

Longitudinal Control Volumes: A Novel Centralized Estimation and Control Framework for Distributed Multi-Agent Sorting Systems

James Maier¹, Prasanna Sriganesh¹, and Matthew Travers¹

Abstract—Centralized control of a multi-agent system improves upon distributed control especially when multiple agents share a common task e.g., sorting different materials in a recycling facility. Traditionally, each agent in a sorting facility is tuned individually which leads to suboptimal performance if one agent is less efficient than the others. Centralized control overcomes this bottleneck by leveraging global system state information, but it can be computationally expensive. In this work, we propose a novel framework called Longitudinal Control Volumes (LCV) to model the flow of material in a recycling facility. We then employ a Kalman Filter that incorporates local measurements of materials into a global estimation of the material flow in the system. We utilize a model predictive control algorithm that optimizes the rate of material flow using the global state estimate in real-time. We show that our proposed framework outperforms distributed control methods by 40-100% in simulation and physical experiments.

Index Terms—Recycling, Process Control, Computer Vision, State Estimation, Kalman Filters, Receding Horizon Model Predictive Control

I. INTRODUCTION

Centralized control approaches improve task performance over distributed approaches by leveraging information about the system state to coordinate agent actions. In practice, however, difficulties with computational complexity when solving the centralized coordination problem can mean that system designers follow distributed control approaches. One compelling example of a problem domain where practitioners often opt for distributed over centralized control approaches is found in modern facilities for sorting recycled material.

Recycling sorting facilities generally are comprised of a series of material sorting agent stations coupled by conveyor belt network. Fig. 1 shows a simple example of a material sorting operation. Sorting system operators improve the performance of individual sorting agents by using computer vision systems to detect, classify, and track individual waste items. However, due to the large number of waste items in a sorting system (measured in millions of individual items), these approaches become computationally intractable when they are applied to optimize the system's operation as a whole.

The control of the flow rate of material through the system is a key aspect of centralized multi-agent sorting system control. Since the rate at which material is presented to a sorting agent directly affects the sorting agent's ability to remove target material from the stream, a centralized authority can affect the efficacy of sorting agents throughout



Fig. 1: Materials on a conveyor being sorted by a human

the system by varying flow rates. Dynamically varying the flow rate throughout the system presents a problem for distributed control systems because the system throughput will be limited to that of the slowest agent. In contrast, a central authority can avoid bottlenecks by dynamically varying material flow rates in response to varying infeed conditions.

In this work, we present a novel modeling framework that divides the entire conveyor into smaller Longitudinal Control Volumes (LCV) which enables a low-dimensional representation and prediction of material flows throughout the sorting facility. We present a method for converting detections of individual waste items from images into sparse measurements in our state space. We use these measurement and prediction framework with a Kalman Filter to compute a global estimate of the system state.

This modeling framework is used in conjunction with the global estimation result in a model predictive control (MPC) module that optimizes the material flow to maximize the overall value generated for successful sorting. The controller reacts to the incoming material over a receding time horizon and changes the speed of the conveyor to allow for efficient sorting of materials.

We demonstrate the efficacy of our centralized estimation and control approach in simulation and on a physical laboratory sorting system. We show improvements on the order of 40-100% with respect to constant belt speed approaches. We later discuss applications of our framework to diverse multi-agent systems including autonomous mobile robots and over-actuated control systems.

II. RELATED WORK

In order to estimate performance of the core sorting system in a material recovery facility (MRF), the authors in [1][2][3] develop tools to estimate sorting process efficiency based on

¹James Maier, Prasanna Sriganesh and Matthew Travers are from The Robotics Institute, Carnegie Mellon University, USA. {jamesmai, pkettava, mtravers}@andrew.cmu.edu




Fig. 2: An illustration of multi-agent distributed sorting system with dynamics coupled by a single conveyor belt

historical data on material flow compositions. These tools leverage long-term historical data and are used to calculate steady-state optimal material flow rates for a given system. In this system-level offline modeling approach, a network flow diagram for the MRF is created by approximating the dynamics of mixed material flows through sorting machines and connecting them via a directed graph. A system of linear equations is then constructed and solved to calculate material stream compositions throughout the system at steady state [4][2].

