This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Optimizing Makespan and Ergonomics in Integrating Collaborative Robots Into Manufacturing Processes

Margaret Pearce, Bilge Mutlu, Julie Shah, and Robert Radwin

***Abstract*— As collaborative robots begin to appear on factory floors, there is a need to consider how these robots can best help their human partners. In this paper, we propose an optimization framework that generates task assignments and schedules for a human–robot team with the goal of improving both *time* and *ergonomics* and demonstrate its use in six real-world manufacturing processes that are currently performed manually. Using the strain index method to quantify human physical stress, we create a set of solutions with assigned priorities on each goal. The resulting schedules provide engineers with insight into selecting the appropriate level of compromise and integrating the robot in a way that best fits the needs of an individual process.**

***Note to Practitioners*—Collaborative robots promise many advantages on the shop and factory floor, including low-cost automation and flexibility in small-batch production. Using this technology requires engineers to redesign tasks that are currently performed by human workers to effectively involve human and robot workers. However, existing quantitative methods for scheduling and allocating tasks to multiple workers do not consider factors, such as differences in skill between human and robot workers or the differential ergonomic impact of tasks on workers. We propose a method to analyze how the inclusion of a collaborative robot in an existing process might affect the makespan of the task and the physical strain the task places on the human worker. The method enables the engineer to prioritize and weigh makespan and worker ergonomics in creating schedules and inspect the resulting task schedules. Using this method, engineers can determine how the addition of a collaborative robot might improve makespan and/or reduce job risk and potential for occupational hazard for human workers, particularly in tasks that involve high physical strain. We apply our method to six real-world tasks from various industries to demonstrate its use and discuss its practical limitations. In our future work, we plan to develop a software tool that will assist engineers in the use of our method.**

***Index Terms*— Human–robot teaming, manufacturing scheduling, occupational ergonomics, optimization methods, work-related musculoskeletal injuries.**

Manuscript received August 6, 2017; accepted November 20, 2017. This paper was recommended for publication by Associate Editor G. Nejat and Editor J. Wen upon evaluation of the reviewers' comments. This work was supported by the National Science Foundation under Award 1426824. *(Corresponding author: Bilge Mutlu.)*

M. Pearce and B. Mutlu are with the Department of Computer Sciences, University of Wisconsin–Madison, Madison, WI 53706 USA (e-mail: mlpearce@cs.wisc.edu; bilge@cs.wisc.edu).

J. Shah is with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: julie_a_shah@csail.mit.edu).

R. Radwin is with the Department of Industrial and Systems Engineering, University of Wisconsin–Madison, Madison, WI 53706 USA (e-mail: rradwin@wisc.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TASE.2018.2789820

I. Introduction

COLLABORATIVE robots are an emerging technology poised to change the way the work is done in manufacturing settings. To ensure safety, traditional robots have been physically caged off from humans, performing work that is isolated from human workers. Collaborative robots, on the other hand, are designed as underactuated systems to limit speed and forces that they can apply and, thus, are safe for humans, to be around humans, and work with humans as true partners [1]–[3]. Despite their rise in popularity, integrating a collaborative robot into a work process poses many challenges [4]. While production time and efficiency have often been used as criteria for assessing the performance of manufacturing teams [5]–[8], worker ergonomics for a given job must also be considered. Repetitive motion injuries and operator fatigue may be prevented by distributing the work in a manner that reduces physical stress on the human worker while optimizing the work allocation so that humans and robots are performing tasks that they are most suitable to perform.

In this paper, we jointly consider *time* and *ergonomics* in the introduction of a collaborative robot into a work cell. Illustrated at a high level in Fig. 1, we present an optimization framework that allocates work between a single human and a single robotic worker while minimizing production "makespan" (task completion time) and human physical "strain." Using real-world tasks, we demonstrate the tradeoffs in makespan and physical strain, as the importance of each factor is varied. The resulting schedule alternatives provide process engineers and other stakeholders involved in integrating collaborative robots with insight into how multiple factors may be compared in process design.

In Section II, we discuss prior work in human–robot teaming. We introduce the strain index (SI) method in Section III and present our approach for capturing time and ergonomics along with our optimization framework in Section IV. In Section V, we describe six real-world processes that cover a range of industries. We discuss the data that we obtained for each process and how it was translated into input for the optimizer. Section VI presents the results of applying our method to the process data followed by a broader discussion of the results in Section VII.

II. Related Work

Prior work has primarily proposed methods for planning and scheduling work for human–robot teams to maximize performance indicators. An emerging body of the literature has considered ergonomic constraints in and consequences of

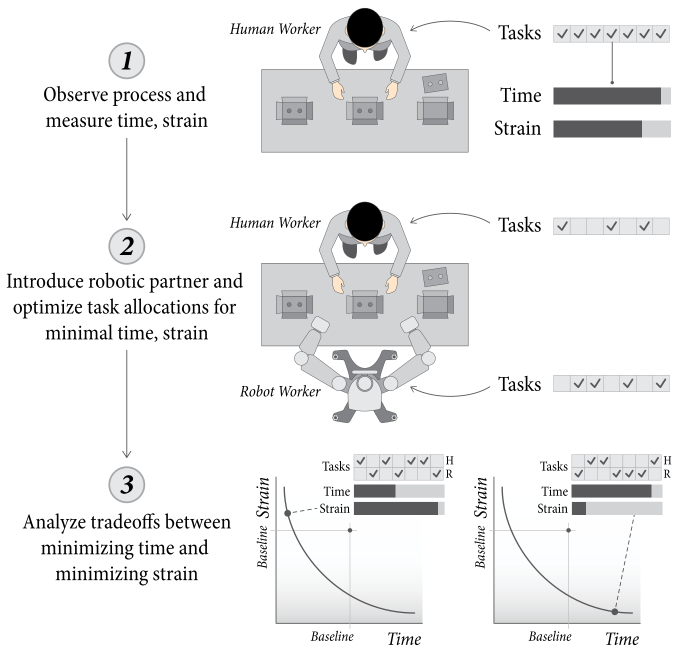

Fig. 1. In this paper, we propose an optimization-based approach to generate task assignments and schedules to integrate a collaborative robot into a currently manual manufacturing process. The optimizer provides a set of schedules that enable process engineers to consider tradeoffs between minimizing makespan and physical strain on the human worker.

task allocation. In this section, we briefly describe existing scheduling approaches and recent work considering human ergonomics.

Several optimization-based approaches have been proposed to find task assignments and schedules by maximizing a global reward and minimizing a cost quantity for multiagent teams. Koes *et al.* [5] describe a mixed-integer linear program (MILP) for multirobot teams that maximizes "utility"—the sum of rewards for tasks evaluated at the time that they were accomplished. In this search-and-rescue scenario, rewards decrease linearly with time. Similarly, Ponda *et al.* [6] present a task-allocation algorithm for human operators and robotic agents that minimizes task execution times and cost while maximizing efficiency. Assembly-time cost and payment cost have also been explored as a tradeoff [7]. A human–robot scheduling algorithm by Korsah [9] maximizes the difference between an overall reward and costs, such as overall travel and waiting/delay costs.

Other approaches involve scalable models that can solve assignment and scheduling problems for a large number of agents in minimal time, such as Tercio—an efficient task assignment and scheduling algorithm that supports on-the-fly replanning for a large number of tasks and agents [8].

The concept of using intelligent assist devices for reducing physical stress in manufacturing operations was introduced by Krüger *et al.* [10]. Robotic task sharing for the purpose of relieving the human operator from excessive physical demands of work, as well as for reducing human errors, was previously considered from an ergonomics perspective by Ogorodnikova [11].

Research on the allocation of manufacturing tasks to human and robot operators has focused mainly on determining and minimizing safety risks, although more recent work has considered the ergonomic benefits of the integration of robots into task plans for human workers [12]–[14]. Farber *et al.* [15] described an approach for considering ergonomics as well as human movement capability to ascertain the optimal assembly sequences for human–robot collaboration. They considered cognitive and physical properties as well as functional-allocation criteria that minimized ergonomically poor work conditions while minimizing the number of human–robot changes in the workflow. Execution time between assembly steps allocated to the human and robot utilized the methods-time measurement predetermined time system, while ergonomic factors were based on a linear combination of the Ovako Working Posture Analysis Systems, which is limited to categorize common postures. This approach was not applied to actual tasks, as not all influencing factors identified could be quantified.

While prior work has recognized the importance of considering ergonomics in task allocation and scheduling for human–robot teams, the literature, to date, focuses primarily on ergonomic *consequences* of integrating robots into task plans. Furthermore, research that has considered the ergonomic *characteristics* of tasks in the planning of tasks for human–robot teams does not offer a flexible, operational framework for quantifying human ergonomics and for exploring the tradeoffs between ergonomic and productivity benefits of possible task plans. This paper closes this gap by presenting an optimization-based approach that utilizes hierarchical modeling to quantify physical stress in assembly operations at the work-element level and generates human–robot task plans optimized to jointly improve production time and ergonomics. The resulting plans for task allocation and scheduling serve as a decision-making aid for engineers to determine whether or not the integration of robots would benefit a given operation and to explore tradeoffs between the productivity of the operation and its ergonomic impact on human workers.

