A Primal-Dual Solver for Large-Scale Tracking-by-Assignment

Stefan Haller[1], **Mangal Prakash**[2], **Lisa Hutschenreiter**[1], **Tobias Pietzsch**[2],
Carsten Rother[1], **Florian Jug**[2], **Paul Swoboda**[3], **Bogdan Savchynskyy**[1]
[1]Visual Learning Lab, Heidelberg University,　　[2]Center for Systems Biology, Dresden,
[3]Max-Planck-Institute for Informatics, Saarbrücken

Abstract

We propose a fast approximate solver for the combinatorial problem known as *tracking-by-assignment*, which we apply to cell tracking. The latter plays a key role in discovery in many life sciences, especially in cell and developmental biology. So far, in the most general setting this problem was addressed by off-the-shelf solvers like Gurobi, whose run time and memory requirements rapidly grow with the size of the input. In contrast, for our method this growth is nearly linear.

Our contribution consists of a new (1) decomposable compact representation of the problem; (2) dual block-coordinate ascent method for optimizing the decomposition-based dual; and (3) primal heuristics that reconstructs a feasible integer solution based on the dual information. Compared to solving the problem with Gurobi, we observe an up to 60 times speed-up, while reducing the memory footprint significantly. We demonstrate the efficacy of our method on real-world tracking problems.

1 Introduction

The tracking problem consists of segmenting images obtained over T time steps and matching the segments in consecutive images to each other. In the case of cell tracking each cell in the image t must be matched to the corresponding cell (or a pair of cells in case of a division) in the image $t + 1$. Tracking problems are important not only for bioimaging (e.g. Kamentsky et al., 2011; Chenouard et al., 2014; Ulman et al., 2017) but also for general computer vision (e.g. Kristan et al., 2017). Cell tracking represents one of the hardest types

of this problem, due to the existence of cell divisions and the indistinguishability of individual cells.

The most successfully deployed tracking models are typically formulated as integer linear programs (ILPs) and are in general NP-hard. This makes the commonly used ILP-based optimizations only amenable to moderately sized tracking instances. With the advent of modern microscopy techniques, this bottleneck became a limiting factor for many real-world applications, leading either to faulty tracking results, or impractical optimization problems. In order to address run time and scalability issues, we propose a new method to solve tracking-by-assignment problems, whose iterations as well as memory footprint scale nearly linearly with the problem size. Its convergence speed enables us to obtain high-quality approximate solutions in a fraction of the time required by the best off-the-shelf solvers.

While this paper focuses on biologically motivated object tracking, similar problems also arise in other vision domains (see e.g. Wang et al., 2014; Tang et al., 2015; Iqbal et al., 2017). We believe that with appropriate modifications our ideas can be applied there as well.

Related work. Any visual tracking contains two key interrelated operations: segmentation and matching. On one side, matching requires segmentation, on the other side, the segmentation quality can often be significantly improved by the matching results. There are approaches addressing these two problems jointly (Jug et al., 2016), however, the resulting algorithms are quite time consuming. Less expensive modeling techniques can be categorized into tracking-by-model-evolution and tracking-by-assignment (Jug et al., 2014c).

In *tracking-by-model-evolution*, objects are detected in the first frame and a model of their properties, e.g. shape and position, is obtained. This model is then updated greedily for pairs of neighboring frames, thereby tracking all detected objects. For sensible results these methods typically require a high temporal resolution (Kausler et al., 2012; Jug et al., 2014c). Payer et al. (2018) proposed a neural network architecture for tracking-by-model-evolution that jointly tackles

segmentation and tracking while having the ability to handle object divisions. Arbelle and Raviv (2018) incorporate temporal information by using LSTMs. However, such neural network based approaches require vast amounts of annotated training data, which is typically not available for biomedical tracking problems.

Tracking-by-assignment first segments potential object candidates in all time frames (Kausler et al., 2012; Schiegg et al., 2013; Jug et al., 2014a; Schiegg et al., 2014; Magnusson et al., 2015; Kaiser et al., 2018). There are two important cases: a single segmentation hypothesis per object (Kausler et al., 2012; Schiegg et al., 2013) and multiple ones (Schiegg et al., 2014; Kaiser et al., 2018). Multiple segmentation hypotheses may correspond to different (typically overlapping) positions of the same object, to parts of a single object looking like separate ones, or to several objects close to each other looking like a single one. Apart from better tracking quality, multiple segmentation approaches facilitate user-driven proofreading of automated results .

Existing optimization methods for tracking-by-assignment fall into two categories: (*i*) local approaches that attempt to overcome scalability issues by decomposing the overall tracking problem into smaller sub-problems (Jaqaman et al., 2008; Xing et al., 2009; Castanon and Finn, 2011), and (*ii*) global approaches, treating the whole spatio-temporal problem jointly (Padfield et al., 2009; Butt and Collins, 2013; Magnusson et al., 2015; Haubold et al., 2016). While the first type of method scales better with the problem size, the second one leads to better solutions. Among approaches of the second type we distinguish the work of Padfield et al. (2009), that couples multiple min-cost-flow networks to handle divisions and finds an approximate solution to the overall problem with an off-the-shelf LP solver. Another notable contribution is the primal heuristics by Haubold et al. (2016), based on sequentially computing shortest paths in an augmented flow graph which accounts for object divisions. This work generalizes the one of Magnusson et al. (2015) which utilizes the Viterbi algorithm. However, these works neither handle overlapping segmentation hypotheses nor do they provide any bounds on the quality of the proposed solutions. Butt and Collins (2013) employ stochastic gradient descent to maximize a Lagrange dual based on multiple min-cost-flow subproblems, and obtains primal solutions by rounding. However, the method does not handle cell divisions.

After all, the most general models including both multiple hypothesis and object divisions have been addressed with off-the-shelf ILP solvers only (Schiegg et al., 2014; Jug et al., 2014a,b; Kaiser et al., 2018). Building a scalable solver able to compete with, for example, Gurobi in this case seems to be a non-trivial task, which has not been addressed in the literature yet.

Contributions. We propose a new approximate optimization method for cell-tracking problems, which favorably compares with Gurobi in terms of run time and memory footprint, while delivering solutions of a comparable quality. Our method is able to handle cell divisions and multiple segmentation hypotheses. Together with an approximate primal solution it provides a lower bound on the optimum. This is achieved by optimizing the Lagrange dual problem constructed from a new compact decomposition. To optimize the dual we propose a specialized fast converging algorithm based on the block-coordinate ascent principle. The approximate primal solution is obtained with a novel primal heuristics based on conflict resolution and greedy elongation of trajectories. While the dual solver simplifies the objective function by reweighting its costs, the primal one reconstructs an integer solution based on these costs. We empirically show advantages of our framework on publicly available instances of the cell-tracking challenge by Ulman et al. (2017), on instances of developing flywing tissue, and on an instance of nuclei tracking in developing drosophila embryos. These datasets represent different biological applications and exhibit diverse characteristics. Therefore, we believe our method to be applicable to a wide spectrum of cell-tracking problems.

The focus of our work is to improve the optimization stage of a typical tracking-by-assignment pipeline, therefore, we do not address modeling and segmentation aspects of the problem here. We assume that for each time step $t \in \{1, \ldots, T\}$ a set of *segmentation hypotheses* is available along with a set of possible transitions and the corresponding costs.

Mathematical proofs and information about our code and models can be found in the supplement.

2 Standard tracking as ILP

The standard modeling approach for tracking-by-assignment is based on its *problem (hyper-)graph* representation (Lou et al., 2011; Kausler et al., 2012; Schiegg et al., 2013, 2014; Jug et al., 2014a,b), see Figure 1 (b). In the following we omit the prefix *hyper-* and use the *hat* superscript (as in $\hat{\mathcal{V}}$ or $\hat{\mathcal{E}}$) for the standard problem graph to distinguish it from the graph we propose later. Nodes $\hat{\mathcal{V}}$ of a problem graph $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$ are associated with *finite-valued variables*, and edges $\hat{\mathcal{E}} \subseteq 2^{\hat{\mathcal{V}}}$ correspond to the *coupling constraints* between the respective nodes. Here, $2^{\hat{\mathcal{V}}}$ denotes the power set of $\hat{\mathcal{V}}$.

