Linear-Time Optimal Deadlock Detection for Efficient Scheduling in Multi-Track Railway Networks

Hastyn Doshi, Ayush Tripathi, Keshav Agarwal, Harshad Khadilkar, Shivaram Kalyanakrishnan

Department of Computer Science and Engineering, IIT Bombay {200070025,harshadk,shivaram}@cse.iitb.ac.in, {kshvgrwal,ayush33143314}@gmail.com

Abstract

The railway scheduling problem requires the com-1 putation of an operable timetable that satisfies con-2 straints involving railway infrastructure and re-3 source occupancy times, while minimising aver-4 age delay over a set of events. Since this problem 5 6 is computationally hard, practical solutions typi-7 cally roll out feasible (but suboptimal) schedules one step at a time, by choosing which train to move 8 next in every step. The choices made by such algo-9 rithms are necessarily myopic, and incur the risk of 10 driving the system to a *deadlock*. To escape dead-11 locks, the predominant approach is to stay away 12 from states flagged as *potentially unsafe* by some 13 fast-to-compute rule R. While many choices of R14 guarantee deadlock avoidance, they are suboptimal 15 in the sense of also flagging some safe states as 16 unsafe. In this paper, we revisit the literature on 17 process scheduling and describe a rule R_0 that is 18 19 (i) necessary and sufficient for deadlock detection when the network has at least two tracks in each 20 resource (station / track section). (ii) computable in 21 linear time, and (iii) yields lower delays when com-22 bined with existing scheduling algorithms on both 23 synthetic and real data sets from Indian Railways. 24

25 **1** Introduction

Railway networks around the world form the backbones of 26 national economies. However, due to the constraints imposed 27 by movement on tracks, delays in railways have particularly 28 large domino effects [Goverde, 2010]. In the US, the esti-29 mated cost of delays ranges from 200 USD to more than 1000 30 USD per train-hour [Schlake et al., 2011; Lovett et al., 2015]. 31 Of the total delays in Britain in the 2000s, 40% was composed 32 of primary delay (random events such as train or infrastruc-33 ture faults), and 60% was secondary delay (caused by subse-34 quent congestion) [Preston et al., 2009]. Khadilkar [2017a] 35 observes that in India, 20% of passenger train services are 36 delayed by at least 10 minutes, and that different scheduling 37 strategies have a significant effect on operational efficiency. 38 This fact motivates us to look at scheduling strategies in de-39 tail, with a particular focus on *deadlocks*, which affect both 40 the computation time and schedule efficiency. 41

The railway scheduling problem is a blocking version 42 [Strotmann, 2007; Liu and Kozan, 2009] of the famous job 43 shop scheduling problem (JSSP) [Manne, 1960]. The JSSP is 44 a class of problems in which a number of jobs or processes 45 (in this case, trains) need to be scheduled to pass through a 46 pre-specified sequence of machines or resources (in this case, 47 stations and inter-station track sections) in some "optimal" 48 way. The *blocking* variant implies that once a job is loaded 49 on a machine (train enters a track), it must be fully processed 50 through that step before the unit (track) becomes available for 51 the next job (train). Various versions of optimality exist in the 52 literature, from makespan (time duration from start of the first 53 job to the end of the last job) and average queueing time to 54 average delay. It has been established that the JSSP in several 55 forms is NP-complete [Mascis and Pacciarelli, 2002]. 56

With solving for optimality ruled out, common approaches 57 for railway scheduling proceed by "rolling out" schedules 58 over time [Khadilkar, 2018; Prasad et al., 2021]. Abstractly, 59 such algorithms begin from an initial state in which trains are 60 located in respective resources. Then, at each time step, a 61 train is chosen from those eligible to move, and moved for-62 ward. This process is continued until (possibly) all the trains 63 complete their journeys. Since the decision of which train to 64 move at each step is made myopically, there is a possibility 65 of reaching a deadlock state, from which no further progress 66 is possible unless some train moves backward-which is ex-67 pensive and induces large delays in practice. Figure 1 shows a 68 network with three resources, with one free track in the mid-69 dle resource. If a train heading right is moved into this free 70 track, there is a deadlock. However, deadlock can be avoided 71 by moving a train heading left into the free track. 72