III. Approach

In this section, we present our approach to compute task schedules and assignments by jointly considering time and ergonomics to integrate a collaborative robot. We first define key terms to describe *work* at different levels. Next, a process for quantifying physical stress is introduced along with our method of quantification. Finally, the optimization framework and its mathematical formulation are described. While humans are the best at performing work that requires fine and adaptive motor adjustments, processing complex information, and responding to unexpected conditions and high flexibility, robots excel in areas that physically stress humans. Collaborative robots are well suited to assume work involving awkward postures and orientations, repetitive motions, and tight grips and sustained forces for long periods of time—all recognized risk factors for musculoskeletal injuries and fatigue [16].

A. Definitions

Throughout this paper, the definitions of "job," "task," and "work element" as described by Radwin *et al.* [17] will be used. Under this terminology, a worker has a *job* that he/she performs during a shift. Each worker has one and only one job (e.g., an individual could have a job as a construction worker). A job is composed of several different *tasks* that may span different operational goals. Furthermore, each task in a job consists of *work elements* that describe the required steps to perform a task from start to finish. As an example, a worker might be assigned to an assembly task for part of the day and a packaging task for the rest of the day. In the packaging task, the worker might retrieve parts and place them in a box. Retrieving parts is one work element that makes up the packaging task.

B. Strain Index

To quantify ergonomics, we leverage a widely used job analysis method called SI. The SI is used to evaluate how hazardous a job is based on factors, such as force exerted, frequency of movement, and wrist posture [18]. The tasks that we evaluated for this paper relied primarily on upper limb movements, such as reaching, grasping, pushing, and pulling, making the SI score a suitable basis for measuring human physical stress.

Originally developed by Moore and Garg [18], SI has been the subject of numerous investigations demonstrating its validity in food processing and manufacturing [19]–[21]. More recently, the SI was used in a multi-institutional prospective study of upper-limb-work-related musculoskeletal disorders conducted between 2001 and 2010 among U.S. production and service workers from a variety of industries [22]–[24]. Although the SI variables are subjectively rated, it has been demonstrated that interrater reliability is relatively concurrent [25] and repeatable [26].

Moore and Garg [18] propose the following thresholds to predict job safety based on the SI.

1) SI $\leq$ 3: Safe.
2) 3 $<$ SI $<$ 7: Moderate risk.
3) SI $\geq$ 7: Hazardous.

C. Scheduler

We consider assigning work elements between a human and a robotic worker as a multiagent coordination problem with temporal and spatial constraints and formulate this problem as an MILP [27]. Although this approach does not scale well into large-scale problems due to its exponential computational complexity, the minimal team composition, involving a single human and a single robot, and the relatively smaller size of the manufacturing operations considered in this paper allowed its efficient use. We report on the performance of the algorithm in Section VI-D4 and discuss how this approach must be extended for larger problems in Section VII-E.

The goal of the MILP scheduler is to allocate work elements to individual workers and determine the sequence of work elements such that time, human physical stress, or both factors are minimized. The input to the MILP scheduler consists of the SI for each possible task breakdown along with the expected time to complete each work element for both workers. The scheduler enforces a set of constraints representing predecessor information between work elements, capabilities of the robotic worker, and basic work-element assignment rules. The output of the MILP is a set of assigned work elements for the human and the robot, the order and times that work elements should be completed by, the resulting makespan, and the SI.

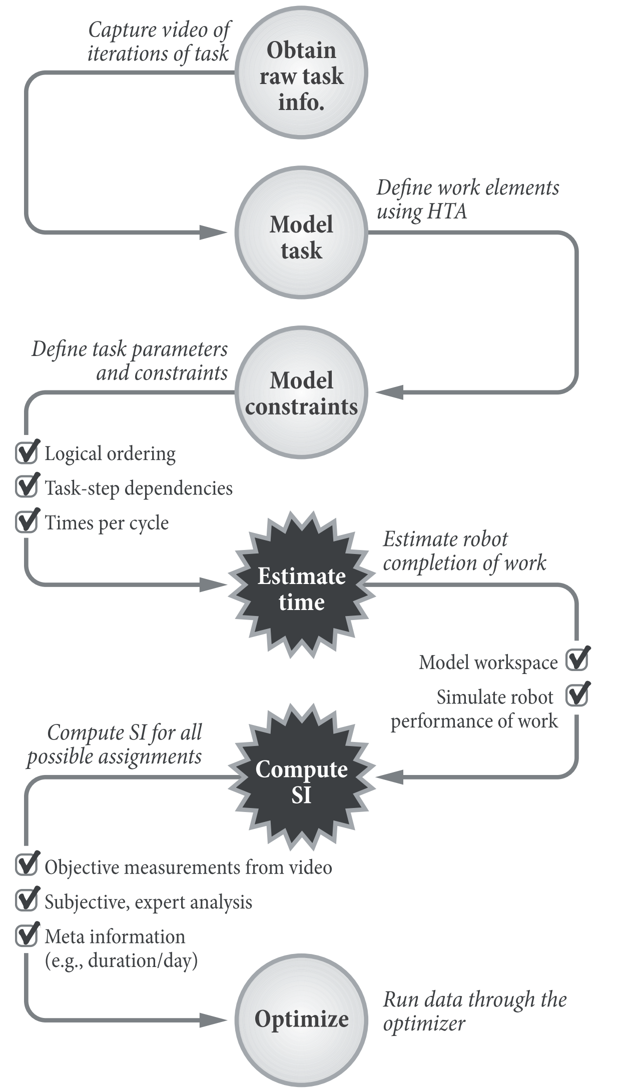

Fig. 2. High-level summary of the process followed by our approach. Capture video showing several iterations of a task. Break down tasks to work elements using HTA through a manual analysis of the videos. Given the work elements, model task parameters, and constraints from video. Estimate times for each work element as completed by the robot. Compute SI for each possible work-element assignment. Send the data to the optimizer.

We note that, while the sequence of work elements for a manual processes is optimized at design time and thus fixed, the scheduler may output schedules that are substantially different from the input sequence depending on given constraints and the complexity of the process. For example, many work elements can be parallelized between human and robot workers, which are not possible in a conventional station involving a single human worker.

IV. Implementation

A high-level summary of our approach is shown in Fig. 2. Central to our approach is two key components (highlighted with "star" shapes in Fig. 2): a method to quantify ergonomics and a task scheduling and allocation algorithm. In this section, we describe in detail our approach to quantifying ergonomics by to modeling task elements and calculating SI for a task, and present our optimizer formulation that assigns work elements to workers given task criteria.[1]

A. Modeling of Work Elements

To turn manufacturing processes into work elements that can be used in computational analysis, we utilized hierarchical task analysis (HTA), a standard task-modeling approach that has been developed and used for over a century [28]. Initially adapted from time and motion studies, HTA has been widely used to model predetermined time systems, such as manufacturing operations [29], and has been applied to assessing the ergonomic impact of these processes on human workers in order to inform job design [30]. Prior work into task allocation for human–robot teams has also utilized HTAs to model and represent manufacturing processes [31].

The HTA process iteratively breaks down plans into subplans until further breakdowns no longer provide meaningful units of analysis [32]. The resulting model includes a hierarchical plan of goals, subgoals, tasks, operations, and rules that can be expressed as hierarchical diagrams, numbered lists, or tables. An example HTA model of one of the tasks considered in this paper is shown in Fig. 4 included in Section V. In our implementation, lowest level operations from HTA models serve as work-element units used for SI analysis and scheduling of work for human–robot teams.

B. Strain Index Calculation

SI is measured by rating six parameters on a scale of 1–5: intensity of exertion (IE), duration of exertion (DE), efforts per minute (EM), hand/wrist posture (HWP), speed of work (SW), and duration per day (DD). Each parameter's rating is matched to a corresponding multiplier value. Parameter rating criteria and multipliers are presented in Table I [18]. Details for determining each parameter rating are as follows.

1) Intensity of Exertion: All work elements are ranked subjectively by their perceived IE using the scale provided in Table I. The IE rating for the entire task is the maximum value among the work-element-level IE ratings for all work elements assigned to the human worker.

2) Duration of Exertion: DE is the percentage of time throughout a cycle when exertion occurs. We sum the duration for each work element assigned to the human worker that consists of an exertion, divide it by the cycle time, and match the resulting value to the ranges provided in Table I.

3) Efforts Per Minute: In addition to duration, we record the number of efforts for each work element. For instance, if a work element consisted of grasping three screws, the number of efforts is one if they are grasped as a group but three if they are grasped one at a time. Similar to DE, we take the sum of efforts for all work elements assigned to the human worker, divide it by the cycle time in minutes, and select the corresponding rating from Table I.

[1]The codebase for our implementation of the optimizer can be found online at https://github.com/Wisc-HCI/human-robot-optimizer.