Node set. For tracking-by-assignment the set of nodes $\hat{\mathcal{V}}$ is divided into disjoint subsets $\hat{\mathcal{V}}^t$ corresponding to each time step $t \in \{1, \ldots, T\}$, i.e. $\hat{\mathcal{V}} = \bigcup_{t=1}^{T} \hat{\mathcal{V}}^t$. In its turn, each subset $\hat{\mathcal{V}}^t$ is subdivided into a set $\hat{\mathcal{V}}^t_{\text{det}}$ representing the *segmentation hypothesis (detections)*
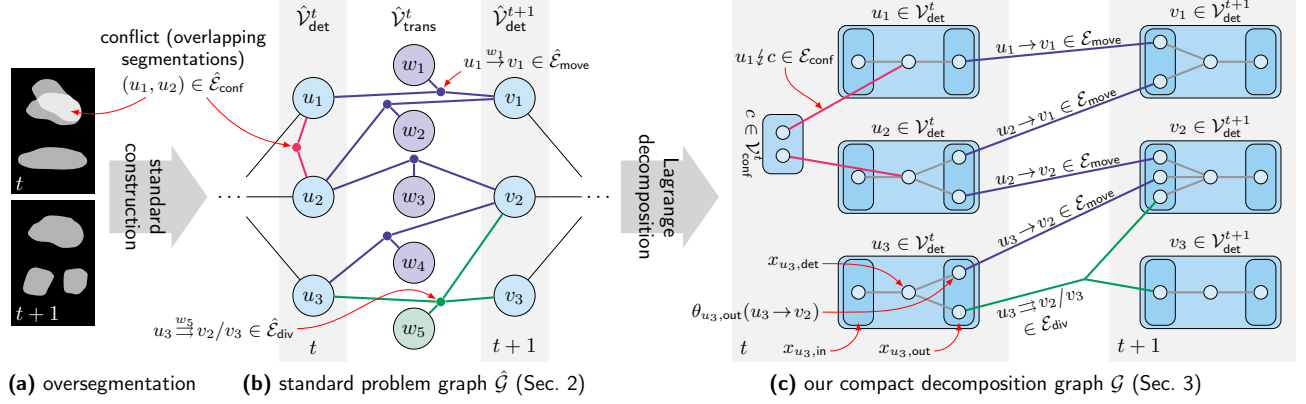
(a) oversegmentation   **(b)** standard problem graph $\hat{\mathcal{G}}$ (Sec. 2)   **(c)** our compact decomposition graph $\mathcal{G}$ (Sec. 3)

Figure 1: Illustration of the standard model and our proposed problem decomposition. Each detection (segmentation hypothesis) in **(a)** is represented by a node in $\hat{\mathcal{V}}_{\mathsf{det}}$ in the standard model **(b)**. Possible transitions between time steps are represented by a node in $\hat{\mathcal{V}}_{\mathsf{trans}}$. Hyper-edges $\hat{\mathcal{E}}$ handle detection and transition coupling constraints. Our Lagrange decomposition **(c)** represents each detection together with its incoming and outgoing transitions by a single node in $\mathcal{V}_{\mathsf{det}}$. Conflicting detections are represented by separate nodes in $\mathcal{V}_{\mathsf{conf}}$. Edges $\mathcal{E}$ correspond to coupling constraints between nodes and refer to transitions (blue), divisions (green), and conflicts (red).

at time step t, and a set $\hat{\mathcal{V}}_{\mathsf{trans}}^t$ representing the possible *transitions (moves, divisions)* from time step t to $t+1$. We will write $\hat{\mathcal{V}}_{\mathsf{det}} = \bigcup_{t=1}^T \hat{\mathcal{V}}_{\mathsf{det}}^t$ and $\hat{\mathcal{V}}_{\mathsf{trans}} = \bigcup_{t=1}^T \hat{\mathcal{V}}_{\mathsf{trans}}^t$ for the sets of all detection and transition nodes.

Each segmentation hypothesis as well as each transition is associated with a *binary variable*, i.e. its value is in the set $\{0,1\}$. We will refer to the variable corresponding to node $v \in \hat{\mathcal{V}}$ as x_v, where $x_v \in \{0,1\}$. A variable is said to be *active* if it assumes value 1.

Edge set. The set of edges $\hat{\mathcal{E}}$ coupling the nodes is divided into subsets $\hat{\mathcal{E}}^t$ corresponding to each time step $t \in \{1,\ldots,T\}$, i.e. $\hat{\mathcal{E}} = \bigcup_{t=1}^T \hat{\mathcal{E}}^t$. In turn, $\hat{\mathcal{E}}^t = \hat{\mathcal{E}}_{\mathsf{move}}^t \cup \hat{\mathcal{E}}_{\mathsf{div}}^t \cup \hat{\mathcal{E}}_{\mathsf{conf}}^t$, where $\hat{\mathcal{E}}_{\mathsf{move}}^t \subseteq \hat{\mathcal{V}}_{\mathsf{det}}^t \times \hat{\mathcal{V}}_{\mathsf{det}}^{t+1} \times \hat{\mathcal{V}}_{\mathsf{trans}}^t$ and $\hat{\mathcal{E}}_{\mathsf{div}}^t \subseteq \hat{\mathcal{V}}_{\mathsf{det}}^t \times (\hat{\mathcal{V}}_{\mathsf{det}}^{t+1})^2 \times \hat{\mathcal{V}}_{\mathsf{trans}}^t$ are edges corresponding to possible moves and divisions of the cells between time steps t and $t+1$, and $\hat{\mathcal{E}}_{\mathsf{conf}}^t \subseteq 2^{\hat{\mathcal{V}}_{\mathsf{det}}^t}$ are the edges prohibiting the activation of conflicting (intersecting) segmentation hypothesis at time step t. Note that for each node in $\hat{\mathcal{V}}_{\mathsf{trans}}^t$ there is exactly one incident edge.

Coupling constraints. Let $(u,v,w) \in \hat{\mathcal{E}}_{\mathsf{move}}^t$ be a possible move from time step t to $t+1$ connecting nodes $u \in \hat{\mathcal{V}}_{\mathsf{det}}^t$ and $v \in \hat{\mathcal{V}}_{\mathsf{det}}^t$ via transition node $w \in \hat{\mathcal{V}}_{\mathsf{trans}}^t$. We write $u \xrightarrow{w} v$ for such edges. Then the set of corresponding coupling constraints is defined as

$$\forall u \xrightarrow{w} v \in \hat{\mathcal{E}}_{\mathsf{move}}: \quad x_w \leq x_u \,\wedge\, x_w \leq x_v, \quad (1)$$

ensuring that if either of the hypothesis is deactivated ($x_u = 0$ or $x_v = 0$), the move is deactivated as well ($x_w = 0$). Analogously, we denote possible divisions $(u,v,v',w) \in \hat{\mathcal{E}}_{\mathsf{div}}^t$ by $u \xrightarrow{w} v/v'$, and obtain the following coupling constraints for divisions:

$$\forall u \xrightarrow{w} v/v' \in \hat{\mathcal{E}}_{\mathsf{div}}: x_w \leq x_u \,\wedge\, x_w \leq x_v \,\wedge\, x_w \leq x_{v'}. \quad (2)$$

For any detection node $v \in \hat{\mathcal{V}}_{\mathsf{det}}$ we denote by $\hat{\mathsf{in}}(v)$ the set of *incoming* transitions, i.e. $\hat{\mathsf{in}}(v) :=$

$\{w \in \hat{\mathcal{V}}_{\mathsf{trans}} \mid \exists u\colon u \xrightarrow{w} v \in \hat{\mathcal{E}}_{\mathsf{move}}$ or $\exists u,v'\colon u \xrightarrow{w} v/v' \in \hat{\mathcal{E}}_{\mathsf{div}} \vee u \xrightarrow{w} v'/v \in \hat{\mathcal{E}}_{\mathsf{div}}\}$. Likewise, for all $u \in \mathcal{V}_{\mathsf{det}}$ we define $\hat{\mathsf{out}}(u) := \{w \in \hat{\mathcal{V}}_{\mathsf{trans}} \mid \exists v\colon u \xrightarrow{w} v \in \hat{\mathcal{E}}_{\mathsf{move}}$ or $\exists v,v'\colon u \xrightarrow{w} v/v' \in \hat{\mathcal{E}}_{\mathsf{div}}\}$ as the set of *outgoing* transitions.