Interestingly, detecting deadlocks in the general JSSP is 73 also NP-complete [Araki *et al.*, 1977; Cocco and Monasson, 2001]. It has consequently been accepted wisdom in 75



Figure 1: Illustration of possible deadlock. The moves corresponding to red (solid) arrows lead to deadlock, while those corresponding to green (dashed) arrows admit a solution without deadlocks.

the railway scheduling community [Törnquist and Persson, 76 2007; Pachl, 2012; Khadilkar, 2017b; Vujanic and Hill, 2022] 77 that deadlock detection for railway scheduling is also NP-78 complete in all its forms. The motivation for our work is a 79 result by Reveliotis et al. [1997] for a variant of JSSP called 80 single-unit resource allocations systems (SURAS), showing 81 that in many reasonable situations and for an arbitrary net-82 work topology, a necessary and sufficient condition for dead-83 lock detection can be computed in polynomial time. 84

We make an explicit connection between the SURAS 85 polynomial-time deadlock detection condition given by Reve-86 liotis et al. [1997] and the railway scheduling problem. In the 87 case where every resource (station and track section) has at 88 least two tracks, we show that this condition can be evaluated 89 for arbitrary network states with a linear complexity in the 90 number of trains (slightly more efficiently than Reveliotis et 91 al. [1997]). Thereafter, we demonstrate the significant benefit 92 of implementing not just sufficient but necessary conditions 93 for deadlock-free movement on real-world railway networks 94 in India. These include single-track as well as multi-track 95 lines, thus showcasing the wide applicability of the algorithm. 96 From an infrastructure perspective, laying two tracks has 97 only modest additional cost compared to laying one track, 98 because the land and utilities infrastructure already exists. 99 Hence the only places where single tracks are typically laid 100 are ones where the traffic is low (in which case sophisticated 101 algorithms are not needed) or where the terrain is tough. In-102 103 deed it is apparent even from data sets used in the academic literature [Pappaterra et al., 2021; Prasad et al., 2021] that 104 nodes in real-world railway lines usually contain more than 105 one track. Even where single tracks exist (we have empir-106 ical results in Section 5), we can handle them so long as a 107 set of feasible moves exist for moving trains to the nearest 108 multi-track resource. At this point, we can drop empty single 109 tracks from the analysis (since they effectively connect two 110 resources, rather than act as independent resources) and the 111 remaining analysis is valid. In the case of scheduling algo-112 rithms, we simply ensure that a train that moves into a single 113 track must be moved on to a multi-track section before an-114 other train move is attempted. 115

After discussing related work in Section 2, we formalise the railway deadlock detection problem in Section 3. In Section 4, we present a novel and conceptually simple interpretation of the detection rule of Reveliotis *et al.* [1997] for multitrack networks. In Section 5, we empirically validate of the utility of this rule in scheduling. We conclude in Section 6.

122 **2 Related Work**

The railway scheduling problem is that of establishing a 123 feasible timetable, given a set of 'services' to be defined from 124 a given origin and destination [Törnquist and Persson, 2007]. 125 Previous literature [Cai and Goh, 1994; Liu and Kozan, 126 2009; Strotmann, 2007] shows that the Job Shop Scheduling 127 Problem (JSSP) and railway scheduling are reducible to 128 each other; their decision variants are both NP-complete. 129 Typically the timetabling (scheduling) problem can be solved 130 at leisure, using a mix of exact and randomised search algo-131 rithms [Higgins et al., 1996; Törnquist and Persson, 2007]. 132

The train rescheduling problem [Sinha et al., 2016] involves 133 a disruption externally imposed on the timetable, from which 134 one must quickly compute a set of recovery actions to return 135 to the original timetable. Delays in this context can be com-136 puted relative to the corresponding event times in the original 137 timetable. Deadlocks are also important in this context since 138 disruptions to the timetable may require changes in the order 139 of train moves, with uncertain implications for operational 140 feasibility. In this paper, we consider both scheduling and 141 rescheduling, with the understanding that computational 142 complexity has more impact on rescheduling (disruption 143 recovery) than on scheduling (timetabling). 144