TABLE I
SI PARAMETERS AND MULTIPLIERS

Rating	*IE**	*DE*	*EM*	*HWP*	*SW*	*DD*
Rating criteria						
1	Light	< 10	< 4	Very good	Very slow	≤ 1
2	Somewhat hard	10-30	4-9	Good	Slow	1–2
3	Hard	30–50	9–15	Fair	Fair	2–4
4	Very hard	50–80	15–20	Bad	Fast	4–8
5	Near maximal	≥ 80	≥ 20	Very bad	Very fast	> 8

Rating	IE'	DE'	EM'	HWP'	SW'	DD'
Multiplier table						
1	1	0.5	0.5	1.0	1.0	0.25
2	3	1.0	1.0	1.0	1.0	0.50
3	6	1.5	1.5	1.5	1.0	0.75
4	9	2.0	2.0	2.0	1.5	1.00
5	13	3.0†	3.0	3.0	2.0	1.50

* *IE* = Intensity of exertion; *DE* = Duration of exertion; *EM* = Efforts per minute; *HWP* = Hand/wrist posture; *SW* = Speed of work; *DD* = Duration per day (h)
† If *DE* is 100%, the multiplier of effort per minute is set to 3.0.

4) Hand/Wrist Posture: The rating for HWP uses the same approach as IE: rank all work elements and then select the maximum rating among all work elements assigned to the human worker.

5) Speed of Work: For this paper, we map movement frequency ranges to the rating criteria to objectively rate SW. First, we compute the total number of movements for each work element. For each work element assigned to the human worker, we sum the number of movements and divide the total value by the cycle time. The result is a movement frequency that dictates the SW rating as follows.

- SW $= 3$ if *movement frequency* < 0.5.
- SW $= 4$ if $0.5 \leq$ *movement frequency* < 1.
- SW $= 5$ if *movement frequency* ≥ 1.

Note that the multiplier SW' has the same value for SW $= 1, 2,$ or 3, and hence, there is no need to differentiate values of SW less than 0.5.

6) Duration Per Day: DD is assessed at the task level. Since it is not directly observable by video, we either request this information separately or we use a default rating value of 4, which represents a worker performing a given task for 4–8 h.

Finally, the SI is determined by taking the product of the multipliers across each parameter

$$SI = IE' \cdot DE' \cdot EM' \cdot HWP' \cdot SW' \cdot DD'. \qquad (1)$$

C. Scheduling Algorithm

The input to the MILP includes a number of sets and parameters as outlined in the following.

1) *Sets:*
   a) a set of tasks $\boldsymbol{T}$ to be scheduled;
   b) a set of work elements $\boldsymbol{e}$ composing each task;
   c) a set of workers (human or robotic) $\boldsymbol{w}$ who will be assigned work elements;
   d) a set of discrete timepoints $\boldsymbol{p}$.

2) *Parameters:*
   a) The expected duration $D_{T,e,w}$ values to complete each work element e of a task T for a worker w (robotic or human).
   b) SI s_T values for each possible work-element assignment.
   c) *Work-Element-Predecessor Relationship Indicators:* One work element may need to be completed before another work element can begin. As an example, a worker cannot use a drill to assemble a part before the part has been brought to the work area. Furthermore, some work elements may need to be completed by the same person to ensure the validity of the current task plan. For instance, a worker may need to scan a label before placing it inside of a box. Distributing these two work elements between two different workers would require an additional work element that is not included in the original task plan, such as setting the label down at an intermediate spot or performing a handover. Instead, we require the two work elements to be assigned to the same worker, using binary parameters $P(T, e', e) \in \{0, 1\}$ and $SW(T, e', e) \in \{0, 1\}$ to indicate the applicability of these conditions, respectively.
   d) *Capability Indicators Defined at the Work-Element Level:* For example, robotic workers cannot reasonably perform certain work elements due to design or physical constraints, such as payload limitations. In addition, the robotic worker may be incapable of using certain tools or machinery. Furthermore, the robot may not have appropriate end-effectors to perform work elements, such as applying tape, tying knots, or placing packaging materials in tight spaces. Given the focus of this paper on integrating collaborative robots into existing manual processes that are currently performed by human workers, we have considered human workers to be capable of performing all work elements, although the design of new processes that may include work elements that human workers cannot perform must define capability indicators for both human and robot workers. Capability information is provided to the optimizer as a binary parameter $B(T, e, w) \in \{0, 1\}$, where $B(T, e, w) = 1$ implies that w is able to perform work element e in task T.
   e) A weight parameter α that represents the importance of minimizing makespan over strain.

Based on the input values and a set of constraints, the optimizer attempts to solve for the variables listed as follows.

3) *Variables:*
   a) a binary decision variable $a_{T,e,w,p} \in \{0, 1\}$ indicating the assignment of work element e in task T to worker w at timepoint p;
   b) C_T, a positive variable representing the completion time of task T based on the values of $a_{T,e,w,p}$;
   c) an upper bound on cycle time M_T;
   d) the resulting SI Γ_T;
   e) an upper bound on the SI S_T.

TABLE II

SCHEDULER SETS, PARAMETERS, AND VARIABLES

Sets	Description
T	Set of tasks to schedule
e	Set of work elements belonging to a task
p	Set of discrete time-points
w	Set of workers (human/robotic)
Parameters	**Description**
$D_{T,e,w}$	Estimated work element duration
s_T	Strain Index for each possible work element allocation
$P(T,e',e)$	Binary predecessor relationships for two work elements
$SW(T,e',e)$	Binary same worker indicator for two work elements
$B(T,e,w)$	Binary indicator if a worker can perform a work element
α	Weighted importance of time over ergonomics
Variables	**Description**
$a_{T,e,w,p}$	Boolean indicating e is assigned to w at p
C_T	Resulting cycle time based on scheduler output
Γ_T	Strain Index based on scheduler output
M_T	Upper bound on cycle time
S_T	Upper bound on Strain Index

All input and variables used by the optimizer are summarized in Table II.

Our implementation utilizes a third-party optimizer, CPLEX, to solve the MILP [33]. We chose to use CPLEX due to its high performance in solving linear and mixed-integer programming problems and flexible use through an application programming interface. The returned solution consists of a time-indexed schedule that describes which worker performs which work element and during what time interval. The solution must adhere to all constraints while minimizing the objective function within 1% optimality. The mathematical formulation of the problem is as follows:

$$\min\ f \triangleq \alpha \cdot \frac{C_T}{M_T} + (1-\alpha) \cdot \frac{\Gamma_T}{S_T}. \tag{2}$$

This problem is subject to the following constraints.

1) Each work element e in task T should be assigned
$$1 \le \Sigma_w \Sigma_p\, a_{T,e,w,p} \le \max(D_{T,e,w'} \forall w'), \qquad \forall T, e. \tag{3}$$
2) Each work element e in task T should not take more than its duration D based on the assigned worker w
$$\Sigma_p\, a_{T,e,w,p} \le D_{T,e,w}, \quad \forall \{T, e, w | B(T, e, w)\}. \tag{4}$$
3) Each worker w can only be assigned one work element e in task T at any given time p
$$\Sigma_e\, a_{T,e,w,p} \le 1, \quad \forall T, w, p. \tag{5}$$
4) Each work element e in task T can be assigned to at most one worker w in a given time p
$$\Sigma_w\, a_{T,e,w,p} \le 1, \quad \forall T, e, p. \tag{6}$$
5) Once started, work elements should not be interrupted or paused
$$\begin{aligned}(a_{T,e,w,p} - a_{T,e,w,p-1}) \cdot D_{T,e,w} \\ \le \Sigma_{p'=p}^{p+D_{T,e,w}} a_{T,e,w,p'}, \quad \forall T, e, w, p.\end{aligned} \tag{7}$$

6) Work elements cannot be shared between workers

$$a_{T,e,w,p} \cdot D_{T,e,w} \leq \Sigma_p \, a_{T,e,w,p}, \quad \forall T, e, w, p. \tag{8}$$

7) If a worker w is assigned work element e' with predecessor e at time p, then e must be completed before time p

$$\Sigma_w \frac{\Sigma_{p'=0}^{p} a_{T,e,w,p'}}{D_{T,e,w}} \geq a_{T,e',w',p}, \quad \forall\{T, e, e', p, w, w' | P(T, e', e)\}. \tag{9}$$

8) Work element e' with predecessor e must be completed by the same worker w if indicated by $SW(T, e', e)$

$$\Sigma_{p'=0}^{p} a_{T,e,w,p'} \geq D_{T,e,w} \cdot a_{T,e',w,p}, \quad \forall\{T, e, e', p, w | P(T, e', e), SW(T, e', e)\}. \tag{10}$$

9) The task T is completed when all its work elements e have been completed

$$p \cdot a_{T,e,w,p} \leq C_T \leq M_T, \quad \forall p, T, e, w. \tag{11}$$

10) Workers should only be assigned work elements that they are capable of performing

$$a_{T,e,w,p} = 0, \quad \forall\{T, e, w, p | !B(T, e, w)\}. \tag{12}$$

11) The SI Γ_T is equal to the precomputed SI value based on the subset of work elements e assigned to the human worker

$$\Gamma_T = s_{T*} \mid a_{T,e,w,p} == a_{T*,e,w,p}, \quad \forall e, w. \tag{13}$$

12) The upper bound on time M_T is greedily approximated by summing the longest duration for each work element e among all workers w

$$M_T = \Sigma_e \max(D_{T,e,w}, \, \forall w). \tag{14}$$

13) The upper bound on the SI S_T is equal to the precomputed SI value for the scenario when all work elements e are assigned to the human worker w

$$S_T = s_{T*} \mid a_{T,e,w,p} == a_{T*,e,w,p}, \quad \forall e. \tag{15}$$

V. Case Studies

In this section, we apply the approach, which we described earlier to six real-world factory tasks, and explore the results of varying the importance of makespan versus ergonomics.