The work presented in [3][5][6] expands on the foundational work by [4] and [2] by estimating the percentage of material in the main stream that a single sort agent can remove. They then set up a centralized optimization problem to find nominal steady-state optimal material flow rates. These offline modeling and optimization techniques represent a method for achieving global system coordination. However, they do not capture dynamics associated with material flow characteristics changing in time. They also rely on assumption of static sorting parameters which are often not representative of real sorting operations [7][8].

A second line of research involves distributed control techniques which can run in real-time. In these distributed control approaches, individual sorting machines leverage local system information to improve their performance. A literature review of these approaches for distributed optimization of sorting systems was conducted by [9].

Of particular relevance to the content of this work, [10][11][12] leverage techniques from computer vision and machine learning along with Kalman Filtering and multi-target tracking. They track each object entering a sorting machine and optimize that sorting machine's operation around real-time flow characteristics. These approaches show promising improvements to the operations of individual sorting machines, but they have not been applied to solving the global coordination problem. We hypothesize that the computational complexity of simultaneously tracking millions of individual pieces of recycled material presents an obstacle to use in centralized control methods.

Furthermore, as noted in [13] and [1], real MRF operations are much less efficient than would be expected. The authors of [14] investigate the discrepancy between theoretically

expected sorting operations and real MRF sorting data and conclude that the high degree of variability in the material stream means that solutions calculated with offline centralized optimization systems do not translate well to real operation. The authors in [15] study the variability in the input material stream and provide statistics on flow variability from real MRF operations, and they establish that online control techniques are necessary to cope for the variations in the input material flows.

In this work, we draw particular inspiration from [14] and [15] who establish that MRF sorting performance suffers in the face of variable input material flows. We develop a state modeling approach which follows ideologically from the systems described in [2] and [3], but which provides a measure of system state with enough fidelity for real-time dynamic control. We use this state estimation result as a basis for real-time centralized control of material flows throughout the multi-agent sorting system.

III. MODELING MATERIAL FLOW COMPOSITIONS

This section introduces a new framework to represent the current state of a sorting system, and describes a model to track material composition flow as it moves across the conveyor.

A. Longitudinal Control Volumes (LCV)

The state representation of a sorting system needs to provide information about the composition of the material flow as it evolves along the conveyor. To compute this, we first longitudinally split the sorting system into small control volumes arranged sequentially along the direction of travel of the belt (Fig. 3). The size of these control volumes is tuned based on the spatial resolution of the perception system and the computational load of having a high number of control volumes. The amount of material in the control volume forms the basis for constructing the state vector.

We construct the state vector by first concatenating smaller vectors \mathbf{x}_k^i , for every material i . Each vector element in \mathbf{x}_k^i corresponds to a control volume and indicates the quantity of a material occupying that control volume at timestep k . Lastly, the current speed of the conveyor belt, r_k is added to the end of the concatenated vector which forms the full



Fig. 3: Control volumes overlaid on conveyor belt serving a team of human and robot sorters

system state. The total amount of a given material in the entire system can be computed by summing the elements of the state vector corresponding to that material. Given a total of m control volumes and n different material types to be sorted, the state vector is of the size $nm + 1$ and can be written as follows

$$\mathbf{X}_k = [\mathbf{x}_k^1, \mathbf{x}_k^2, \dots, \mathbf{x}_k^n, r_k]^T \quad (1)$$

$$\text{Total amount of material } i \text{ in the system} = \sum \mathbf{x}_k^i$$

B. State Transition Model

The motion of the conveyor belt and the sorting process change the composition of material in the control volumes. To predict this change in the system state, we decompose the state transition matrix to model both the dynamics associated with the conveyor motion ($\mathcal{L}(r_k)$) and the dynamics associated with the sorting process ($\mathcal{F}(\mathbf{X}_k)$) to model a linear approximation of these processes. The control input to the system is the change in conveyor speed Δr . This makes the matrix B a $mn + 1$ long zero vector with the last element equal to 1 as the control input affects only the last element of the state. This state transition model is written as follows,

$$\begin{aligned} \mathbf{X}_{k+1} &= A(\mathbf{X}_k)\mathbf{X}_k + Bu_k \\ &= \mathcal{F}(\mathbf{X}_k)\mathcal{L}(r_k)\mathbf{X}_k + B\Delta r \end{aligned} \quad (2)$$