To guarantee that each hypothesis at time step t is matched to at most one at time steps $t+1$ and $t-1$, *uniqueness constraints* are introduced as follows:

$$\forall v \in \hat{\mathcal{V}}_{\mathsf{det}}: \sum_{w \in \hat{\mathsf{in}}(v)} x_w \leq 1 \quad \text{and} \quad \sum_{w \in \hat{\mathsf{out}}(v)} x_w \leq 1. \quad (3)$$

Finally, conflicting segmentation hypothesis are connected via similar constraints:

$$\forall c \in \hat{\mathcal{E}}_{\mathsf{conf}}: \sum_{v \in c} x_v \leq 1. \quad (4)$$

Objective function. Let $\hat{\mathcal{X}} \subseteq \{0,1\}^{|\hat{\mathcal{V}}|}$ be the set of binary vectors x satisfying all coupling constraints defined by (1)-(4). Each coordinate x_v, $v \in \hat{\mathcal{V}}$, is associated with a *cost* $\theta_v \in \mathbb{R}$ based on image data (for segmentation hypothesis) and geometric priors (for transitions). The goal of tracking is to find an assignment $x \in \hat{\mathcal{X}}$ that minimizes the cost of the active binary variables, i.e. which solves

$$\min_{x \in \hat{\mathcal{X}}} \langle \theta, x \rangle, \quad (5)$$

where $\theta = (\theta_v)_{v \in \hat{\mathcal{V}}}$. Problem (5) is the standard ILP representation of the tracking problem. In this form it is usually addressed (see e.g. Lou et al., 2011; Kausler et al., 2012; Schiegg et al., 2013, 2014; Jug et al., 2014a,b) by off-the-shelf solvers like Gurobi (2018) or CPLEX (2017). However, the run time and memory requirements of these solvers rapidly grow with the size of the input. Moreover, even solving an LP relaxation of the above problem, i.e. considering a vector in $[0,1]^{|\hat{\mathcal{V}}|}$ satisfying (1)-(4), requires a significant time

using standard solvers, as they are based on simplex or interior point methods with a super-linear iteration complexity. Note that Kappes et al. (2015) shows that first order subgradient-based methods perform even slower than the standard solvers.

3 Our decomposable representation

Efficiency of large-scale approximate optimization methods heavily depends on the problem decomposition used to build a dual problem. A good decomposition should contain small number of easily tractable subproblems. Consider a trivial decomposition of (5), when every binary variable corresponds to a separate subproblem. It would satisfy the tractability condition, but the large number of subproblems would significantly slow down the optimization. Therefore, below we give an alternative representation of the problem (5), which leads to a natural decomposition with a much smaller number of easily tractable subproblems.

Lagrange decomposition idea. Assume we want to minimize a function $F(x)$ representable as $F(x) = F_1(x) + F_2(x)$. The Lagrange decomposition (Guignard and Kim, 1987a,b; Guignard, 2003; Sontag et al., 2011) duplicates the variable x and introduces the equality constraint $x_1 = x_2$, i.e. $\min_x F(x) = \min_{x_1, x_2 \colon x_1 = x_2} \big(F_1(x_1) + F_2(x_2) \big)$. Dualization of the constraint $x_1 = x_2$ leads to the *Lagrange dual problem*, which forms a lower bound for the original problem:

$$\min_x F(x) \geq \max_\lambda \min_{x_1, x_2} \big(F_1(x_1) + F_2(x_2) + \langle \lambda, x_1 - x_2 \rangle \big)$$

$$= \max_\lambda \Big[\min_{x_1} \big(F_1(x_1) + \langle \lambda, x_1 \rangle \big) + \min_{x_2} \big(F_2(x_2) - \langle \lambda, x_2 \rangle \big) \Big]. \quad (6)$$

Tightness of the lower bound as well as the efficiency of its maximization depend on the decomposition of F into F_1 and F_2. Ideally, the minimization subproblems over x_1 and x_2 are solvable in closed form, and the coupling constraint $x_1 = x_2$ is only violated in a small subset of coordinates of the subproblem minima. The dual vector λ allows to reweight the functions associated with the duplicated variables during optimization in order to reduce violations of coupling constraints.

Decomposed graph. We will apply the Lagrange decomposition idea to problem (5). To this end, we first duplicate all binary variables and then regroup them. Each group corresponds to a new graph node. This leads to considerably less nodes. Although each node is associated with a non-binary variable, its minimal value can still be efficiently found. All coupling constraints turn into simple equalities as in the general scheme (6).

Graph structure. Our graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, see Figure 1 (c), contains only two types of nodes: detection and conflict nodes. The transition variables are duplicated (tripled for divisions) and their copies are assigned to the corresponding detection nodes. Since

each detection corresponds to a large number of transitions, this significantly decreases the graph size. The detection variables are duplicated as well and their copies are assigned to the detection and conflict nodes. Below we give the formal definitions.

The set of graph nodes is defined as $\mathcal{V} = \mathcal{V}_{\mathsf{det}} \cup \mathcal{V}_{\mathsf{conf}}$, where $\mathcal{V}_{\mathsf{det}} := \hat{\mathcal{V}}_{\mathsf{det}}$ and $\mathcal{V}_{\mathsf{conf}} := \hat{\mathcal{E}}_{\mathsf{conf}}$, i.e. the detection nodes and conflict edges in the standard model correspond to detection and conflict nodes in our model. As in the standard model, $\mathcal{V}_{\mathsf{det}}^t$ and $\mathcal{V}_{\mathsf{conf}}^t$ denote the nodes at time step $t \in \{1, \ldots, T\}$.

Each edge in the edge set $\mathcal{E} = \mathcal{E}_{\mathsf{move}} \cup \mathcal{E}_{\mathsf{div}} \cup \mathcal{E}_{\mathsf{conf}}$ corresponds to either a transition or conflict. The transition edges divide into $\mathcal{E}_{\mathsf{move}} := \hat{\mathcal{E}}_{\mathsf{move}}$, $\mathcal{E}_{\mathsf{move}} \subseteq (\mathcal{V}_{\mathsf{det}})^2$, corresponding to moves, and $\mathcal{E}_{\mathsf{div}} := \hat{\mathcal{E}}_{\mathsf{div}}$, $\mathcal{E}_{\mathsf{div}} \subseteq (\mathcal{V}_{\mathsf{det}})^3$, corresponding to divisions. As for the standard model, we will denote an edge $(u, v) \in \mathcal{E}_{\mathsf{move}}$ by $u \to v$, and an edge $(u, v, w) \in \mathcal{E}_{\mathsf{div}}$ by $u \rightrightarrows v/w$.

Conflict edges (u, c), denoted by $u \nmid c$, are introduced between any detection node $u \in \mathcal{V}_{\mathsf{det}}$ and conflict node $c \in \mathcal{V}_{\mathsf{conf}}$ as soon as $\hat{u} \in \hat{c}$ for the corresponding $\hat{u} \in \hat{\mathcal{V}}_{\mathsf{det}}$ and $\hat{c} \in \hat{\mathcal{E}}_{\mathsf{conf}}$. Note, $\mathcal{E}_{\mathsf{conf}} \subseteq \mathcal{V}_{\mathsf{det}} \times \mathcal{V}_{\mathsf{conf}}$. When considering a conflict node $c \in \mathcal{V}_{\mathsf{conf}}$, we also use c to refer to all detection nodes that are part of the conflict, i.e. $u \in c$ if and only if $u \nmid c \in \mathcal{E}_{\mathsf{conf}}$. Furthermore, for any $u \in \mathcal{V}_{\mathsf{det}}$ we define $\mathsf{conf}(u) := \{u' \nmid c \in \mathcal{E}_{\mathsf{conf}} \mid u = u'\}$ as the set of all conflicts concerning u.

Detection variables. As noted above, nodes of the graph $\mathcal{G}$ correspond to variables having more than two states. These states are represented by binary vectors.

To define the state space of the detection variables we first introduce the sets $\mathsf{in}(u)$ and $\mathsf{out}(u)$, corresponding to $\hat{\mathsf{in}}(\hat{u})$ and $\hat{\mathsf{out}}(\hat{u})$ in $\hat{\mathcal{G}}$. For any detection node $u' \in \mathcal{V}_{\mathsf{det}}$ we denote by $\mathsf{in}(u')$ the set of all incoming transitions, i.e. $\mathsf{in}(u') := \{u \to v \in \mathcal{E}_{\mathsf{move}} \mid v = u'\} \cup \{u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}} \mid v = u' \text{ or } w = u'\}$. Analogously, we define $\mathsf{out}(u') := \{u \to v \in \mathcal{E}_{\mathsf{move}} \mid u = u'\} \cup \{u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}} \mid u = u'\}$ as the set of all outgoing transitions.