A. Tasks

The case studies consist of six real-world processes that are currently manually performed by human workers, including three observed during a factory site visit at Steelcase, Inc., a furniture manufacturer-based in Grand Rapids, MI, USA [35], and three chosen from a video task database developed by the U.S. National Institute for Occupational Safety and Health (NIOSH) [36]. These tasks were chosen from among larger set of candidate tasks for including work elements that can be performed by currently available collaborative robot platforms.

Fig. 3. Worker performing the packaging task (T2) at a Steelcase, Inc., plant in Grand Rapids, MI, USA.

The three tasks currently performed at Steelcase included a quality-control task (T1), a packaging task (T2), and an assembly task (T3). Each task involved a single human worker performing all work elements. In T1, a worker gathers parts as indicated by an order, uses specialized equipment to validate that the correct selection and quantity of parts have been retrieved, and places them in a package. Task T2 involves a worker packaging supplies in a bag and securing the items inside of a box. Finally, T3 is an assembly task where a worker connects a smaller part to a larger piece that is being built. Specific details on tasks are excluded from this paper to protect confidentiality of the processes. Fig. 3 shows a worker performing T2 at the Steelcase facility.

In addition, we examined three tasks recorded on video and included in the NIOSH task database: a stocking task (T4), a parts-assembly task (T5), and a metal-cutting task (T6). The stocking task (T4) video shows a worker unpacking items from a crate and placing them on shelfs, occasionally stepping away to organize empty crates. In T5, a worker fits a rubber piece around a large part and stacks it to the side of the workspace. The metal-cutting task (T6) requires positioning a metal sheet on a cutting machine, running the machine, setting aside the cut piece, and discarding the scrap parts.

B. Data Collection

We observed and documented tasks T1–T3 at a site visit to a Steelcase facility carried out March 2015. Multiple cycles of each task were observed and captured on video during the site visit that also included interviews with workers. Existing video was used for tasks T4–T6 from the NIOSH task database. Unlike the tasks we observed at the Steelcase

Task # Plans/Operations

0. STOCK SHELVES
Plan 0: Do 1.–2. or 5.–6. When crate is empty, do 3. When all crates are empty, do 4. EXIT
1. Stock products on top shelf
1.1. Grab item from crate
1.2. Place item on top shelf
2. Organize item on top shelf
2.1. Examine item to find right place
2.2. Position item(s) on top shelf
3. Move crate
3.1. Pick up crate
3.2. Set crate off to the side
4. Organize crates
4.1. Grab crate from side
4.2. Stack on top of last empty crate
4.3. Wheel crates outside work area
5. Stock products on bottom shelf
5.1. Grab item from crate
5.2. Place item on bottom shelf
6. Organize item on bottom shelf
6.1. Examine item to find right place
6.2. Position item(s) on bottom shelf

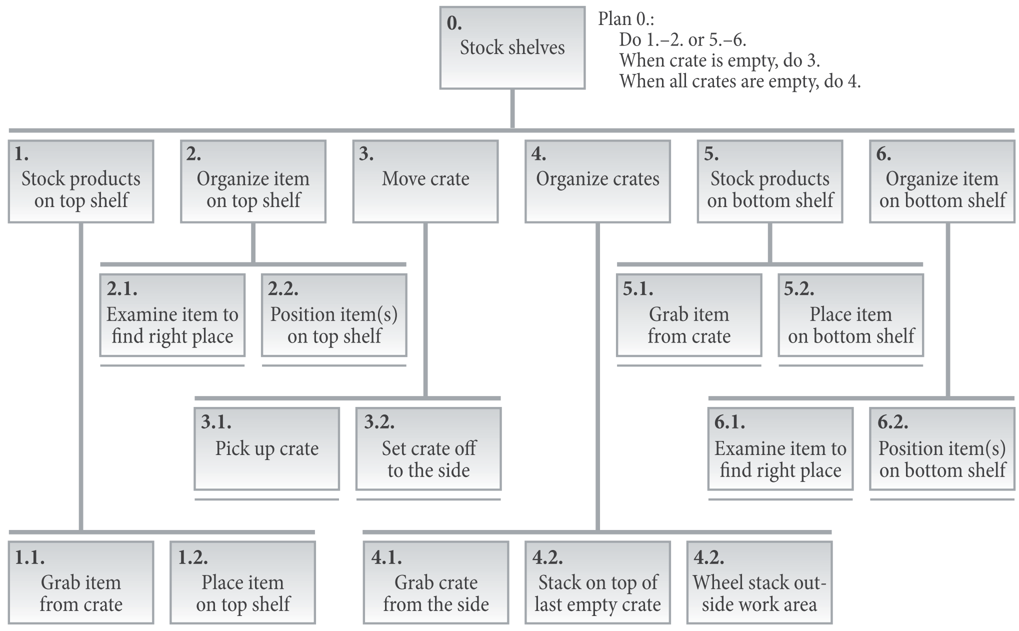

Fig. 4. HTA model of T4 presented in numbered-list and hierarchical-diagram formats based on standard conventions of HTA modeling [29], [34]. In the task, a worker restocks the pantry of a commercial kitchen one crate at a time. The full HTA model, in addition to the goals, tasks, and operations shown earlier, includes task parameters, such as the number of repetitions, duration, and cycle time.

facility, the NIOSH database did not include additional data on the tasks beyond what we could observe from video.

Task videos were analyzed to define the task at the work-element level using the approach described in Section IV-A. Using HTA, plans were expanded to lower levels on a hierarchy until plan steps cannot be further subdivided in a meaningful way. For this analysis, each task was broken down into work elements such that each work element could not be assigned to more than one worker. Fig. 4 shows the resulting HTA model for T4 in list and diagrammatic formats.

Within each task, the work elements that a robotic worker could feasibly perform were highly dependent on the capabilities of the robot. In this paper, we considered the capabilities of the Baxter robot developed and manufactured by Rethink Robotics [3]. We chose to evaluate Baxter as a robotic partner due to its widespread use as a collaborative robot. We classified each work element of the tasks according to whether or not it could be reasonably completed by Baxter based on hardware limitations [37], known use cases, and available standard end-effectors.

Manually annotated video data were the source of many inputs required by the scheduler. Variations of the task from one repetition to the next in addition to logical orderings were used to determine work-element predecessor information. For instance, if two work elements were performed in alternating orders within multiple cycles, it is inferred that neither work element is dependent on the other to be completed. In some cases, predecessors were evident based on how a task was broken down into work elements. As an example, a worker on a packaging task should not close a package box shut until all necessary parts have been placed inside the box.

Work elements were timed for each recorded cycle. The recorded times were averaged to obtain the expected duration of each work element for a human worker. The tasks that we observed were performed by a single human worker at the time of our site visit, and thus, directly measuring the expected duration of each work element as performed by Baxter was not possible. Instead, durations for Baxter were estimated using MoveIt!, a motion planner for robots [38]. A 3-D mockup of each work cell was created to scale based on measurements taken on site or estimated from video. Next, a model of Baxter was added to the work cell. The location of the robot needed to be out of the way of the human worker. This ensured that the human worker could still perform a task in the same way regardless of the presence of the robot, and therefore, the durations measured for human workers remained valid. For T1–T6, it is assumed that a duplicate set of all relevant parts and supplies could be placed within reach of the robot. Within each mockup, points of interaction corresponding to work elements were marked as waypoints. Finally, an RRT-Connect [39] planner within MoveIt! (`RRTConnectkConfigDefault`) was used to estimate the duration of each motion for Baxter.

Finally, the SI was computed for each work-element assignment combination in T1–T6. Objective time- and movement-based measures, such as DE, EM, and SW, were calculated directly from the video data. Subjective measures, such as IE and HWP, were estimated from video data and reviewed with experienced job analysts. For T1–T3, we knew the DD directly from Steelcase. Because these data were not available for T4–T6, we assumed that the DD was between 4–8 h/day. This range results in a multiplier value of 1, and thus, it does not have an effect on the resulting SI but limits our ability to make comparisons across jobs.