The motion matrix $\mathcal{L}(r_k)$ models the motion of material across different control volumes in the system carried by conveyor belts. This is dependent on the current speed of the conveyor belt. To construct this matrix for any arbitrary belt speed, we compute these motion matrices for integral belt speeds and perform linear interpolation between these integral speed motion matrices. Given an integer belt speed \hat{r}_k , we can construct a motion matrix $\mathcal{L}(\hat{r}_k)$ as follows,

$$\mathcal{L}(\hat{r}_k) = \begin{bmatrix} L(\hat{r}_k) & 0_{m \times m} & \dots & 0_{m \times m} & 0 \\ 0_{m \times m} & L(\hat{r}_k) & \dots & 0_{m \times m} & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0_{m \times m} & 0_{m \times m} & \dots & L(\hat{r}_k) & 0 \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}$$

where $L(\hat{r}_k)$ is a $m \times m$ modified shift matrix given by

$$L(\hat{r}_k)_{ij} = \delta_{i, j + \hat{r}_k}$$

where $\delta_{i,j}$ is the Kronecker delta

This motion matrix can only be defined for an integral belt speed due to the Kronecker delta. Hence we do a linear interpolation of the lower and upper integral motion matrices to compute the motion matrix for an arbitrary belt speed. Given a belt speed r_k , we can compute the motion matrix $\mathcal{L}(r_k)$ as follows,

$$\mathcal{L}(r_k) = \frac{\lceil r_k \rceil - r_k}{\lceil r_k \rceil - \lfloor r_k \rfloor} \mathcal{L}(\lfloor r_k \rfloor) + \frac{r_k - \lfloor r_k \rfloor}{\lceil r_k \rceil - \lfloor r_k \rfloor} \mathcal{L}(\lceil r_k \rceil)$$

where $\lceil x \rceil = \text{ceil}(x)$ and $\lfloor x \rfloor = \text{floor}(x)$

We model the dynamics associated with removing material at sorting agents using the sort matrix $\mathcal{F}(\mathbf{X}_k)$. As each sorting agent only spans a few control volumes in the entire system, they can sort a limited amount of material in each time step. To capture this behavior, we define a separation parameter p_j^i for each sort agent [16][4]. This separation parameter is defined per material for every control volume based on the maximum number of items that can be picked up by the corresponding sort agent.

If α^i represents the subset of control volumes spanned by a sort agent and γ_i represents the maximum amount of material that can be picked up in a single timestep, we can compute the separation parameters based on ratio of maximum pick rate and amount of material present on the control volumes of a sort agent (η). We can also represent imperfections inherent in the sorting process by a discount factor q^i with values ranging from 0 (i.e. no successful sorting) to 1 (perfect sorting).

$$\eta^i = q^i \frac{\gamma^i}{\sum_{\alpha^i} \mathbf{x}^i} \quad p_j^i = \begin{cases} 1, & \text{if } j \notin \alpha^i \\ \max(0, 1 - \eta^i), & \text{otherwise} \end{cases}$$

The sorting $\mathcal{F}(\mathbf{X}_k)$ matrix consists of smaller diagonal matrices $F(\mathbf{x}_k^i)$ that capture the sorting process for every material in each control volume. This allows us to represent the dynamics associated with all sorting machines in the system in a single linear element. The matrix $F(\mathbf{x}_k^i)$ is a $m \times m$ diagonal matrix with the diagonal elements equal to the separation parameter p_j^i . We can write this as follows:

$$\mathcal{F}(\mathbf{X}_k) = \begin{bmatrix} F(\mathbf{x}_k^1) & 0_{m \times m} & \dots & 0_{m \times m} & 0 \\ 0_{m \times m} & F(\mathbf{x}_k^2) & \dots & 0_{m \times m} & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0_{m \times m} & 0_{m \times m} & \dots & F(\mathbf{x}_k^n) & 0 \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}$$

$$F(\mathbf{x}_k^i) = \begin{bmatrix} p_1^i & 0 & \dots & 0 \\ 0 & p_2^i & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & p_m^i \end{bmatrix}$$

IV. METHODOLOGY

We create a system to make an optimal tradeoff between sorting task completion and system throughput in the context of a recycling sorting operation. As the material flows through the conveyor, there is uncertainty in the state induced by the sorting process and material moving relative to the conveyor belt. To incorporate this uncertainty, we present

a framework to use a computer vision system to detect materials and estimate the system's state using an Kalman Filter based on the LCV state transition model. We then use this estimated state as input for a centralized model predictive control module, which maximizes the difference between the reward associated with successful sorting and the cost associated with failing to sort potentially valuable material. The controller outputs an optimized conveyor belt speed that maximizes the system's throughput while maintaining high rewards for sorting. A block diagram of this system is shown in Fig. 4.