Consider $u \in \mathcal{V}_{\mathsf{det}}$. The set of states $\mathcal{X}_u$, which models whether the detection u is active, and if so, which incoming and outgoing edge is active, is defined as

$$\mathcal{X}_u = \left\{ \begin{pmatrix} x_{\mathsf{det}} \in \{0,1\} \\ x_{\mathsf{in}} \in \{0,1\}^{|\mathsf{in}(u)|} \\ x_{\mathsf{out}} \in \{0,1\}^{|\mathsf{out}(u)|} \end{pmatrix} \middle| \begin{array}{l} \langle \mathbb{1}, x_{\mathsf{in}} \rangle \leq x_{\mathsf{det}}, \\ \langle \mathbb{1}, x_{\mathsf{out}} \rangle \leq x_{\mathsf{det}} \end{array} \right\}. \quad (7)$$

The scalar products $\langle \mathbb{1}, x_{\mathsf{in}} \rangle$ and $\langle \mathbb{1}, x_{\mathsf{out}} \rangle$ express the number of activated incoming and outgoing transitions. Note that they can only be non-zero if the detection is active, i.e. if $x_{\mathsf{det}} = 1$. Below we use $x_{\mathsf{in}}(e)$ for any incoming edge $e \in \mathsf{in}(u)$ to refer to the value of this edge in the current state. Analogously, we use $x_{\mathsf{out}}(e)$ for edges $e \in \mathsf{out}(u)$. Recalling the standard

model, for any $w \in \hat{\mathcal{V}}_{\mathsf{trans}}$ associated with $u \overset{w}{\to} v \in \hat{\mathcal{E}}_{\mathsf{move}}$ ($u \overset{w}{\rightrightarrows} v/v' \in \hat{\mathcal{E}}_{\mathsf{div}}$) the binary variable x_w is split into two (three) variables, one belonging to $\mathsf{out}(u)$ and another to $\mathsf{in}(v)$ (and $\mathsf{in}(v')$).

With each detection $u \in \mathcal{V}_{\mathsf{det}}$ we associate a cost vector $\theta_u = (\theta_{\mathsf{det}}, \theta_{\mathsf{in}}, \theta_{\mathsf{out}})$ consisting of the cost $\theta_{\mathsf{det}} \in \mathbb{R}$ for activating the detection, and costs $\theta_{\mathsf{in}} \in \mathbb{R}^{|\mathsf{in}(u)|}$ and $\theta_{\mathsf{out}} \in \mathbb{R}^{|\mathsf{out}(u)|}$ associated with the incoming and outgoing edges. $\theta_{\mathsf{det}} = \theta_{\hat{u}}$, where $\hat{u} \in \hat{\mathcal{V}}_{\mathsf{det}}$ is the corresponding detection node in the standard model. θ_{in} and θ_{out} are obtained by splitting the given transition costs between incoming and outgoing variables, i. e. $\theta_{\mathsf{in}}(e) = \frac{1}{|e|}\theta_{\hat{w}}$ for all $e \in \mathsf{in}(u)$, and, analogously, $\theta_{\mathsf{out}}(e) = \frac{1}{|e|}\theta_{\hat{w}}$ for all $e \in \mathsf{out}(u)$, where $\hat{w} \in \hat{\mathcal{V}}_{\mathsf{trans}}$ is the transition node corresponding to e in the standard model. So each admissible state $x = (x_{\mathsf{det}}, x_{\mathsf{in}}, x_{\mathsf{out}}) \in \mathcal{X}_u$ has a linear cost $\langle \theta_u, x \rangle = \langle \theta_{\mathsf{det}}, x_{\mathsf{det}} \rangle + \langle \theta_{\mathsf{in}}, x_{\mathsf{in}} \rangle + \langle \theta_{\mathsf{out}}, x_{\mathsf{out}} \rangle$.

Conflict variables. Let $c \in \mathcal{V}_{\mathsf{conf}}$. The associated set of states $\mathcal{X}_c$, that models which of the conflicting detections, if any, is active, can be written as

$$\mathcal{X}_c = \left\{ x \in \{0,1\}^{|c|} \mid \langle \mathbb{1}, x \rangle \leq 1 \right\}. \tag{8}$$

For any detection node $u \in c$ we write $x(u)$ to refer to the value of this detection in the current state. With c we associate a cost vector $\theta_c \in \mathbb{R}^{|c|}$. These costs are initially zero, but may change during optimization. So each admissible state $x \in \mathcal{X}_c$ has a linear cost $\langle \theta_c, x \rangle$.

Coupling constraints. The semantics of each single move, division and detection activation is split between the states of multiple nodes in our problem graph. Consider a move $u \to v$ from detection node u to v. To obtain a consistent solution, we require moves to be consistent in u and v, i. e. $x_{u,\mathsf{out}}(u \to v) = x_{v,\mathsf{in}}(u \to v)$ for any feasible combination of states $x_u \in \mathcal{X}_u$, $x_v \in \mathcal{X}_v$. Analogous considerations for divisions and conflicts result in the following coupling constraints for $\mathcal{G}$:

$$\begin{aligned}
\forall e = u \to v \in \mathcal{E}_{\mathsf{move}}: & \quad x_{u,\mathsf{out}}(e) = x_{v,\mathsf{in}}(e), \\
\forall e = u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}}: & \quad x_{u,\mathsf{out}}(e) = x_{v,\mathsf{in}}(e), \\
\forall e = u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}}: & \quad x_{u,\mathsf{out}}(e) = x_{w,\mathsf{in}}(e), \\
\forall u \nleftrightarrow c \in \mathcal{E}_{\mathsf{conf}}: & \quad x_{u,\mathsf{det}} = x_c(u).
\end{aligned} \tag{9}$$

The set of all state assignments satisfying all coupling constraints can then be written as

$$\mathcal{X} = \left\{ \left(\{x_v \in \mathcal{X}_v\}_{v \in \mathcal{V}_{\mathsf{det}}}, \{x_c \in \mathcal{X}_c\}_{c \in \mathcal{V}_{\mathsf{conf}}} \right) \mid (9) \right\}.$$

Minimization problem. Our graph decomposition naturally gives rise to the minimization problem

$$\min_{x \in \mathcal{X}} \left[E(\theta, x) := \sum_{u \in \mathcal{V}_{\mathsf{det}}} \langle \theta_u, x_u \rangle + \sum_{c \in \mathcal{V}_{\mathsf{conf}}} \langle \theta_c, x_c \rangle \right]. \tag{10}$$

The goal is to find an optimal state assignment that satisfies all coupling constraints given costs θ.

Dualization of coupling constraints. Let us return to the general idea of the Lagrange decomposition (6).

Assume $F_1(x) = \langle \theta_1, x \rangle$. Then $F_1(x) - \langle \lambda, x \rangle = \langle \theta_1 - \lambda, x \rangle$. Similarly, if $F_2(x) = \langle \theta_2, x \rangle$, then $F_2(x) + \langle \lambda, x \rangle = \langle \theta_2 + \lambda, x \rangle$. In other words, λ shifts the costs between parts of the decomposed problem. Since the value of the objective $F(x) = \langle \theta_1 - \lambda, x_1 \rangle + \langle \theta_2 + \lambda, x_2 \rangle$ remains the same for any value of λ if $x_1 = x_2$, this is also referred to as a *reparametrization* of the problem.

Dualizing all coupling constraints (9) in problem (10) results in the following reparametrized cost vectors:

Reparametrization. A *reparametrization* is a vector $\lambda \in \Lambda := \mathbb{R}^{|\mathcal{E}_{\mathsf{move}}| + 2|\mathcal{E}_{\mathsf{div}}| + |\mathcal{E}_{\mathsf{conf}}|}$. Its coordinates will be indexed with edges of the graph $\mathcal{G}$. That is, $\lambda(e) \in \mathbb{R}$ is the dual variable corresponding to the constraint $x_{u,\mathsf{in}}(e) = x_{v,\mathsf{out}}(e)$ if $e = u \to v \in \mathcal{E}_{\mathsf{move}}$, and $x_{u,\mathsf{det}} = x_c(u)$ if $e = u \nleftrightarrow c \in \mathcal{E}_{\mathsf{conf}}$. The only exception are divisions, since two constraints must be dualized for each $e = u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}}$, namely $x_{u,\mathsf{out}}(e) = x_{v,\mathsf{in}}(e)$ and $x_{u,\mathsf{out}}(e) = x_{w,\mathsf{in}}(e)$, cf. (9). The corresponding dual variables are denoted as $\lambda_v(e)$ and $\lambda_w(e)$ respectively. The *reparametrized costs* θ^λ are defined as

$$\forall c \in \mathcal{V}_{\mathsf{conf}} \, \forall u \in c: \quad \theta_c^\lambda(u) := \theta_c(u) + \lambda(u \nleftrightarrow c),$$