TABLE III

BASELINE VALUES FOR MAKESPAN AND TOTAL STRAIN, PERCENT CHANGE IN THESE VALUES, AND PERCENTAGE OF TOTAL IDLE TIME FOR HUMAN AND ROBOT WORKERS

Task		*Baseline*	*Alpha*										
			0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Task T1	Makespan (s)	38	90	76	76	59	59	35	35	35	31	28	28
	Strain Index	18	2.25	2.25	2.25	3.375	3.375	6.75	6.75	6.75	9	12	12
	Change in makespan (%)		136.84	100	100	55.26	55.26	-7.89	-7.89	-7.89	-18.42	-26.32	-26.32
	Change in SI (%)		-87.5	-87.5	-87.5	-81.25	-81.25	-62.5	-62.5	-62.5	-50	-33.33	-33.33
	Human idle time (%)		87.78	85.53	85.53	71.19	71.19	17.14	31.43	31.43	12.9	3.57	3.57
	Robot idle time (%)		15.56	0	0	3.39	3.39	37.14	11.43	11.43	0	0	0
Task T2	Makespan (s)	25	65	39	39	39	39	39	39	30	25	25	25
	Strain Index	18	9	9	9	9	9	9	9	13.5	18	18	18
	Change in makespan (%)		160	56	56	56	56	56	56	20	0	0	0
	Change in SI (%)		-50	-50	-50	-50	-50	-50	-50	-25	0	0	0
	Human idle time (%)		73.85	56.41	56.41	56.41	56.41	56.41	56.41	30	0	0	0
	Robot idle time (%)		53.85	23.08	23.08	23.08	23.08	23.08	23.08	43.33	100	100	100
Task T3	Makespan (s)	55	49	49	49	49	49	49	49	36	36	36	36
	Strain Index	3.375	0.375	0.375	0.375	0.375	0.375	0.375	0.375	1.5	1.5	1.5	1.5
	Change in makespan (%)		-10.91	-10.91	-10.91	-10.91	-10.91	-10.91	-10.91	-34.55	-34.55	-34.55	-34.55
	Change in SI (%)		-88.89	-88.89	-88.89	-88.89	-88.89	-88.89	-88.89	-55.56	-55.56	-55.56	-55.56
	Human idle time (%)		53.06	53.06	53.06	53.06	53.06	53.06	53.06	0	0	0	0
	Robot idle time (%)		10.2	10.2	10.2	10.2	10.2	10.2	10.2	50	50	50	50
Task T4	Makespan (s)	422	820	500	500	500	500	500	500	370	370	370	370
	Strain Index	9	6	6	6	6	6	6	6	9	9	9	9
	Change in makespan (%)		94.31	18.48	18.48	18.48	18.48	18.48	18.48	-12.32	-12.32	-12.32	-12.32
	Change in SI (%)		-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	0	0	0	0
	Human idle time (%)		58.54	32	32	32	32	32	32	0	0	0	0
	Robot idle time (%)		56.1	28	28	28	28	28	28	40.54	40.54	40.54	40.54
Task T5	Makespan (s)	38	53	53	53	53	53	53	42	42	42	38	38
	Strain Index	60.75	27	27	27	27	27	27	40.5	40.5	40.5	60.75	60.75
	Change in makespan (%)		39.47	39.47	39.47	39.47	39.47	39.47	10.53	10.53	10.53	0	0
	Change in SI (%)		-55.56	-55.56	-55.56	-55.56	-55.56	-55.56	-33.33	-33.33	-33.33	0	0
	Human idle time (%)		49.06	49.06	49.06	49.06	49.06	49.06	28.57	28.57	28.57	7.89	0
	Robot idle time (%)		50.94	50.94	50.94	50.94	50.94	50.94	64.29	64.29	64.29	84.21	100
Task T6	Makespan (s)	65	73	57	57	57	57	57	57	57	57	57	57
	Strain Index	27	18	18	18	18	18	18	18	18	18	18	18
	Change in makespan (%)		12.31	-12.31	-12.31	-12.31	-12.31	-12.31	-12.31	-12.31	-12.31	-12.31	-12.31
	Change in SI (%)		-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33	-33.33
	Human idle time (%)		32.88	19.3	19.3	19.3	19.3	19.3	19.3	19.3	19.3	19.3	19.3
	Robot idle time (%)		79.45	45.61	45.61	45.61	45.61	45.61	45.61	45.61	45.61	45.61	45.61

C. Baseline Data

We use the measured work-element times from video data as the baseline makespan for each task prior to introducing Baxter. Similarly, we calculate the SI based on all work elements being assigned to the human worker, as currently performed in these processes, to establish a baseline for the SI. These two starting points enable us to compare the benefit of introducing a robotic worker from both time and ergonomics perspectives.

VI. RESULTS

For each task, we varied α, the weight of time versus ergonomics, to range from 0 to 1 in increments of 0.1. We obtain a Pareto bounded-optimal set of schedules for each task based on the relative importance of the two factors. Table III presents time, total strain, and idle time information.

A. Makespan Versus Total Strain

Makespan and total strain roughly vary inversely in the Pareto results; as makespan increases, total strain decreases. At the limits of α, the gains for one factor may or may not justify the losses in the other. In general, values at $\alpha = 0$ and $\alpha = 1$ tend to be skewed in one direction over the other, and values of α in the middle tend to show subtler tradeoffs. Solutions that offer significant improvement in makespan and/or total strain are highlighted with circles in Fig. 5. The flexibility in this approach is that many scheduling options are presented for process engineers to choose from. It is unrealistic to expect time and ergonomics to be the only factors considered when designing a process, and it is also impractical to attempt to consider all possible factors. Having a set of schedules enables process engineers to select an ideal match or a subset of schedules that suit their goals for the processes in consideration.

B. Changes From Baseline Data

The makespan of a task may increase or decrease after introducing Baxter depending on α, but the total strain cannot increase. We consider the strain of a work element on the human worker to be 0 if it is completed by a robot, and

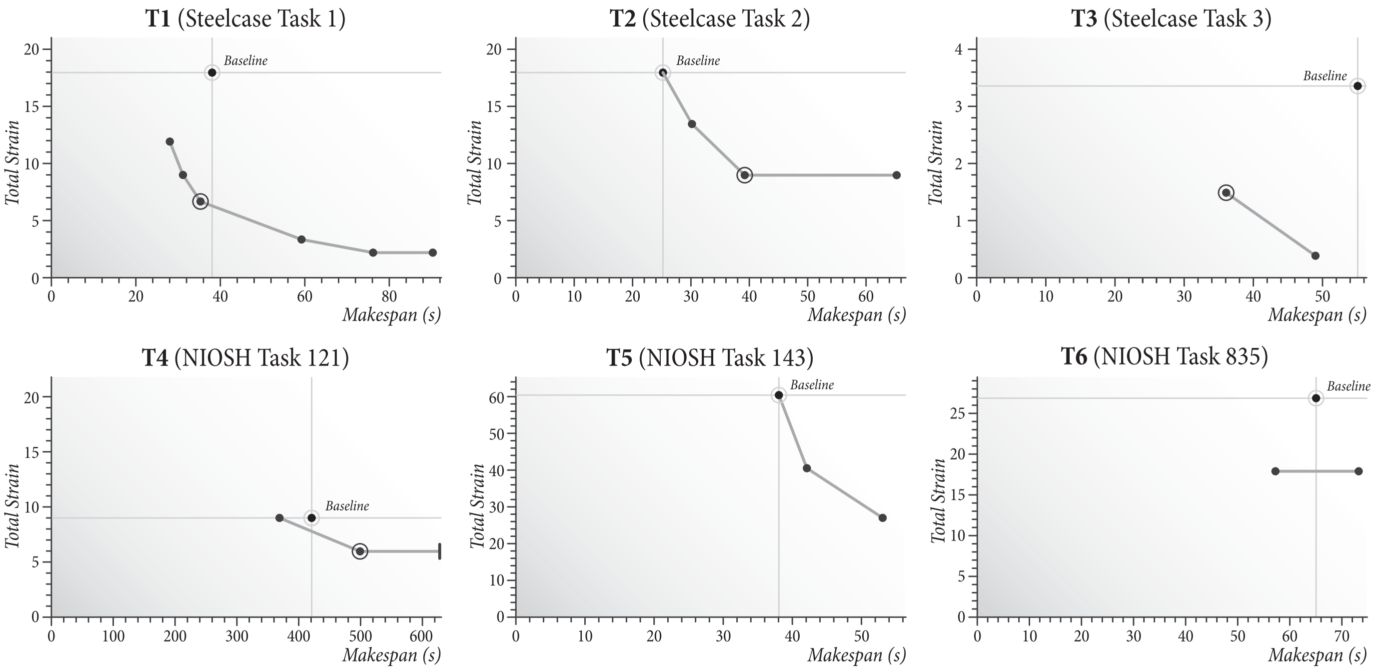

Fig. 5. Makespan and SI results for each task, illustrating the tradeoffs between the two factors. As work elements are assigned to the robot, strain on the human worker reduces, although makespan can increase due to the slower pace in which collaborative robots perform tasks compared with human workers. Baseline values for each task represent makespan and strain measurements taken, as a single human worker performs the task. The circled solutions in T1–T4 involve significant gains in makespan and/or strain. Gains in T5–T6 are not considered significant.

thus, any work elements assigned to Baxter will decrease the total strain. Therefore, we see noticeable improvements in total strain, as α approaches 0. For the tasks we evaluated, decreases in makespan compared with the baseline are less extreme due to Baxter's physical limitations. A collaborative robot inherently compromises speed and force capabilities in favor of the safety of nearby workers. Although Baxter may take longer duration to perform work elements, having a second worker functioning in parallel with a human worker can still speed up the task as a whole. In cases where Baxter was incapable of completing certain work elements or numerous same-worker predecessor restrictions applied, there may not have been as many opportunities for the process to benefit from a robot collaborator. This case can be seen in Task T2, which saw no improvements in makespan and only modest improvements in total strain. Fig. 5 shows the changes in total strain and makespan over the initial baseline for all tasks analyzed in the case study.

TABLE IV

SI Risk Classifications

Task	*Baseline*	*Alpha*										
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
T1	High	Safe			Moderate					High		
T2	High	High										
T3	Moderate	Safe										
T4	High	Moderate							High			
T5	High	High										
T6	High	High										

Strain Index can be used to understand how hazardous a job is by predicting the risk of injury. Using results from Table III, we map SI values to their risk classifications and discuss how successful the collaborative robot was for reducing the risk of injury.