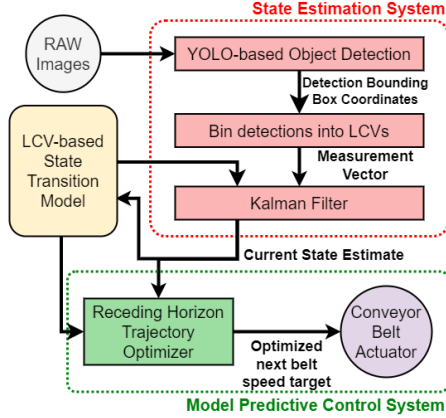


Fig. 4: A block diagram of the proposed system to optimize belt speeds for sorting

A. State Estimation

In order to incorporate real-time data into our state estimate, we devise a computer vision system to detect individual objects on a subset of control volumes. An Intel Realsense D435 RGB camera provides top-down images to a YOLO-based object detection system that classifies different materials and returns bounding boxes that describe the location and type of individual items [17]. The detection from every frame is incorporated as a measurement vector.

To compute the measurement vector, we project the visible control volumes into the camera image and add the counts of detected objects into the control volumes associated with their locations. If a detected object spans multiple control volumes, we add portions of the detected object to each control volume spanned by the object. If λ is the number of control volumes in the camera's view and W^i is the set of detections of material i , we can compute the measurement vector \mathbf{z}_k as follows,

$$\mathbf{z}_k = [\bar{\zeta}^0 \quad \bar{\zeta}^1 \quad \dots \quad \bar{\zeta}^n] \quad \bar{\zeta}^i = [\beta_0^i \quad \beta_1^i \quad \dots \quad \beta_\lambda^i]$$

$$\beta_j^i = \sum_{w \in W^i} \frac{\text{Area of bounding box for detection } w \text{ spanning control volume } j}{\text{Total area of bounding box for detection } w}$$

This process is shown in Fig. 5. The measurement vector is of the size $n\lambda$ as it comprises the control volumes visible to the camera. We also associate a measurement




Fig. 5: Control volume boundaries overlaid on image with bounding box detections from object detection system. The red and blue bars at the bottom help illustrate binning for different materials and the measurement vector generated from this input image.

covariance matrix to incorporate the noise inherent to the object detection system.

We adopt a Kalman Filter (KF) that can estimate the entire state of the system using the state transition model and the measurements coming from the computer vision system [18]. The uncertainty induced by the sorting process and movement of the conveyor belt is modeled as the process noise for the KF. This filter enables the prediction of the full-state with an associated covariance that represents the uncertainty in the predicted state.

B. Model Predictive Control

We setup a Model Predictive Control (MPC) problem which uses the state-transition model described in III-B and the estimated state to optimize sorting operations and control the conveyor belt speed in real-time. We choose to maximize an objective function that can maximize the value of sorted material which is obtained using a Value matrix V and an Opportunity Cost matrix O . The Value matrix maps the price of a material ρ_i onto the sort matrix ($F^i(\mathbf{x}_k)$). The Opportunity Cost matrix encodes the material that we fail to sort and are forced to sell for a lower price as mixed material. These matrices are computed as shown below.

$$V = [\bar{v}^0 \quad \bar{v}^1 \quad \dots \quad \bar{v}^n \quad 0] \quad \bar{v}^i = \rho^i F(\mathbf{x}_k^i) \mathbf{1}_m$$

$$O = [\bar{o}^0 \quad \bar{o}^1 \quad \dots \quad \bar{o}^n \quad 0] \quad \bar{o}^i = \rho^i (I - F(\mathbf{x}_k^i)) \mathbf{1}_m$$

where $\mathbf{1}_m$ is a 1-vector of size $m \times 1$

For the optimization, we set the horizon to the time it would take material to fully transit the system with all belts at their slowest speed. Eq. 3 describes our objective function that seeks to maximize the difference between the total sale price of material sorted over the horizon and the value lost by selling a material with contamination.