$$\forall u \in \mathcal{V}_{\mathsf{det}}: \quad \theta_{u,\mathsf{det}}^\lambda := \theta_{u,\mathsf{det}} - \sum_{e \in \mathsf{conf}(u)} \lambda(e),$$

$$\forall u \in \mathcal{V}_{\mathsf{det}} \, \forall e \in \mathsf{in}(u):$$

$$\theta_{u,\mathsf{in}}^\lambda(e) := \begin{cases} \theta_{u,\mathsf{in}}(e) + \lambda(e), & e \in \mathcal{E}_{\mathsf{move}} \\ \theta_{u,\mathsf{in}}(e) + \lambda_u(e), & e \in \mathcal{E}_{\mathsf{div}} \end{cases},$$

$$\forall u \in \mathcal{V}_{\mathsf{det}} \, \forall e \in \mathsf{out}(u):$$

$$\theta_{u,\mathsf{out}}^\lambda(e) := \begin{cases} \theta_{u,\mathsf{out}}(e) - \lambda(e), & e \in \mathcal{E}_{\mathsf{move}} \\ \theta_{u,\mathsf{out}}(e) - \sum_{v' \in \{v,w\}} \lambda_{v'}(e), & e = u \rightrightarrows v/w \in \mathcal{E}_{\mathsf{div}} \end{cases}.$$

Each element of λ shifts the cost between two copies of variables in different subproblems coupled by an edge in $\mathcal{E}$. This shift does not influence the optimization objective as long as the coupling constraints (9) hold:

Proposition 1. $\forall x \in \mathcal{X}, \ \lambda \in \Lambda: E(\theta, x) = E(\theta^\lambda, x)$.

Our dual is built similarly to the general scheme (6):

Proposition 2. *Dualizing all coupling constraints* (9) *in the objective* (10) *yields the* Lagrange dual problem $\max_{\lambda \in \Lambda} D(\lambda)$, *where*

$$D(\lambda) := \sum_{u \in \mathcal{V}_{\mathsf{det}}} \min_{x_u \in \mathcal{X}_u} \langle \theta_u^\lambda, x_u \rangle + \sum_{c \in \mathcal{V}_{\mathsf{conf}}} \min_{x_c \in \mathcal{X}_c} \langle \theta_c^\lambda, x_c \rangle. \tag{11}$$

Obviously, $D(\lambda)$ is concave and piecewise linear, i. e. non-smooth. By construction, $\max_{\lambda \in \Lambda} D(\lambda) \leq \min_{x \in \mathcal{X}} E(\theta, x)$. The dual objective $D(\lambda)$ is a sum of small-sized minimization problems. Due to the structure of (7) and (8), each subproblem related to a graph node can be solved in linear time. As we show in the supplement, the maximization of the dual (11) yields the same value as the natural LP relaxation of (10).

Algorithm 1: Dual optimization

$T' \leftarrow \{1, \ldots, T\}; \lambda \leftarrow 0$
while *not converged* **do**
    for $t \in T'$ **do**
        for $v \in \mathcal{V}_{\mathsf{det}}^t$ **do**
            compute $\Delta_v^{\uparrow}$ on θ^λ ; $\lambda \leftarrow \lambda + \Delta_v^{\uparrow}$
        for $c \in \mathcal{V}_{\mathsf{conf}}^t$ **do**
            compute $\Delta_c^{\uparrow}$ on θ^λ; $\lambda \leftarrow \lambda + \Delta_c^{\uparrow}$
        estimate assignment (Alg. 2, optional)
        for $v \in \mathcal{V}_{\mathsf{det}}^t$ **do**
            compute $\Delta_v^{\rightarrow}$ on θ^λ ; $\lambda \leftarrow \lambda + \Delta_v^{\rightarrow}$
            (backward direction: use $\Delta_v^{\leftarrow}$)
    reverse the order of T'

---

4 Dual block-coordinate ascent (BCA)

In order to maximize (11), we developed an algorithm based on the BCA principle, as such methods perform competitively for similar relaxations of large-scale combinatorial problems. These techniques received a lot of attention in connection with the local polytope relaxation of the discrete energy minimization problem (Kappes et al., 2015). BCA methods like TRW-S (Kolmogorov, 2006), SRMP (Kolmogorov, 2015) or the recently proposed DMM (Shekhovtsov et al., 2016), MPLP++ (Tourani et al., 2018) and SPAM Tourani et al. (2020) notably outperform off-the-shelf solvers as well as dedicated subgradient-based methods. However, they are inapplicable in our case, due to a substantially different problem structure. Swoboda et al. (2017) partially fills this gap by proposing a general framework for constructing dual BCA algorithms for a substantial subclass of combinatorial problems, which covers our problem as well. However, our experiments with this framework did not lead to an improvement over Gurobi. Therefore, we constructed a new algorithm, which resembles the local polytope technique (Kolmogorov, 2006, 2015; Shekhovtsov et al., 2016; Tourani et al., 2018; Savchynskyy, 2019) and at the same time uses the results of Swoboda et al. (2017) to guarantee monotonicity of the dual improvement.[1]

Dual BCA algorithm. Algorithm 1 is a realization of the BCA principle for the Lagrange dual (11) and guarantees its monotonous improvement. It contains four types of updates (also referred to as passing *messages*). The first two, $\Delta_u^{\uparrow}$ and $\Delta_c^{\uparrow}$, called *conflict updates*, reweight the detection variable costs $\theta_{u,\mathsf{det}}$ and conflict variable costs $\theta_c(u)$. The second two, $\Delta_u^{\rightarrow}$ and $\Delta_u^{\leftarrow}$, called *transition updates*, reweight costs $\theta_{u,\mathsf{in}}$ and $\theta_{v,\mathsf{out}}$ across consecutive time steps. All update vectors $\Delta_u^{\uparrow}$, $\Delta_c^{\uparrow}$, $\Delta_u^{\rightarrow}$ and $\Delta_u^{\leftarrow}$ are in Λ, with zero assigned to the unaffected coordinates. We process the time steps sequentially. We first perform all conflict updates

within the current time step, then we propagate the costs to the next time step via transition updates.

Conflict updates. In the dual problem $D(\lambda)$ a detection will favor activation ($x_{\mathsf{det}} = 1$) if its locally minimal state has negative cost. Ideally, only a single detection connected to a particular conflict node should be active. Otherwise, a coupling constraint on at least one of the edges in $\mathcal{E}_{\mathsf{conf}}$ or the constraint in $\mathcal{X}_c$ is violated, see (8), (9). The following updates encourage agreement of the local minimizers of the conflict nodes and the associated detection nodes. We define for all $u \in \mathcal{V}_{\mathsf{det}}$, $e \in \mathsf{conf}(u)$:

$$\Delta_u^{\uparrow}(e) := \min_{x \in \mathcal{X}_u : \, x_{\mathsf{det}} = 1} \frac{\langle \theta_u, x \rangle}{|\mathsf{conf}(u)|}.$$

Intuitively, $\Delta_u^{\uparrow}$ redistributes as much of the cost as possible to the connected conflict nodes while preserving the locally optimal state in the detection node. Similarly, we define for all $c \in \mathcal{V}_{\mathsf{conf}}$, $u \in c$:

$$\Delta_c^{\uparrow}(e) := -\theta_c(u) + \frac{1}{2} \big[\langle \theta_c, z_c^{\star} \rangle + \langle \theta_c, z_c^{\star\star} \rangle \big],$$

where $e = u \, \mathord{\not\downarrow} \, c$, and $z_c^{\star} = \arg\min_{x \in \mathcal{X}_c} \langle \theta_c, x \rangle$ is the best and $z_c^{\star\star} = \arg\min_{x \in \mathcal{X}_c \setminus \{z_c^{\star}\}} \langle \theta_c, x \rangle$ the second-best state. $\Delta_c^{\uparrow}$ shifts the cost back such that only the most promising detection ends up with a negative activation cost.

Transition updates. To propagate information across time steps, we introduce the dual updates $\Delta_u^{\rightarrow}$ and $\Delta_u^{\leftarrow}$. We first define for all $u \in \mathcal{V}_{\mathsf{det}}$:

$$x_u^{\star} := \mathop{\arg\min}_{x \in \mathcal{X}_u : \, x_{\mathsf{det}} = 1} \langle \theta_u, x \rangle, \qquad y_u^{\star} := \mathop{\arg\min}_{\substack{x \in \mathcal{X}_u : \, x_{\mathsf{det}} = 1, \\ x_{\mathsf{in}} \neq x_{u,\mathsf{in}}^{\star}, \, x_{\mathsf{out}} \neq x_{u,\mathsf{out}}^{\star}}} \langle \theta_u, x \rangle.$$

Intuitively, $x_u^{\star}$ is the best state of u under the assumption that u is active, while $y_u^{\star}$ is the best state differing from $x_u^{\star}$ in the incoming and outgoing edge.