C. Job-Risk Categorization

In Table IV, we note the job risk based on the SI for each task. For tasks T2, T5, and T6, although we saw reductions in the SI numbers, the overall risk did not change. The other tasks saw more success; T3 was brought down from moderate risk ($3 < \text{SI} < 7$) to low risk ($\text{SI} \leq 3$) regardless of the α value. Thus, greater time improvements can be achieved without a strong compromise to physical stress. Task T4 could be brought from high ($\text{SI} \geq 7$) to moderate risk but only within certain values of α that all result in an increased makespan compared with the baseline value. Finally, T1 presents the most variety in job-risk categorization results. The task could be brought down from high risk to low risk, again at the cost of time. However, for α within 0.5–0.7, we are able to bring the risk down to a moderate level while seeing modest improvements in time.

D. Schedules and Assignments

1) Task Allocations: Fig. 6 shows the effective task allocations for each value of α considered in our analysis for T1. The schedules illustrate several important points about the optimization output. First, looking at the work elements that are always assigned to the human worker tells us, which tasks are not worth assigning to the robot. These work elements may not be possible for the robot to perform, or they might simply take the robot too much time without enough gain in SI. Next, we can see how the task structure affects the variety of scheduling options that the optimizer has to choose from. Furthermore, some tasks may find different scheduling

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

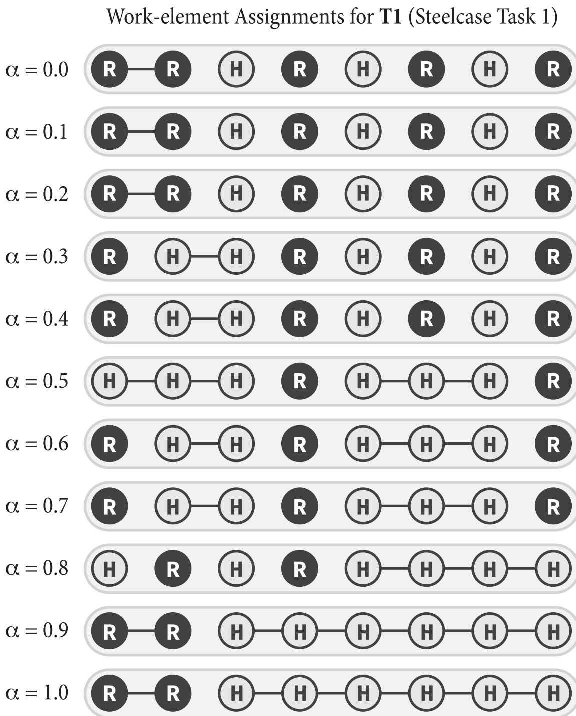

Fig. 6. Work-element allocations for human (H) and robot (R) workers for T1 under different values of α.

options to offer little or no change in makespan or SI. For instance, task T1 has the same makespan and SI results for $\alpha = 0.5$ or 0.6, but the task allocations are not identical. Comparing the task allocations may give process engineers insight into which option is preferable based on their own criteria and preferences.

2) Working and Idle Time: An additional metric, *idle time* of each worker, is explored for each schedule, as it is a common consideration in human–robot teaming [40], [41]. Fig. 7 shows idle time calculations at every α level for human and robot worker for each task. While it is undesirable for either worker to be idle for a large percentage of the task duration, a certain range of idle times may be acceptable. It may even be preferable to include a minimum amount of idle time to provide the human worker with an opportunity to recover from fatigue. The idle time metric is another example of data that is helpful to present to a process engineer to help them distinguish between acceptable and unacceptable schedules.

Idle time may indicate that there is opportunity for either the human or the robotic worker to multitask or absorb additional responsibilities outside of the current process, particularly for the robot. If additional work is assigned to the human worker, process engineers must ensure that SI gains resulting from the collaborative robot partner are not canceled out. However, if the robot is able to do more work during its downtime for a process, the robot could be more productive overall and could yield a higher return on investment.

3) Impact on Production: By reducing SI and makespan, it may be possible to increase the number of cycles of a job that can be accomplished in an 8-h period. For instance, consider task T3 where $\alpha = 0.8$. Assuming an 8-h production shift with a baseline makespan of 55 s, the task can be performed 523 times a day without any unexpected stops. By reducing the makespan by 34.55% down to 36 s, production can increase up to 800 times a day over the same 8-h period. However, realizing such benefits will require the application of the proposed approach at the level of a production line and the consideration of several sources of uncertainty, such as errors in coordination between human and robot workers, in the estimation of makespan.

4) Algorithm Performance: For each task that we evaluated, the optimization algorithm must find an assignment for all workers, all work elements, and for all 1-s time intervals at each desired value of alpha. Therefore, the problem that the optimizer solves scales linearly with the upper bound on makespan, the number of workers, and the number of work elements. In practice, the number of time intervals that the algorithm must solve for is the biggest factor that increases the size of the problem. Although this limits the scalability of the algorithm to shorter tasks, it was acceptable for our purposes, because we were interested in short, repetitive tasks.

Each of the case study tasks had 20 or fewer work elements, and the tasks were modified to plan for two workers (one human and one robot). The baseline times for the tasks ranged from 25 up to 422 s. Thus, in the worst case, the optimizer was solving for under 20 000 variables in the case studies presented. Solving within 1% of optimality using CPLEX, the time to solve the optimization problem ranged from as low as 0.32 s to as high as 37.88 s (task 1–6 took an average of 22.77, 1.46, 12.20, 4.44, 1.47, and 0.49 s for each alpha value, respectively). The high variability of the run time depends on the computational complexity of the problem and the ability of the optimizer to quickly eliminate infeasible solutions, such as work elements, that would drastically increase the makespan if they were assigned to the robot.

VII. Discussion

In this section, we explore the key takeaways from the case study results and discuss the overall promise of our approach. We identify that this approach will benefit some tasks but not others and describe the characteristics for both scenarios. We examine the observed behavior of the SI. Finally, we recognize the limitations of this method and highlight the opportunities for future research.

A. Tasks That Benefit From Our Approach

As demonstrated in the case studies, some tasks saw more improvements in time, SI, or both compared with other processes. Tasks that enabled for parallel work, had sufficient opportunities for the robot to participate, were repetitive, and included poor HWP benefited more from our method.

Consider a task scenario where a large percentage of the work elements cannot be done in parallel. If all work elements must be done sequentially, the task will not see strong benefits from collaboration of any sort. There will be minimal improvements on time, though improvements to the SI are still possible. If more work can be done in parallel, then both workers can be active for a larger percentage of the time,

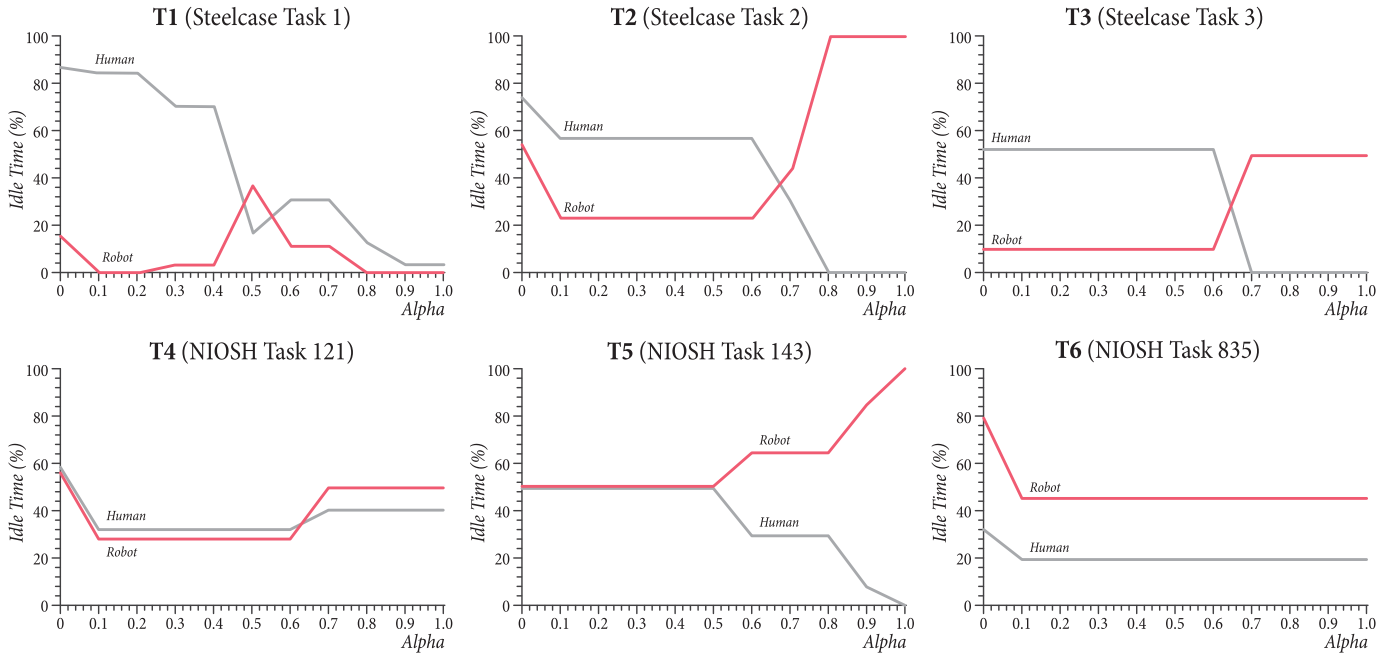

Fig. 7. Idle times for the human (gray) and robot (red) workers for each task. Idle time allows the human worker to recover from fatigue. Based on the physical demands of the task, an acceptable idle time for the human worker may vary. Similarly, idle time of the robot may provide it with opportunities to multitask.

which can decrease total makespan. It is also beneficial for the optimizer when the robot is capable of performing a more diverse set of work elements. For example, if the collaborative robot is able to perform the majority of the work elements composing the task, the optimizer will have more options when evaluating schedules and may find solutions that are more optimal.