$$\max_{\mathbf{X}_{1:T}} \sum_{l=1}^T [V - O] \mathbf{X}_l \quad (3)$$

s.t.

$$\mathbf{X}_{l+1} = \mathcal{F}(\mathbf{X}_l) \mathcal{L}(r_l) \mathbf{X}_l + B u_l \quad \text{for } l = 1, \dots, T-1$$

$$u_{min} \leq u_l \leq u_{max}$$

We set up a quasi-Newton solver with a back-tracking line search to maximize our objective function. We compute the gradients by calculating finite-difference derivatives for each control variable, and use the gradient for to pick the update direction. The update step is then scaled using a BFGS approximation of the inverse Hessian matrix in the traditional Newton update equation [19][20][21][22]. The optimization is terminated based on fulfillment of the Armijo Condition[23]. The resulting optimized control input (conveyor belt speed) corresponding to the next timestep is then applied to the system. This process is repeated at every timestep.

V. EXPERIMENTS

We conduct experiments in simulation and on a physical conveyor belt sorting system and evaluate our system’s performance by comparing it to that of a baseline approach. For the baseline, we run the conveyor at a constant speed, similar to approaches currently used in industrial MRFs. This optimal constant speed is estimated based on the historical average material flow through the system. We compare the total value generated by this baseline to the value generated by running the system at the optimized variable belt speed generated by the MPC module.

A. Simulation-based evaluation of MPC performance

We create a simulated environment where infeed material flow is procedurally generated in a manner representative of actual material flows in an MRF according to material variability statistics described in [15]. The LCV-based state transition model described previously is used to forward-simulate the sorting operation. We generate 30 runs that each represent an hour of continuous operation. We take the average optimized belt speed produced by the MPC module for each run and repeat the experiment with a constant belt speed equal to that average optimized speed. We compare the final accumulated value at the end of each optimized run with the accumulated value from its corresponding constant speed run. We repeat this process for sorting systems with 2, 3, and 4 distinct materials to be sorted.

B. Full-system Experiments on Physical Sorting Equipment

We also conducted experiments on a physical conveyor belt sorting system to evaluate the performance of our end-to-end system. Different material streams can be fed to the conveyor belt using a hopper on one end. A computer with a ROS framework captured overhead images from the Intel RealSense D435 camera, executed the state estimation and MPC modules in real-time, and controlled the speed of the conveyor belt using a servo motor. At the end of this conveyor, a human sorter was tasked with removing one material from the flow. This setup is shown in the Fig. 6.

In these physical experiments, we use two materials: empty green bottles and empty red cans. In order to investigate the ability of our framework to react to changing infeed conditions during these short runs, we loaded the hopper with varying distributions of the material flows as the MPC module optimized the conveyor belt speed in real-time. For each of these runs, we computed the value generated for



Fig. 6: Physical conveyor belt experimental setup showing the RealSense D435 overhead camera, human sorter and the sorted materials

sorting and compared this to the value generated using a constant speed baseline experiment. The constant speed was equal to the average of the optimized speeds generated by the MPC module.

VI. RESULTS

The results of our proposed framework for a three-material sorting simulation is shown in the Fig. 7. We compare the total profit generated by our method to that of the constant speed baseline model. Fig. 7(a) shows the instantaneous distribution of the materials in their respective sort stations, along with their price. Fig. 7(b) shows shows the target conveyor belt speed by the MPC module (cyan line) and the target belt speed by the constant speed baseline (magenta line). Fig. 7(c) shows the total profits accumulated for successfully sorted materials for both versions.

From the plots shown in the figures 7(a) and 7(b), it is evident that the MPC module is reacting to the amount of material coming into the system. As more Material 0 (which has a higher price) enters the system, the controller slows down the conveyor to allow the sorters to pick up more material, increasing the profits accumulated. For the same case, since the constant speed model does not change the belt speeds, the plant incurs a loss as the sorters fail to sort this high-value material, and the impure mixture is worth much less. We see the same behavior even in the 2-material runs as shown in the Fig. 8.