Similar to the conflict updates, the transition updates propagate as much information as possible by setting all the outgoing respectively incoming costs to the same value. We define for all $u \in \mathcal{V}_{\mathsf{det}}$, $e \in \mathsf{out}(u)$:

$$\Delta_u^{\rightarrow}(e) := \min_{\substack{x \in \mathcal{X}_u : \\ x_{\mathsf{out}}(e) = 1}} \langle \theta_u, x \rangle - \Theta_{u,\mathsf{out}}, \text{ if } e \in \mathcal{E}_{\mathsf{move}}$$

$$(\Delta_u^{\rightarrow})_v(e) := \frac{1}{2} \Big[\min_{\substack{x \in \mathcal{X}_u : \\ x_{\mathsf{out}}(e) = 1}} \langle \theta_u, x \rangle - \Theta_{u,\mathsf{out}} \Big], \text{ if } e = u \rightrightarrows v/w$$

where

$$\Theta_{u,\mathsf{out}} := \min\Big\{ 0, \tfrac{1}{2} \big[\langle \theta_u, x_u^{\star} \rangle + \langle \theta_u, (1, x_{u,\mathsf{in}}^{\star}, y_{u,\mathsf{out}}^{\star}) \rangle \big] \Big\},$$

is either 0 or the mean value between the cost of $x_u^{\star}$ and the next-best state with a different outgoing edge.

Similarly, for the updates in the opposite direction, we define for all $v \in \mathcal{V}_{\mathsf{det}}$, $e \in \mathsf{in}(v)$:

$$\Delta_v^{\leftarrow}(e) := -\min_{x \in \mathcal{X}_v, x_{\mathsf{in}}(e) = 1} \langle \theta_v, x \rangle + \Theta_{v,\mathsf{in}}, \text{ if } e \in \mathcal{E}_{\mathsf{move}}. \quad (12)$$

Otherwise, if $e \in \mathcal{E}_{\mathsf{div}}$, $(\Delta_v^{\leftarrow})_v(e)$ is assigned the right-

---

[1]By Swoboda et al. (2017, Lem.1) our updates are *admissible*, but do not satisfy maximality condition (eq. 15).

Algorithm 2: Primal heuristics for time step t

$$\mu^\star \leftarrow \underset{\mu \in \{0,1\}^{|\mathcal{V}_{\mathsf{det}}^t|}}{\arg\min} \langle s, \mu \rangle \text{ s.t. } \sum_{v \in c} \mu_v \leq 1 \ \forall c \in \mathcal{V}_{\mathsf{conf}}^t \quad (13)$$

for $v \in \mathcal{V}_{\mathsf{det}}^t$ **do**
$\quad \lfloor \ x_v \leftarrow \mathsf{OFF} \ \textbf{if} \ \mu_v^\star = 0$
order elements v of $\mathcal{V}_{\mathsf{det}}^t$ by their score $s(v)$
for $v \in \mathcal{V}_{\mathsf{det}}^t$ with $x_v \neq \mathsf{OFF}$ **do**
$\quad \mathcal{X}_v' \leftarrow \left\{ x \in \mathcal{X}_v \ \middle| \ \begin{matrix} x \text{ does not violate} \\ \text{coupling constraints} \end{matrix} \right\}$
$\quad$ assign $x_{v,\mathsf{in}}, x_{v,\mathsf{det}}$ according to $\arg\min_{x \in \mathcal{X}_v'} \langle \theta_v^\lambda, x \rangle$
$\quad$ (backward direction: use $x_{v,\mathsf{out}}$ instead)
$\quad$ propagate state x_v across edges in $\mathcal{E}$

hand-side of (12). Here

$$\Theta_{v,\mathsf{in}} := \min\left\{0, \tfrac{1}{2}\left[\langle \theta_v, x_v^\star \rangle + \langle \theta_v, (1, y_{v,\mathsf{in}}^\star, x_{v,\mathsf{out}}^\star) \rangle\right]\right\},$$

is either 0 or the mean value between the cost of $x_v^\star$ and the next-best state with a different incoming edge.

Proposition 3. *Dual updates* $\Delta \in \{\Delta_u^\leftarrow, \Delta_u^\rightarrow, \Delta_u^\uparrow \mid u \in \mathcal{V}_{\mathsf{det}}\} \cup \{\Delta_c^\uparrow \mid c \in \mathcal{V}_{\mathsf{conf}}\}$ *monotonically increase the dual function, i. e.* $\forall \lambda \in \Lambda \colon D(\lambda) \leq D(\lambda + \Delta)$.

Since the optimal dual value is bounded from above by the optimum of (5), the monotonicity implies convergence of the sequence of dual values. The limit value of the sequence, though, need not be the dual optimum. This is a well-known property of block-coordinate-ascent methods, which may get stuck when applied to non-smooth functions.

5 Primal heuristics

For solving the tracking problem (10) computing a reparametrization λ is not enough. The main goal is to obtain a feasible primal assignment $x \in \mathcal{X}$ corresponding to a low objective value. Generally, this is non-trivial, since solutions to the node subproblems are usually inconsistent, even if an optimal dual solution $\lambda^\star = \arg\max_{\lambda' \in \Lambda} D(\lambda')$ is considered. As mentioned in Section 1, existing techniques either do not handle overlapping segmentation hypotheses or do not allow for cell divisions. Below, we propose a new primal heuristics that handles the considered general case and produces high-quality feasible primal assignments even for non-optimal dual vectors λ.

Temporal direction. Similar to the dual optimization algorithm our primal heuristics works in both temporal directions. For the sake of simplicity we restrict ourselves to the explanation of the forward direction. After obtaining the primal assignments for each direction, we keep the best of the two.

Incremental estimates. We estimate primal solution for each time frame sequentially one by one, starting from $t = 0$ and ending with $t = T$. We assume that primal assignments $x_u \in \mathcal{X}_u$ for $u \in \mathcal{V}_{\mathsf{det}}^{t'}$, $t' \in \{1, \ldots, t-1\}$, are available when applying Algorithm 2 for time step t.
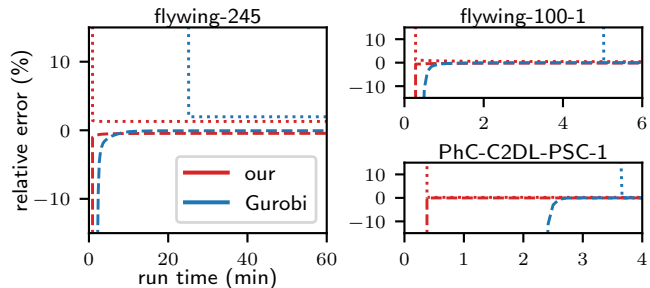

Figure 2: Lower- (＝＝) and upper-bound (⋯⋯) for our solver and Gurobi on selected instances. We obtain high-quality solutions after only a few iterations.

Conflict resolution. Before we look at edges that connect different time steps, we first resolve the conflicts within the current time step t. Overall, we want to activate the most promising detection hypotheses while still obeying all coupling constraints in $\mathcal{E}_{\mathsf{conf}}$. We score the individual detections $v \in \mathcal{V}_{\mathsf{det}}$ by their locally cost-optimal state conditioned on actually activating v, i. e. $s(v) := \min_{x \in \mathcal{X}_v : x_{\mathsf{det}} = 1} \langle \theta_v^\lambda, x \rangle$. To resolve the conflicts of the current time step t, we solve with Gurobi the *weighted independent set problem* of all detections in t:

$$\min_{\mu \in \{0,1\}^{|\mathcal{V}_{\mathsf{det}}^t|}} \langle s, \mu \rangle \quad \text{s.t.} \sum_{v \in c} \mu_v \leq 1 \quad \forall c \in \mathcal{V}_{\mathsf{conf}}^t \ , \quad (13)$$

where $s = (s(v))_{v \in \mathcal{V}_{\mathsf{det}}^t}$. Note that even though the problem (13) is NP-hard, its encountered instances are small and can be solved almost instantly. In comparison to cheaper approaches, using (13) yields significantly better tracking solutions.