Tasks involving poor HWP can benefit from our approach. Of the six parameters that make up the SI, HWP is generally a good indicator that the offending work element(s) should be assigned to the robot whenever possible. Jobs that mandate frequent pinching or overextended wrists cause the rating of HWP to increase. By contrast, collaborative robots, such as Baxter, can have a wider range of motion for the hands compared with a human worker and can perform pinching motions repetitively and consistently. Reallocating work elements attributed to poor HWP to the robot removes the higher rating from the SI calculation. Finally, tasks with a high degree of repetition may benefit from the integration of robots, as robots are designed to perform repetitive motions for long periods of time. The EM parameter can be reduced for the overall task by allocating repetitive work elements away from the human worker and assigning them to the robot.

B. Tasks That Do Not Benefit From Our Approach

We have also identified scenarios that yield poor results from our approach. Specifically, we do not see strong improvements for time and SI in tasks that feature high force exertions, occur in settings that are unsafe for the robot to operate, or require precise dexterity or complex tools.

Depending on the collaborative robot model, lifting capabilities of the robot are limited. For instance, Baxter's maximum payload (including the end-effector) is 5 lb (2.2 kg). Other robots can support larger weights, such as Universal Robots UR10 industrial robot arm with a 22-lb (10 kg) payload limit [42]. As collaborative robots by design cannot apply as much force as a traditional industrial robot, tasks with high SI due to hard exertions will not see significant benefits based on IE alone. Next, it is important to consider the environment where the task is performed. Collaborative robots have safe operating temperature ranges and are not waterproof. A dishwashing task, for example, may not offer many opportunities for the robot to contribute due to its safety requirements. Such scenarios may also not include many work elements that the robot is able to perform. Finally, tasks that require fine motions, specialized tools, or rely strongly on the sense of touch are unlikely to benefit from this approach given the capabilities of currently available robot platforms and the need to equip them with specialized end-effectors. Again, these tasks will limit the number of work elements that the robot can perform, therefore limiting the variety of feasible scheduling solutions and restricting results.

C. Flat Strain Index Versus Makespan

For several case study tasks, we see that the SI versus makespan curve remains flat even as α increases. This result can occur when the SI value is largely controlled by a small number of factors or when SI is heavily influenced by one or two work elements. For example, consider a task that has uniformly low SI parameter ratings for each parameter except for IE. The IE rating is determined by finding the maximum exertion throughout the entire job. Thus, even if the robot takes over all of the work elements aside from the exertion, the SI will not be significantly affected. Likewise, if every work element is ergonomically low risk aside from a single work element, the SI will remain stagnant.

D. Extensions of the Presented Framework

In this paper, we presented a framework for considering worker ergonomics, in addition to work performance, in task allocation and scheduling for human–robot teams. The application of this framework to a number of existing manual processes illustrates the promise of the approach in task schedules that are overall more efficient and less strenuous for human workers, or both. We expect, however, this approach to be even more beneficial in the design of new processes for human–robot teams, as it can be applied at the level of a multistation work cell in which multiple humans and multiple robots work together; robots with specialized tooling can serve multiple stations to perform unergonomic work elements; and more capable collaborative robot platforms can enable the introduction of work elements that are currently considered unsafe or unergonomic for human workers. However, the success of the proposed approach at the level of a work cell or a production line will require a careful consideration of several sources of uncertainty, such as individual differences in responding to physical strain and errors in coordination between human and robot workers. In addition, while this paper explored the use of the framework at design time, task allocation and schedules can also be further optimized at performance time through real-time monitoring and measurement of changes in task performance and ergonomic strain.

E. Limitations

Our approach has a number of limitations. First, because our optimizer uses a time-index approach, tasks with a long makespan increase the overall size of the MILP. As a result, the optimizer takes more time to converge to a solution within the optimality bounds. A separate optimization problem must be solved for each value of α, compounding the problem. In the case studies, task T4 took significantly longer duration to find a solution compared with the other tasks. Addressing this problem and considering more complex team compositions, involving multiple human and robot workers, and more complex tasks would require a more scalable approach. Prior work has explored hybrid algorithms that couple the MILP formulation with constraint programming to limit the search space [43] or heuristic schedulers to perform more effective search [44]. Other approaches achieve efficiency and scalability by using a multiagent sequencer to decompose the MILP [8] or utilizing a multiabstraction search within a multilevel optimization formulation [45]. Another limitation of this approach is that it does not account for spatial limitations. For the case studies, we designed the work-cell layouts integrated with the robot by adding a duplicate robot working area, removing the possibility of a spatial conflict. However, this solution will not necessarily work for all tasks, or it may be impractical to include the additional space, parts, and supplies. Adding spatial constraints to the optimizer and location information for each object may be a potential solution in such scenarios. Furthermore, we do not consider any costs for using the robot in the optimization objective function. Cost factors, such as wear and tear and energy consumption, would be important to consider for future research. In addition, our approach relies on estimations of SI from videos and thus needs validation through a systematic analysis of actual strain placed on workers and of injury reduction over time. A final limitation is that reducing makespan in a single station is not sufficient to improve the productivity of multistage production. Future extension of the proposed approach to task allocation and scheduling at a work-cell level is necessary to realize these benefits of the proposed approach beyond small-scale, single-station production.

VIII. Conclusion

In this paper, we propose and demonstrate an approach to consider time and ergonomics simultaneously for integrating a collaborative robot to a manufacturing process. We present a variation of the SI as a measure of human physical stress of each work element in a task and use its values in conjunction with task times to generate a task schedule. We discuss tradeoffs in assigning different priorities to time and ergonomics, including completion time, total strain, and worker idle time. The resulting set of schedules enable process engineers to balance these factors in a way that is the best suited to the task at hand.

Acknowledgment

The authors would like to thank C. Bodden, C.-M. Huang, J. Gagliardi, F. Golden, M. Gombolay, E. Jesse, J. Kim, J. Perez, A. Sauppé, and S. Wright, and Steelcase, Inc., for their contributions to this paper.

References

[1] (2015) *KUKA LBR iiwa*. [Online]. Available: http://www.kuka-robotics.com/en/products/industrial_robots/sensitiv/

[2] (2015). *UR3 Robots*. [Online]. Available: http://www.universal-robots.com/products/ur3-robot/

[3] (2015). *Rethink Robotics*. [Online]. Available: http://www.rethinkrobotics.com/

[4] A. Sauppe, "Designing communication strategies for human-robot collaboration," Ph.D. dissertation, Dept. Comput. Sci., Univ. Wisconsin-Madison, Madison, WI, USA, 2015.

[5] M. Koes, I. Nourbakhsh, and K. Sycara, "Heterogeneous multirobot coordination with spatial and temporal constraints," in *Proc. AAAI*, vol. 5. 2005, pp. 1292–1297.

[6] S. Ponda, H.-L. Choi, and J. How, "Predictive planning for heterogeneous human-robot teams," in *Proc. AIAA Infotech Aerosp.*, 2010, pp. 1–10.

[7] F. Chen, K. Sekiyama, F. Cannella, and T. Fukuda, "Optimal subtask allocation for human and robot collaboration within hybrid assembly system," *IEEE Trans. Autom. Sci. Eng.*, vol. 11, no. 4, pp. 1065–1075, Oct. 2014.

[8] M. C. Gombolay, R. Wilcox, and J. A. Shah, "Fast scheduling of multirobot teams with temporospatial constraints," in *Robotics: Science and Systems*. Berlin, Germany, 2013.

[9] G. A. Korsah, "Exploring bounded optimal coordination for heterogeneous teams with cross-schedule dependencies," Ph.D. dissertation, Robot. Inst., Carnegie Mellon Univ., Pittsburgh, PA, USA, 2011.

[10] J. Krüger, R. Bernhardt, D. Surdilovic, and G. Spur, "Intelligent assist systems for flexible assembly," *CIRP Ann.-Manuf. Technol.*, vol. 55, no. 1, pp. 29–32, 2006.

[11] O. Ogorodnikova, "Human weaknesses and strengths in collaboration with robots," *Periodica Polytech. Eng. Mech. Eng.*, vol. 52, no. 1, pp. 25–33, 2008.