We notice similar trends across all the different runs for varying amounts of materials in the system. When different materials are distributed evenly, the constant speed baseline and our proposed framework perform similarly. In contrast, the profit rate of our method goes up quickly when the distribution varies quickly. This was true across all the 90 runs we conducted in simulation. Fig. 9(a) shows the average profit rate along with the variance in profit rate for the proposed method and the constant speed baseline for varying number of simulated material runs. In the 2-material case, our optimized speeds improved the margin by around 40 percent while we see around two times the profit in the 3-material and 4-material cases. In the physical experiments with 2-materials, the optimized speed run improves the profit margins by two times as shown in Fig. 9(b).

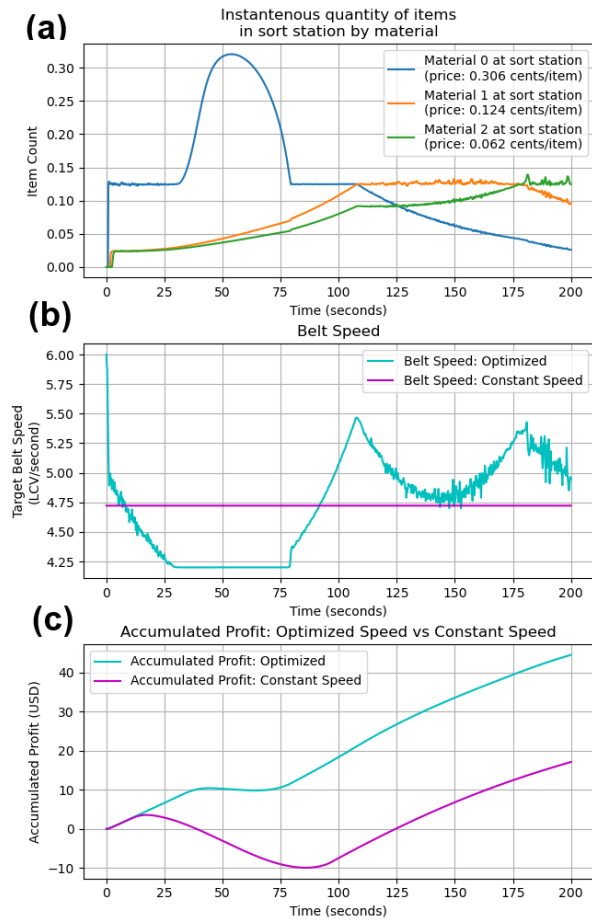


Fig. 7: Result of a simulated 3-material sorting run. (a) Plot showing the instantaneous amount of material passing through a sort agent station (b) Plot showing the output belt speeds for the MPC and the baseline. (c) Plot showing the accumulated value as a result of successfully sorting.

VII. CONCLUSION AND FUTURE WORK

In the paper, we presented a novel framework that modeled the movement of material along a conveyor belt in a sorting system. We presented an end-to-end system that successfully detected and estimated these material flows while optimizing the conveyor belt speed for maximizing value in the sorting process. We conducted experiments both in simulation and on a physical conveyor belt with varying numbers of materials and compositions. We achieved around 40 to 100% increase in the profit rates using the proposed model predictive controller that accounted for the entire state of the system.

Although the presented work shows dramatic improvements to process profitability, we identify several potential avenues for further research. As the MPC module does not leverage the uncertainty information from the state estimation module, we intend to explore incorporating the uncertainty in the optimization process along with trying to model an objective function derived from operations research. We would like to investigate learning-based approaches to directly infer measurement vectors from camera images. Further, we are in the process of scaling our experiments at an operational