Transition assignments. After resolving all conflicts for the current time step t, we look for a consistent assignment for $x_{v,\mathsf{in}}$ for all $v \in \mathcal{V}_{\mathsf{det}}^t$ with $\mu_v^\star = 1$, where $\mu^\star$ is a solution of (13). Ideally, we would like to set them to a locally optimal state given their current cost θ_v^λ. Unfortunately, this will in general violate coupling constraints in $\mathcal{E}_{\mathsf{move}}^{t-1}$ and $\mathcal{E}_{\mathsf{div}}^{t-1}$. Therefore, we use a greedy approach and process all possible nodes, i. e. all $v \in \mathcal{V}_{\mathsf{det}}^t$ with $\mu_v^\star = 1$, sequentially. While optimizing $x_{v,\mathsf{in}}$ for a given node v we ignore all options that would violate coupling constraints with already fixed nodes. Obviously, such a procedure heavily depends on the node ordering. As we want to address the most promising states first, we order all nodes by their score $s(v)$ starting with the state with the best, i. e. lowest, score.

Propagation. When $x_{v,\mathsf{in}}$ is assigned we propagate this across incident coupling constraints in $\mathcal{E}_{\mathsf{move}}^{t-1} \cup \mathcal{E}_{\mathsf{div}}^{t-1}$. Eventually, this results in all variables $x_{u,\mathsf{out}}$ being set for all nodes $u \in \mathcal{V}_{\mathsf{det}}^{t-1}$ of the previous time step. This procedure leads to a consistent variable assignment up to and including time step t. Continuing the procedure for all time steps leads to a consistent variable assignment for the whole problem. We run the primal heuristics repeatedly every 25 dual iterations since the assignment usually improves with dual improvement.

Instance	Our Solver			TRA	Gurobi				Improvement	
	time (s)	mem. (MiB)	err. (%)		time relax. (s)	/ equ. (s)	/ opt. (s)	mem. (MiB)	time	mem.
drosophila	**2.7**	**157**	0.00	0.9999	0.9	9.2	9.3	1154	**3.4**×	**7.3**×
flywing-100-1	**82.4**	**474**	0.65	0.9867	270.6	301.8	554.9	5522	**3.7**×	**11.6**×
flywing-100-2	**78.5**	**490**	0.98	0.9826	243.0	275.2	2311.5	8527	**3.5**×	**17.4**×
flywing-245	**159.8**	**1192**	1.29[1]	—[2]	1159.0	9540.6	—[3]	20756	**59.7**×	**17.4**×
Fluo-C2DL-MSC-1	1.9	**53**	0.38	0.9922	0.0	**0.5**	**0.6**	90	0.3×	**1.7**×
Fluo-C2DL-MSC-2	2.0	**50**	0.08	0.9863	0.0	**0.1**	**0.1**	61	0.1×	**1.2**×
Fluo-N2DH-GOWT1-1	**0.5**	**58**	0.00	1.0000	0.0	0.6	0.6	153	**1.2**×	**2.6**×
Fluo-N2DH-GOWT1-2	**0.3**	**65**	0.00	1.0000	0.0	1.1	1.1	196	**3.3**×	**3.0**×
PhC-C2DL-PSC-1	**22.7**	**930**	0.23	0.9952	40.9	219.7	267.6	13199	**9.7**×	**14.2**×
PhC-C2DL-PSC-2	**17.9**	**708**	0.14	0.9969	22.2	127.5	156.9	9870	**7.1**×	**13.9**×

[1]opt. unknown, shows rel. primal/dual-diff. instead   [2]opt. unknown, no reference available   [3]did not terminate within 8 h

Table 1: Quantitative comparison of our solver and the ILP solver Gurobi. We display run time, maximal memory consumption (mem.), relative error of (5) compared to optimum (err.) and TRA score. For Gurobi time of root relaxation (relax.), finding a comparable solution (equ.) and finding an optimal solution (opt.) is reported.

6 Experimental evaluation

We evaluate the performance of our solver on cell and nucleus tracking datasets from different biomedical research projects. As our contributions are exclusively concerned with the optimization of the problem instances, we are not discussing any aspects of potential model mismatch. We used reasonable segmentations obtained by different segmentation routines and selected appropriate costs for the different tracking events without excessive fine-tuning (see supplement). Existing methods that allow for overlapping segmentation hypotheses and cell divisions, like Jug et al. (2014a,b); Kaiser et al. (2018); Schiegg et al. (2014), offload the optimization to off-the-shelf ILP solvers like Gurobi (Gurobi, 2018), which we use as baseline for our comparison.

Datasets and tracking instances. In total, we use 10 problem instances from 3 biomedical data domains, see supplement for additional information.

Drosophila embryo data: We have one problem instance for tracking nuclei in a developing *Drosophila* embryo. The tracking model consists of 252 frames, each containing about 320 detection hypotheses. With about 160 actual objects per time step, we typically observe two conflicting hypotheses per real object in the data.

Flywing data: We use 3 problem instances for tracking membrane-labelled cells in developing *Drosophila* flywing tissue. Two of these consist of 100 frames, each containing about 2 100 detection hypotheses. The third instance is larger, consisting of 245 frames with more than 3 300 detection hypotheses each. In contrast to the embryo data, the segmentation hypotheses in these problem instances are very dense, leading to considerably larger sets of (transitively) conflicting detections.

Cell Tracking Challenge (CTC) data: Finally, we use the publicly available cell tracking datasets *Fluo-C2DL-MSC*, *Fluo-N2DH-GOWT1*, and *PhC-C2DL-PSC* (Ulman et al., 2017) to evaluate our solver. The CTC data consists of 48, 92, and 426 frames, where each frame contains on average 88, 186, and 1 400 detection hypotheses, respectively, with conflict set sizes of about 10, 7, and 3. Each dataset consists of two time-lapse movies, allowing us to generate six tracking instances.

Evaluation criteria. We evaluate the performance using two metrics, namely, total *run time* and overall *memory consumption*. To ensure a fair run time comparison between our solver and Gurobi, we measure not only the time it takes Gurobi to compute the optimal solution, *time (opt.)*, but also the time Gurobi needs to surpass the quality of our best primal solution, *time (equ.)*. Additionally, we compute the *relative error* $\frac{|E(\theta,x)-E(\theta,x^\star)|}{|E(\theta,x^\star)|}$ of our final assignment x with respect to the optimal solution $x^\star$. We finally also compute the *TRA* (Ulman et al., 2017) score, a commonly used tracking scoring function. TRA values are in $[0,1]$, where 1 means that the final assignment is identical to the reference solution, while 0 occurs if the compared solutions have nothing in common. We use TRA to compare our tracking results to the optimal solution obtained with Gurobi. Times show the median result of 5 single-threaded runs on an Intel i7-4770 3.40GHz CPU.

Results and conclusions. In Table 1 we show the results for our solver and Gurobi on all instances. Especially for larger problems, i.e. the *PhC-C2DL-PSC* and *flywing* instances, our solver obtains near-optimal solutions in a fraction of the time (sped-up by factor between 3 and 60). Due to the compact decomposition, our solver consistently requires considerably less memory than Gurobi (up to 17x). Figure 2 shows that our solver converges to small relative errors after only a few iterations. The high quality of our solutions is confirmed by the TRA scores. Compared to existing techniques our solver is scalable and quickly provides high quality solutions even for large-scale real-world problems, so far beeing practically intractable. Therefore it is even applicable in low-latency settings like user-driven proofreading of automated tracking results.

References

Assaf Arbelle and Tammy Riklin Raviv. Microscopy cell segmentation via convolutional LSTM networks. *arXiv:1805.11247*, 2018.

Asad A. Butt and Robert T. Collins. Multi-target tracking by Lagrangian relaxation to min-cost network flow. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

Gregory Castanon and Lucas Finn. Multi-target tracklet stitching through network flows. In *Proceedings of IEEE Aerospace Conference*, 2011.

Nicolas Chenouard, Ihor Smal, Fabrice de Chaumont, Martin Maška, Ivo F. Sbalzarini, Yuanhao Gong, Janick Cardinale, Craig Carthel, Stefano Coraluppi, Mark Winter, Andrew R. Cohen, William J. Godinez, Karl Rohr, Yannis Kalaidzidis, Liang Liang, James Duncan, Hongying Shen, Yingke Xu, Klas E.G. Magnusson, Joakim Jaldén, Helen M. Blau, Perrine Paul-Gilloteaux, Philippe Roudot, Charles Kervrann, François Waharte, Jean-Yves Tinevez, Spencer L. Shorte, Joost Willemse, Katherine Celler, Gilles P. van Wezel, Han-Wei Dan, Yuh-Show Tsai, Carlos Ortiz-de Solorzano, Jean-Christophe Olivo-Marin, and Erik Meijering. Objective comparison of particle tracking methods. *Nature Methods*, 11(3), 2014.