[12] G. Michalos, S. Makris, J. Spiliotopoulos, I. Misios, P. Tsarouchi, and G. Chryssolouris, "ROBO-PARTNER: Seamless human-robot cooperation for intelligent, flexible and safe operations in the assembly factories of the future," *Procedia CIRP*, vol. 23, pp. 71–76, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212827114011366

[13] G. Michalos, S. Makris, P. Tsarouchi, T. Guasch, D. Kontovrakis, and G. Chryssolouris, "Design considerations for safe human-robot collaborative workplaces," *Procedia CIRP*, vol. 37, pp. 248–253, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212827115008550

[14] P. Tsarouchi, S. Makris, and G. Chryssolouris, "On a human and dual-arm robot task planning method," *Procedia CIRP*, vol. 57, pp. 551–555, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2212827116312495

[15] M. Faber, S. Kuz, A. Mertens, and C. M. Schlick, "Model-based evaluation of cooperative assembly processes in human-robot collaboration," in *Advances in Ergonomics of Manufacturing: Managing the Enterprise of the Future*. Cham, Switzerland: Springer, 2016, pp. 101–112.

[16] J. Barondess *et al.*, "Musculoskeletal disorders and the workplace: Low back and upper extremities," Washington, DC, USA: National Academy Sciences, 2001, pp. 1–512.

[17] R. G. Radwin, M. L. Lin, and T. Y. Yen, "Rapid communication. exposure assessment of biomechanical stress in repetitive manual work using frequency-weighted filters," *Ergonomics*, vol. 37, no. 12, pp. 1984–1998, 1994.

[18] J. S. Moore and A. Garg, "The strain index: A proposed method to analyze jobs for risk of distal upper extremity disorders," *Amer. Ind. Hygiene Assoc. J.*, vol. 56, no. 5, pp. 443–458, 1995.

[19] J. S. Moore, N. P. Rucker, and K. Knox, "Validity of generic risk factors and the strain index for predicting nontraumatic distal upper extremity morbidity," *Amer. Ind. Hygiene Assoc.*, vol. 62, no. 2, pp. 229–235, 2001.

[20] K. Knox and J. S. Moore, "Predictive validity of the Strain Index in turkey processing," *J. Occupat. Environ. Med.*, vol. 43, no. 5, pp. 451–462, 2001.

[21] N. Rucker and J. S. Moore, "Predictive validity of the strain index in manufacturing facilities," *Appl. Occupational Environ. Hygiene*, vol. 17, no. 1, pp. 63–73, 2002.

[22] A. Garg *et al.*, "The strain index (SI) and threshold limit value (TLV) for hand activity level (HAL): Risk of carpal tunnelsyndrome (CTS) in a prospective cohort," *Ergonomics*, vol. 55, no. 4, pp. 396–414, 2012.

[23] J. M. Kapellusch, A. Garg, K. T. Hegmann, M. S. Thiese, and E. J. Malloy, "The strain index and ACGIH TLV for HAL: Risk of trigger digit in the WISTAH prospective cohort," *Hum. Factors, J. Hum. Factors Ergonom. Soc.*, vol. 56, no. 1, pp. 98–111, 2014.

[24] A. Garg *et al.*, "The strain index and TLV for HAL: Risk of lateral epicondylitis in a prospective cohort," *Amer. J. Ind. Med.*, vol. 57, no. 3, pp. 286–302, 2014.

[25] E. M. Stevens, G. A. Vos, J.-P. Stephens, and J. S. Moore, "Inter-rater reliability of the strain index," *J. Occupat. Environ. Hygiene*, vol. 1, no. 11, pp. 745–751, 2004.

[26] J.-P. Stephens, G. A. Vos, E. M. Stevens, Jr., and J. S. Moore, "Test–retest repeatability of the Strain Index," *Appl. Ergonom.*, vol. 37, no. 3, pp. 275–281, 2006.

[27] P. Brucker, *Scheduling Algorithms*, vol. 3. Berlin, Germany: Springer, 2007.

[28] J. Annett, "Hierarchical task analysis," in *Handbook of Cognitive Task Design*, vol. 2. Mahwah, NJ, USA: Lawrence Erlbaum Associates, 2003, pp. 17–35.

[29] A. Shepherd, *Hierarchical Task Analysis*. Boca Raton, FL, USA: CRC Press, 2000.

[30] T. J. Armstrong, R. G. Radwin, D. J. Hansen, and K. W. Kennedy, "Repetitive trauma disorders: Job evaluation and design," *Hum. Factors*, vol. 28, no. 3, pp. 325–336, 1986.

[31] J. T. C. Tan, F. Duan, R. Kato, and T. Arai, *Collaboration Planning by Task Analysis in Human-Robot Collaborative Manufacturing System*. Rijeka, Croatia: INTECH, 2010.

[32] N. A. Stanton, "Hierarchical task analysis: Developments, applications, and extensions," *Appl. Ergonom.*, vol. 37, no. 1, pp. 55–79, 2006.

[33] IBM. (2015). *CPLEX Optimizer*. [Online]. Available: http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/

[34] B. Kirwan and L. K. Ainsworth, *A Guide to Task Analysis: The Task Analysis Working Group*. Boca Raton, FL, USA: CRC Press, 1992.

[35] (2015). *Steelcase*. [Online]. Available: http://steelcase.com

[36] (2016). *The National Institute for Occupational Safety and Health (NIOSH)*. [Online]. Available: https://www.cdc.gov/niosh/

[37] (2015) *Baxter Tech Specs*. [Online]. Available: http://www.rethinkrobotics.com/baxter/tech-specs/

[38] I. A. Sucan and S. Chitta. (2015). *Moveit!* [Online]. Available: http://moveit.ros.org

[39] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, vol. 2, Apr. 2000, pp. 995–1001.

[40] S. Nikolaidis and J. Shah, "Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy," in *Proc. 8th ACM/IEEE Int. Conf. Human-Robot Interaction*, Mar. 2013, pp. 33–40.

[41] C.-M. Huang, M. Cakmak, and B. Mutlu, "Adaptive coordination strategies for human-robot handovers," in *Robotics: Science and Systems*. Rome, Italy, 2015.

[42] (2016). *Universal Robots UR10*. [Online]. Available: https://www.universal-robots.com/products/ur10-robot/

[43] V. Jain and I. E. Grossmann, "Algorithms for hybrid MILP/CP models for a class of optimization problems," *INFORMS J. Comput.*, vol. 13, no. 4, pp. 258–276, 2001.

[44] E. Castro and S. Petrovic, "Combined mathematical programming and heuristics for a radiotherapy pre-treatment scheduling problem," *J. Scheduling*, vol. 15, no. 3, pp. 333–346, 2012.

[45] C. Zhang and J. A. Shah, "Co-optimizating multi-agent placement with task assignment and scheduling," in *Proc. IJCAI*, 2016, pp. 3308–3314.

Margaret Pearce received the B.S. degree in computer science and computational mathematics from the University of Michigan at Dearborn, Dearborn, MI, USA, in 2011, and the M.S. degree in computer science from the University of Wisconsin–Madison, Madison, WI, USA, in 2016.

Her research interests include human-robot teaming and optimization.

Bilge Mutlu received the Ph.D. degree from the Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, USA, in 2009.

He is currently an Associate Professor of Computer Science with the University of Wisconsin–Madison, Madison, WI, USA, where he is also the Director of the Wisconsin Human-Computer Interaction Laboratory. His research program bridges human–computer interaction and robotics, focusing on the development of human-centered methods and principles for designing robotic technologies.

Dr. Mutlu has served as an elected member of the Human-Robot Interaction (HRI) Steering Committee. He was a recipient of the Fulbright Scholarship, the National Science Foundation's CAREER Award, and several paper awards and nominations. He has served as the Program Co-Chair of the International Conference on Social Robotics, HRI, and IEEE International Conference on Robot and Human Interactive Communication Conferences, in 2012, 2015, and 2016, respectively, and the General Co-Chair of the HRI Conference in 2017. His research has been covered by national and international press, including the New Scientist, MIT Technology Review, Discovery News, Science Nation, and Voice of America.

Julie Shah received the S.B. and S.M. degrees in aeronautics and astronautics and the Ph.D. degree in autonomous systems from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2004, 2006, and 2010, respectively.

She has developed innovative methods for enabling fluid human-robot teamwork in time-critical and safety-critical domains, ranging from manufacturing to surgery to space exploration. Her group draws on expertise in artificial intelligence, human factors, and systems engineering to develop interactive robots that emulate the qualities of effective human team members to improve the efficiency of human-robot teamwork. She is currently an Associate Professor with the Department of Aeronautics and Astronautics, MIT, where she leads the Interactive Robotics Group, Computer Science and Artificial Intelligence Laboratory.

Dr. Shah was recognized with the NSF CAREER Award and by the MIT Technology Review TR35 list as one of the world's top innovators under the age of 35.

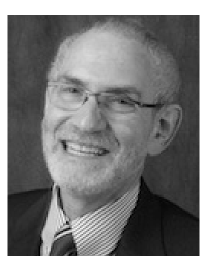

Robert Radwin received the B.S. degree from the New York University Polytechnic School of Engineering, Brooklyn, NY, USA, in 1975, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, MI, USA, in 1979 and 1986, respectively.

He is currently a Professor of Biomedical Engineering, Industrial and Systems Engineering, and Orthopedics and Rehabilitation with the University of Wisconsin–Madison, Madison, WI, USA.