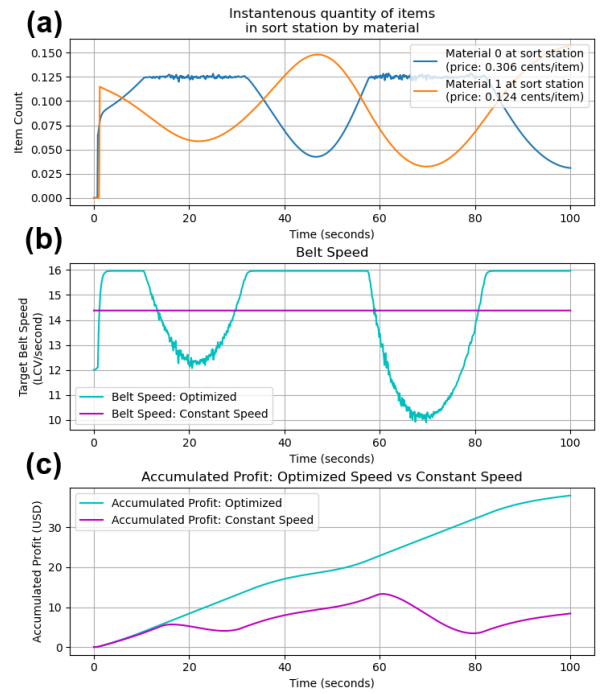
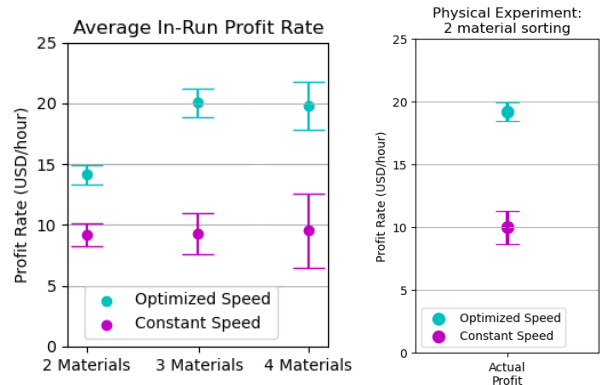


Fig. 8: Result of a 2-material sorting run. (a) Plot showing the instantaneous amount of material passing through a sort agent (b) Plot showing the output belt speeds for the MPC and the baseline. (c) Plot showing the accumulated value as a result of successfully sorting.



(a) Simulated material flow

(b) Physical experiment with 2 materials

Fig. 9: Result plot showing profit rate for optimized speed vs constant speed control

MRF sorting facility. Future work also includes integrating this system with a fully mechanised sorting line in order to eliminate any inconsistency associated with using a human sorting worker.

Beyond applications in conveyor based sorting, we firmly believe that use of control volumes is more broadly applicable to systems in settings that have a team of agents team collaborating on a task. Specifically, we intend to apply this approach to problems associated with a team of autonomous mobile robots exploring an unknown area [24][25][26]. We also intend to investigate this approach in the control of micro-robot swarms for medical applications [27][28].

