plot horizontal lines at about** $\log O_d \approx 11.9$ **and** $\log O_d \approx 18.6$ **as references, where the average distance a section moves is about two meters and 86 meters, respectively. (EC.1a) Logarithm of the relocation distance objective (**O_d**) as a function of contact hours (**O_h**) for fixed values of tolerances on the number of overall mode preferences satisfied** β_p**. (EC.1b)** $\log O_d$ **as a function of the number of mode preferences satisfied (**O_p**) for fixed values of the tolerance on contact hours**β_h**.**

distance jumps from 76 meters to about 1.5 kilometers. There is a clear cost to achieving a high number of contact hours, where the relocation distance objective increases exponentially as the number of contact hours increases. In other words, to deliver a high number of contact hours, planners would have to allocate sections to distant buildings relative to the pre-existing allocation, and few sections would be allocated to their home department. Moreover, we note that relaxing the number of mode preferences O_p satisfied has little effect on the relocation distance objective, unless we also relax the contact hours objective to be less than about 87,500 hours, at which point relaxing O_p leads to a decrease in O_d. This suggests that to decrease the relocation distance, it is more effective to first relax the tolerance on contact hours, as it leads to the greatest decrease. This can also be seen in Figure EC.1b, where we plot $\log O_d$ as a function of O_p for fixed values of contact hour tolerances β_h. We observe that for low values of β_h, the relocation distance is almost constant for all values of O_p.

We show the benefits of considering a multi-objective framework by comparing the solution obtained from applying hierarchical optimization to one obtained from individually optimiz-

Table EC.4 Comparison of objective values between individually optimizing objectives and applying hierarchical optimization, where O_p is the number of sections whose mode preferences are satisfied, O_h is the number of contact hours, and O_d is the total relocation distance in meters. Hierarchical optimization runtimes are reported for solving all models in the hierarchy, excluding and including the final weighted objective, along with the resulting objective values.

Objective Optimized	O_p	O_h	O_d	Runtime (sec)
O_p	1,406	33,861	5.12×10^8	9
O_h	1,347	96,232	4.67×10^8	11
Hierarchical Optimization	1,392	86,609	2.48×10^6	97
Hierarchical + Weighted	1,395	86,610	2.48×10^6	208

ing contact hours and number of mode preferences satisfied. Note that we do not consider stability objectives individually, as without constraints on other objectives, the pre-existing assignment is optimal. We report the values of the different objectives and computation times in Table EC.4. We note that optimizing any individual objective takes less than 12 seconds. The total runtime of solving the models in the hierarchy excluding finding a non-dominated solution is about 97 seconds. After solving (7) where we optimize a weighted average of the objectives, the total runtime is about 208 seconds. As before, we use the weights of 60, 1 and 10^{-3} in the weighted objective for O_p, O_h and O_d, respectively. By considering multiple objectives, we are able to reduce the total relocation distance by two orders of magnitude, at only a minor sacrifice to the number of mode preferences satisfied and contact hours. More specifically, when comparing the resulting objective values from the hierarchical optimization (including solving the weighted objective) to their counterpart of individually optimizing the objective, the overall number of mode preferences O_p is reduced by less than 0.8%, and the number of contact hours O_h is reduced by about 10%. Moreover, considering multiple objectives leads to an increase in contact hours of about 155% compared to only optimizing O_p. Finally, we note that by solving the weighted combination of the objectives, the value of O_p increases from 1392 to 1395.