145

184

Rescheduling and deadlock avoidance in railways. While 146 exact formulations of rescheduling as an optimisation prob-147 lem are available [Higgins et al., 1996; Törnquist and Pers-148 son, 2007], they are not scalable. Practical approaches instead 149 use heuristics [Higgins et al., 1997; Chen et al., 2015] or pre-150 trained policies [Šemrov et al., 2016; Khadilkar, 2018; Prasad 151 et al., 2021] to move trains forward through the network (in 152 a manner analogous to checkers pieces) until they reach their 153 destinations. The order of train movements, their timings, and 154 track allocations vary by the algorithm. However, as illus-155 trated in Figure 1, the risk in all such finite-lookahead meth-156 ods is that of deadlock, or a situation where some or all trains 157 are unable to move forward because of a circular dependence 158 on each other [Pachl, 2012]. 159

Pachl [2012] proposes four conditions which are necessary 160 to create deadlocks. If a scheduling strategy ensures that at 161 least one of these conditions is not met, it is sufficient to avoid 162 deadlocks. Similarly, Mackenzie [2010] proposes sufficient 163 conditions for avoiding deadlock, including the conservative 164 path-to-destination approach. Khadilkar [2017b] proposes 165 the critical-first approach for railway lines, which focuses on 166 prioritising occupants of the most constrained resources in or-167 der to avoid bottlenecks. The condition in this case is to keep 168 moving trains forward up to a point where at least one addi-169 tional free track is available for other trains to pass. Vujanic 170 and Hill [2022] make this more concrete by defining the no-171 tion of a safe state as one where all nodes (resources) have an 172 unoccupied slot. If initialised from a compliant initial state, 173 this procedure takes polynomial time for scheduling. 174

The basis for all these studies is that the NP-completeness 175 of the general deadlock detection problem makes it hard 176 to detect deadlocks in instances encountered in practice. 177 Therefore, deadlock detection is typically performed by 178 rules that provably detect deadlocks, but might also flag 179 false positives. The novelty of our paper is in identifying 180 the applicability of an optimal polynomial-time deadlock 181 detection algorithm in JSSP to the railway context (the notion 182 of optimality is formally defined below). 183

Deadlock avoidance in JSSP. We shall first define the various terms used in this paper. The problem of evaluating an arbitrary state of the railway network, for the presence/absence of present/future deadlock, is *deadlock detection.* Any subsequent scheduling policies that reduce the probability of deadlock (but not eliminate it) are called *deadlock avoiding policies*, while scheduling policies that guaran-191

tee the absence of deadlock are called *deadlock free policies*. 192 Similar to the railway scheduling case, it is well known 193 that optimal deadlock detection in JSSP is also NP-complete. 194 Araki et al. [1977] consider the question: "given a state S, 195 is S safe?" They reduce the 3-SAT problem [Cocco and 196 Monasson, 2001] to optimal deadlock detection in JSSP, 197 thereby proving the latter to be NP-complete. Fanti et al. 198 [1997] derive necessary and sufficient conditions for dead-199 lock in production systems with resource sharing, and then 200 propose a 'restriction policy' that is tractable and provably 201 202 correct (in the sense of sufficiency). Gold [1978] considered the question from a more practical perspective, examining 203 under what restrictions on state S one can detect deadlock in 204 polynomial time. They derived some conditions under which 205 deadlock detection can be solved in polynomial time. 206

207

Optimal deadlock detection. Previously published studies 208 also consider two forms of optimality in the present context. 209 The first definition of optimal implies the minimisation of de-210 lays in the schedule with respect to a reference timetable, or 211 the minimisation of the makespan of the schedule if no ref-212 erence timetable is available [Törnquist and Persson, 2007]. 213 The second definition of optimality [Reveliotis et al., 1997] 214 refers to the removal of unsafe transitions from the current 215 state, with the objective of identifying the smallest (hence 216 optimal) set of unsafe transitions that ensures the absence of 217 present or future deadlocks. In this paper, by optimal rule we 218 refer to the second definition: to a rule that characterises nec-219 essary and sufficient conditions of states or transitions to be 220 safe, and hence can be used for deadlock-free scheduling. 221