REFERENCES

- [1] F. Tanguay-Rioux, R. Legros, and L. Spreutels, "On the limits of empirical partition coefficients for modeling material recovery facility unit operations in municipal solid waste management," *Journal of Cleaner Production*, vol. 293, p. 126016, 2021.
- [2] M. I. Wolf, "Modeling and design of material separation systems with applications to recycling," Ph.D. dissertation, Massachusetts Institute of Technology, 2011.
- [3] M. Testa, "Modeling and design of material recovery facilities: Genetic algorithm approach," Ph.D. dissertation, Massachusetts Institute of Technology, 2015.
- [4] T. Gutowski and J. Dahmus, "Product recycling systems," in *2005 NSF Design Service and Manufacturing Grantees and Research Conference*, 2005.
- [5] K. Ip, M. Testa, A. Raymond, S. C. Graves, and T. Gutowski, "Performance evaluation of material separation in a material recovery facility using a network flow model," *Resources, Conservation and Recycling*, vol. 131, pp. 192–205, 2018.
- [6] K. Kleinhans, M. Hallemans, S. Huysveld, *et al.*, "Development and application of a predictive modelling approach for household packaging waste flows in sorting facilities," *Waste Management*, vol. 120, pp. 290–302, 2021.
- [7] A. C. Caputo and P. M. Pelagagge, "Rdf production plants: I design and costs," *Applied Thermal Engineering*, vol. 22, no. 4, pp. 423–437, 2002.
- [8] C. A. Velis, S. Wagland, P. Longhurst, *et al.*, "Solid recovered fuel: Materials flow analysis and fuel property development during the mechanical processing of biodried waste," *Environmental Science & Technology*, vol. 47, no. 6, pp. 2957–2965, 2013.
- [9] N. Kroell, X. Chen, K. Greiff, and A. Feil, "Optical sensors and machine learning algorithms in sensor-based material flow characterization for mechanical recycling processes: A systematic literature review," *Waste Management*, vol. 149, pp. 259–290, 2022.
- [10] F. Pfaff, G. Kurz, C. Pieper, *et al.*, "Improving multitarget tracking using orientation estimates for sorting bulk materials," in *2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, IEEE, 2017, pp. 553–558.
- [11] G. Maier, M. Reith-Braun, A. Bauer, *et al.*, "Machine learning-based multiobject tracking for sensor-based sorting," in *FORUM BILDVERARBEITUNG 2022 IMAGE PROCESSING FORUM 2022*, 2022, p. 115.
- [12] F. Pfaff, *Multitarget Tracking Using Orientation Estimation for Optical Belt Sorting*. KIT Scientific Publishing, 2019, vol. 22.
- [13] A. G. Raymond, "Modeling of material recovery facility performance with applications for life cycle assessment," Ph.D. dissertation, Massachusetts Institute of Technology, 2017.
- [14] F. Tanguay-Rioux, L. Spreutels, M. Héroux, and R. Legros, "Mixed modeling approach for mechanical sorting processes based on physical properties of municipal solid waste," *Waste Management*, vol. 144, pp. 533–542, 2022.
- [15] A. Curtis, B. Küppers, S. Möllnitz, K. Khodier, and R. Sarc, "Real time material flow monitoring in mechanical waste processing and the relevance of fluctuations," *Waste management*, vol. 120, pp. 687–697, 2021.
- [16] F. Tanguay-Rioux, M. Héroux, and R. Legros, "Physical properties of recyclable materials and implications for resource recovery," *Waste management*, vol. 136, pp. 195–203, 2021.
- [17] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," *CoRR*, vol. abs/1506.02640, 2015. arXiv: 1506.02640. [Online]. Available: <http://arxiv.org/abs/1506.02640>.
- [18] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [19] C. G. BROYDEN, "The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations," *IMA Journal of Applied Mathematics*, vol. 6, no. 1, pp. 76–90, Mar. 1970, ISSN: 0272-4960. DOI: 10.1093/imamat/6.1.76. eprint: <https://academic.oup.com/imamat/article-pdf/6/1/76/2233756/6-1-76.pdf>. [Online]. Available: <https://doi.org/10.1093/imamat/6.1.76>.
- [20] R. Fletcher, "A new approach to variable metric algorithms," *The Computer Journal*, vol. 13, no. 3, pp. 317–322, Jan. 1970, ISSN: 0010-4620. DOI: 10.1093/comjnl/13.3.317. eprint: <https://academic.oup.com/comjnl/article-pdf/13/3/317/988678/130317.pdf>. [Online]. Available: <https://doi.org/10.1093/comjnl/13.3.317>.
- [21] D. Goldfarb, "A family of variable-metric methods derived by variational means," *Mathematics of Computation*, vol. 24, no. 109, pp. 23–26, 1970, ISSN: 00255718, 10886842. [Online]. Available: <http://www.jstor.org/stable/2004873> (visited on 09/16/2023).
- [22] D. F. Shanno, "Conditioning of quasi-newton methods for function minimization," *Mathematics of Computation*, vol. 24, no. 111, pp. 647–656, 1970, ISSN: 00255718, 10886842. [Online]. Available: <http://www.jstor.org/stable/2004840> (visited on 09/16/2023).
- [23] L. Armijo, "Minimization of functions having Lipschitz continuous first partial derivatives.," *Pacific Journal of Mathematics*, vol. 16, no. 1, pp. 1–3, 1966.
- [24] N. Bagree, C. Noren, D. Singh, M. Travers, and B. Vundurthy, "Distributed optimal control framework for

high-speed convoys: Theory and hardware results,” *arXiv preprint arXiv:2211.06287*, 2022.

- [25] P. Sriganesh, N. Bagree, B. Vundurthy, and M. Travers, “Fast staircase detection and estimation using 3d point clouds with multi-detection merging for heterogeneous robots,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 9253–9259.
- [26] S. Scherer, V. Agrawal, G. Best, *et al.*, “Resilient and modular subterranean exploration with a team of roving and flying robots,” *Field Robotics*, 2022.
- [27] E. Benjaminson, M. Travers, and R. Taylor, “Steering magnetic robots in two axes with one pair of maxwell coils,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2020, pp. 6332–6338.
- [28] E. Benjaminson, T. Imamura, A. Lorenz, S. Bergbreiter, M. J. Travers, and R. E. Taylor, “Buoyant magnetic milliswimmers reveal design rules for optimizing microswimmer performance,” *Nanoscale*, 2023.