IBM ILOG CPLEX. V12.8: User's manual for CPLEX. 2017.

Monique Guignard. Lagrangean relaxation. *Top*, 11(2), 2003.

Monique Guignard and Siwhan Kim. Lagrangean decomposition: A model yielding stronger lagrangean bounds. *Mathematical programming*, 39(2), 1987a.

Monique Guignard and Siwhan Kim. Lagrangean decomposition for integer programming: theory and applications. *RAIRO-Operations Research*, 21(4), 1987b.

Gurobi. Gurobi optimizer reference manual, 2018. URL http://www.gurobi.com.

Carsten Haubold, Janez Aleš, Steffen Wolf, and Fred A. Hamprecht. A generalized successive shortest paths solver for tracking dividing targets. In *Proceedings of European Conference on Computer Vision*. Springer, 2016.

Umar Iqbal, Anton Milan, and Juergen Gall. Posetrack: Joint multi-person pose estimation and tracking. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

Khuloud Jaqaman, Dinah Loerke, Marcel Mettlen, Hirotaka Kuwata, Sergio Grinstein, Sandra L. Schmid, and Gaudenz Danuser. Robust single-particle tracking in live-cell time-lapse sequences. *Nature Methods*, 5(8), 2008.

Florian Jug, Tobias Pietzsch, Dagmar Kainmüller, Jan Funke, Matthias Kaiser, Erik van Nimwegen, Carsten Rother, and Gene Myers. Optimal joint segmentation and tracking of Escherichia coli in the mother machine. In *Proceedings of Bayesian and Graphical Models for Biomedical Imaging*. Springer, 2014a.

Florian Jug, Tobias Pietzsch, Dagmar Kainmüller, and Gene Myers. Tracking by assignment facilitates data curation. In *Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention, IMIC Workshop*, 2014b.

Florian Jug, Tobias Pietzsch, Stephan Preibisch, and Pavel Tomancak. Bioimage informatics in the context of drosophila research. *Methods*, 68(1), 2014c.

Florian Jug, Evgeny Levinkov, Corinna Blasse, Eugene W. Myers, and Bjoern Andres. Moral lineage tracing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

Matthias Kaiser, Florian Jug, Thomas Julou, Siddharth Deshpande, Thomas Pfohl, Olin K. Silander, Gene Myers, and Erik van Nimwegen. Monitoring single-cell gene regulation under dynamically controllable conditions with integrated microfluidics and software. *Nature Communications*, 9(1), 2018.

Lee Kamentsky, Thouis R. Jones, Adam Fraser, Mark-Anthony Bray, David J. Logan, Katherine L. Madden, Vebjorn Ljosa, Curtis Rueden, Kevin W. Eliceiri, and Anne E. Carpenter. Improved structure, function and compatibility for CellProfiler: modular high-throughput image analysis software. *Bioinformatics*, 27(8), 2011.

Jörg H. Kappes, Bjoern Andres, Fred A. Hamprecht, Christoph Schnörr, Sebastian Nowozin, Dhruv Batra, Sungwoong Kim, Bernhard X. Kausler, Thorben Kröger, Jan Lellmann, Nikos Komodakis, Bogdan Savchynskyy, and Carsten Rother. A comparative study of modern inference techniques for structured discrete energy minimization problems. *International Journal of Computer Vision*, 2015.

Bernhard X. Kausler, Martin Schiegg, Bjoern Andres, Martin Lindner, Ullrich Koethe, Heike Leitte, Jochen Wittbrodt, Lars Hufnagel, and Fred A. Hamprecht. A discrete chain graph model for 3d+ t cell tracking with high misdetection robustness. In *Proceedings of European Conference on Computer Vision*. Springer, 2012.

Vladimir Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10), 2006.

Vladimir Kolmogorov. A new look at reweighted message passing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(5), 2015.

Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Luka Cehovin Zajc, Tomas Vojir, Gustav Hager, Alan Lukezic, Abdelrahman Eldesokey, et al. The visual object tracking vot2017 challenge results. In *Proceedings of IEEE International Conference on Computer Vision*, 2017.

Xinghua Lou, Frederik O. Kaster, Martin S. Lindner, Bernhard X. Kausler, Ullrich Köthe, Burkhard Höckendorf, Jochen Wittbrodt, Heike Jänicke, and Fred A. Hamprecht. Deltr: Digital embryo lineage tree reconstructor. In *Proceedings of IEEE International Symposium on Biomedical Imaging*, 2011.

Klas E.G. Magnusson, Joakim Jaldén, Penney M. Gilbert, and Helen M. Blau. Global linking of cell tracks using the viterbi algorithm. *IEEE Transactions on Medical Imaging*, 34(4), 2015.

Dirk Padfield, Jens Rittscher, and Badrinath Roysam. Coupled minimum-cost flow cell tracking. In *Proceedings of International Conference on Information Processing in Medical Imaging*. Springer, 2009.

Christian Payer, Darko Štern, Thomas Neff, Horst Bischof, and Martin Urschler. Instance segmentation and tracking with cosine embeddings and recurrent hourglass networks. In *Proceedings ofInternational Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018.

Bogdan Savchynskyy. Discrete graphical models — an optimization perspective. *Foundations and Trends in Computer Graphics and Vision*, 11(3-4), 2019.

Martin Schiegg, Philipp Hanslovsky, Bernhard X. Kausler, Lars Hufnagel, and Fred A. Hamprecht. Conservation tracking. In *Proceedings of IEEE International Conference on Computer Vision*, 2013.

Martin Schiegg, Philipp Hanslovsky, Carsten Haubold, Ullrich Köthe, Lars Hufnagel, and Fred A. Hamprecht. Graphical model for joint segmentation and tracking of multiple dividing cells. *Bioinformatics*, 31(6), 2014.

Alexander Shekhovtsov, Christian Reinbacher, Gottfried Graber, and Thomas Pock. Solving dense image matching in real-time using discrete-continuous optimization. In *Proceedings of 21st Computer Vision Winter Workshop*, 2016.

David Sontag, Amir Globerson, and Tommi Jaakkola. Introduction to dual composition for inference. In *Proceedings of Optimization for Machine Learning*. MIT Press, 2011.

Paul Swoboda, Jan Kuske, and Bogdan Savchynskyy. A dual ascent framework for lagrangean decomposition of combinatorial problems. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

Siyu Tang, Bjoern Andres, Miykhaylo Andriluka, and Bernt Schiele. Subgraph decomposition for multi-target tracking. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

Siddharth Tourani, Alexander Shekhovtsov, Carsten Rother, and Bogdan Savchynskyy. MPLP++: Fast, parallel dual block-coordinate ascent for dense graphical models. In *Proceedings of European Conference on Computer Vision*, 2018.

Siddharth Tourani, Alexander Shekhovtsov, Carsten Rother, and Bogdan Savchynskyy. Taxonomy of dual block-coordinate ascent methods for discrete energy minimization. In *Proceedings of Conference on Artificial Intelligence and Statistics*, 2020.

Vladimír Ulman, Martin Maška, Klas E.G. Magnusson, Olaf Ronneberger, Carsten Haubold, Nathalie Harder, Pavel Matula, Petr Matula, David Svoboda, Miroslav Radojevic, Ihor Smal, Karl Rohr, Joakim Jaldén, Helen M. Blau, Oleh Dzyubachyk, Boudewijn Lelieveldt, Pengdong Xiao, Yuexiang Li, Siu-Yeung Cho, Alexandre C. Dufour, Jean-Christophe Olivo-Marin, Constantino C. Reyes-Aldasoro, Jose A. Solis-Lemus, Robert Bensch, Thomas Brox, Johannes Stegmaier, Ralf Mikut, Steffen Wolf, Fred A. Hamprecht, Tiago Esteves, Pedro Quelhas, Ömer Demirel, Lars Malmström, Florian Jug, Pavel Tomancak, Erik Meijering, Arrate Muñoz-Barrutia, Michal Kozubek, and Carlos Ortiz-de Solorzano. An objective comparison of cell-tracking algorithms. *Nature Methods*, 14(12), 2017.

Xinchao Wang, Engin Türetken, Francois Fleuret, and Pascal Fua. Tracking interacting objects optimally using integer programming. In *Proceedings of European Conference on Computer Vision*. Springer, 2014.

Junliang Xing, Haizhou Ai, and Shihong Lao. Multi-object tracking through occlusions by local tracklets filtering and global tracklets association with detection responses. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2009.