Reveliotis et al. [1997] develop necessary and sufficient 222 conditions for deadlock prevention in "single-unit sequential 223 resource allocation systems" (SURAS). They show that dead-224 lock detection in polynomial-time is possible in the special 225 case where every resource in the system has a minimum ca-226 pacity of 2 units. If the number of resources is m and C227 is the maximum capacity among these nodes, their detection 228 condition has $O(m^2 \overline{C})$ complexity. The intuition behind this 229 number is that a search algorithm makes m passes through the 230 set of m resources, eliminating one eligible resource in each 231 pass. Our observation is that this result applies to the case of 232 arbitrary railway network topologies (branching and straight 233 lines) as long as there are at least two tracks in each node 234 (in railway terminology, at stations and inter-station track 235 sections). Furthermore, (i) the result can actually be imple-236 mented in linear (and not quadratic) complexity, and (ii) we 237 can handle single-resource nodes under reasonable assump-238 tions, as explained in Section 1. 239

240 3 Deadlock Detection

In this section, we specify the problem of (optimal) deadlock detection. We begin from the broader context of railway
scheduling, within which this problem arises.

244 3.1 Railway scheduling problem

Railway infrastructure. A railway network is made up of
a number of *resources*, each containing some number of
parallel tracks running from one end of the resource to the

other. Tracks admit traffic in both directions. Stations (where 248 trains may halt) as well as the inter-station track sections 249 between them (where trains do not have scheduled halts) are 250 modeled as resources. A resource connects to other resources 251 through one of its ends. Typically, terminal resources have 252 all their connections only from one end, but in general we 253 could have cycles in the network topology. Figure 2 shows an 254 illustrative railway network with branching and a cycle; the 255 example in Figure 1 has a linear topology (often called a *line*). 256

Desired schedule. The dynamic aspect of the scheduling problem arises from the movement of a set of trains through resources. The target is to meet a desired schedule S_{desired} , 260 which may be represented as a set of N events: 261

257

281

$$S_{\text{desired}} = \{e[i], 1 \le i \le N\}, \text{where}$$
$$e[i] = (train[i], start_res[i], next_res[i], time[i]).$$

Event e[i] specifies that train train[i] must be moved from 262 resource $start_res[i]$ to the adjoining resource $next_res[i]$ at 263 time time[i]. Now, it may not be possible to execute S_{desired} , 264 due to constraints imposed by the railway infrastructure. For 265 instance, if three events all mean to push trains into the same 266 resource at the same time, but this resource only has two free 267 tracks, then at least one of the events will have to be *delayed*. 268 The goal of scheduling is to compute an *operable* schedule 269 $S_{operable}$ that is feasible to execute, but at the expense of de-270 laying a subset of events in S_{desired} . For each event e[i], the 271 operable schedule has a replacement $\overline{e}[i]$ with a new time 272 $\overline{time}[i] > time[i]:$ 273

$$S_{\text{operable}} = \{\overline{e}[i], 1 \le i \le N\}, \text{where}$$

$$\overline{e}[i] = (train[i], start_res[i], next_res[i], \overline{time}[i]).$$

Formally, the objective function to be minimised while computing S_{operable} is the average departure delay 275

$$ADD = \frac{1}{N} \sum_{i=1}^{N} (\overline{time}[i] - time[i]). \tag{1}$$

Since the problem of computing an operable schedule that minimises ADD is NP-hard [Mascis and Pacciarelli, 2002], 277 one practical alternative is to roll out schedules over time, 278 ensuring operabilty, while making greedy choices to reduce 279 delays [Khadilkar, 2018; Prasad *et al.*, 2021]. 280

3.2 Roll-out algorithms

A roll-out algorithm executes the set of events $\{e[i], 1 \le i \le 282 N\}$ one by one. The algorithm begins with a counter τ set 283 to the earliest event time, with state s_{τ} associating each train 284 with its initial resource. An event *i* is said to be executed (and 285 inserted into S_{operable}) when the algorithm sets $\overline{time}[i]$. 286

Figure 2: Example of a network topology. Two trains are shown.

At each counter value τ , the algorithm compiles the list of 287 events that are *eligible*: these are events $e[i], 1 \le i \le N$, such 288 that (i) train[i] is in $start_res[i]$ in state s_{τ} ; (ii) there is a free 289 track in resource $next_res[i]$ in s_{τ} ; and (iii) $time[i] \geq \tau$. 290 If, indeed, there are eligible events, one of these events i is 291 selected and executed by setting $time[i] = \tau$. The updated 292 event $\overline{e}[i]$ is moved into S_{operable} , and e[i] is no longer eligible. 293 As long as there are eligible events at τ , these are repeatedly 294 executed, until there are no eligible events at τ ; in this case 295 au is incremented and the procedure continues. Since train 296 journeys are a sequence of contiguous resources, any train 297 can be in at most one eligible event at any time step. Hence, 298 it is sometimes convenient to view the set of eligible events at 299 τ as the set of trains that are eligible to be moved at τ . 300

By construction, the set of events that have already been 301 executed by a roll-out algorithm have no internal conflicts. 302 Hence, if all N events in S_{desired} get executed, we are guar-303 anteed an operable schedule $S_{\rm operable}.$ However, the ADD of 304 S_{operable} depends on the delays introduced while executing the 305 events. The choice of which event among the eligible ones 306 to execute at any step also has the long-term consequence 307 of which events become eligible in subsequent time steps. 308 By and large, roll-out algorithms make this choice greed-309 310 ily [Khadilkar, 2018; Prasad et al., 2021]. An unfortunate consequence is the possibility of a deadlock, wherein there 311 remain events to execute, but these cannot become eligible at 312 the current counter value τ or anything larger. 313

314 3.3 Deadlock detection problem

Abstractly, the progress of a roll-out algorithm for generat-315 ing a schedule can be viewed as a sequence of state transi-316 tions. The background data from the problem instance, which 317 guide and constrain these transitions, are (1) the set of re-318 sources U; (2) resource capacities encoded by $\overline{C}: U \to \mathbb{N}$; 319 (3) the set of trains T; and (4) the set of train journeys $D = \{(t, u^1, u^2, \dots, u^{l_t}), t \in T\}$. In D, each journey $(t, u^1, u^2, \dots, u^{l_t})$ contains a train $t \in T$ and the identities 320 321 322 of some $l_t \geq 1$ resources through which t must pass in se-323 quence. Exact event times are not needed for deadlock de-324 tection. As motivated in Section 1, we make the following 325 "multi-track" assumption while devising and analysing our 326 algorithm, which is presented in Section 4. 327

328 Assumption 1. For all $u \in U$, $\overline{C}(u) \ge 2$.

However, the problem statement presented below does not depend on this assumption.

331

States, actions, transitions. Each state s in our system con-332 tains a subset of trains $T'_s \subseteq T$ that are yet to complete 333 their journeys. In s, each train $t \in T'_s$ is in some resource 334 $u \in U$. Hence, a state can be represented as a set of pairs 335 $(t, u) \in T \times U$. Let S denote the set of all states. The de-336 sired terminal state $s_{\top} \in S$ is the one in which all trains have 337 reached their destinations. Destinations typically connect to 338 "yards" with effectively infinite capacity, so trains do not oc-339 cupy regular tracks at their destinations. Hence $s_{\top} = \emptyset$. 340

A useful quantity to associate with each state $s \in S$ is its "potential" $\phi(s)$, which we define to be the sum of the distances of the trains present in s to their respective termini. Concretely, suppose $s = \{(t_i, u_i), 1 \le i \le m\}$, where the 344 remaining sequence of resources for train t_i to visit after departing u_i is $u_i^1, u_i^2, \ldots, u_i^{\ell_i}$ for some $\ell_i \ge 1$. Then we have 346 $\phi(s) = \sum_{i=1}^m (1 + \ell_i)$. Observe that $\phi(s_{\top}) = 0$. 347

The set of *actions* available from state $s \in S$ is denoted A(s). Each action $a \in A(s)$ corresponds to moving some train from its current resource to the next one on its journey. Naturally, only moves corresponding to eligible events are present as actions in A(s).

When an action from A(s) is performed on state $s \in S$, 353 we denote the resulting state s + a. Suppose action $a \in A(s)$ 354 moves train t in s, where t is in resource u, to its next resource 355 u'. If u' is the terminal resource for t, then $s+a = s \setminus \{(t, u)\}$; 356 otherwise $s + a = (s \setminus \{(t, u)\}) \cup \{(t, u')\}$. Notice that when 357 an action from A(s) is performed on s, progress is achieved 358 in the sense that $\phi(s+a) = \phi(s) - 1$. Since trains cannot 359 move backwards, this progress cannot be undone. However, 360 as we see next, an action may lead to a state from which 361 further progress is not possible. 362

363

381

Safe and unsafe states. By definition, the desirable terminal state $s_{\top} = \emptyset$ is a safe state; so also is every state for which there exists a sequence of actions to reach s_{\top} . We may write down recursively: for $s \in S$,

 $SAFE(s) \iff (s = s_{\top}) \lor (\exists a \in A(s) : SAFE(s + a)).$

This recursive definition gives rise to a straightforward procedure to *compute* SAFE(s), since any state s + a in the RHS has a potential value $\phi(\cdot)$ one unit lower than s. However, there is branching by a factor of |A(s)|, implying exponential complexity for a naive implementation. Our main observation, described in the next section, is that SAFE(s) can be computed in time that is only linear in the size of s. 370

An unsafe state is a state that is not safe. A deadlocked 371 state is a state containing trains, but on which no valid action 372 can be performed. For $s \in S$, 373

$$UNSAFE(s) \iff \neg SAFE(s);$$

Deadlock(s) $\iff (s \neq s_{\top}) \land (A(s) = \emptyset)$

Since our system contains a finite number of trains, and their journeys are also finite, it follows that $\phi(\cdot)$ has a finite upper bound. Since $\phi(\cdot)$ decreases by 1 unit after each action, the length of any action sequence starting from any initial state s_0 is also guaranteed to be finite. It follows that if s is unsafe, then any sequence of actions starting from s will lead to a deadlocked state, from which no further actions are available. 380

Computational problem. We require a procedure that can 382 efficiently identify whether a given state $s \in S$ is safe or not. 383 It is convenient to view this procedure as a rule or proposition 384 R(s), which evaluates to a boolean value. R(s) may depend 385 both on s and on the journey details of the trains in D, but 386 must be efficient to compute. Several "sufficient" rules R387 from the literature guarantee that $R(s) \implies SAFE(s)$. We 388 require a "necessary and sufficient" rule, also called an opti-389 mal rule, which satisfies $R(s) \iff SAFE(s)$. 390

As described in Section 1, optimal rules are computationally hard on unrestricted problem instances. On the other hand, we show next that if Assumption 1 is satisfied, then an optimal rule can be implemented in only linear time. 391

395 4 Linear-Time Algorithm

404

434

The algorithm presented here is due to Reveliotis et al. 396 [1997], who proposed it in the context of resource allocation 397 and implemented it with quadratic complexity. We describe 398 this algorithm from the perspective of deadlock detection 399 in railway networks, using the vocabulary introduced in 400 Section 3. We provide a concise proof of correctness based 401 on a graph-theoretic model, and also show that a linear-time 402 implementation is possible. 403

Next-stop graph. The main data structure involved in specifying R_0 is a directed graph constructed based on input state s. The construction also requires the resource capacities encoded by \overline{C} and the set of train journeys D. We denote this directed graph $G_s = (V, E, C)$, where V is the set of vertices, E the set of edges, and $C : V \rightarrow \{\text{red, black}\}$ a function that associates a colour with each vertex.

Recall that T_s is the set of trains in s. A resource $u \in U$ is 412 a vertex $v \in V$ in G_s if and only if u is the current resource 413 for some train t in s, or it is the next resource for some train 414 t in s (as specified in t's journey in D). The colour C(v) of a 415 vertex $v \in V$ is red if the corresponding resource has at least 416 one free track (that is, the number of trains in this resource 417 is strictly smaller than the capacity). Fully-occupied vertices 418 are coloured black. Each train t in state s gives rise to an edge 419 from the vertex of its current resource u to the vertex of its 420 next resource u'. Hence, each edge $e \in E$ corresponds to one 421 or more trains in s. Notice that every vertex $v \in V$ must have 422 at least one edge, either incoming or outgoing (but possibly 423 one or more of each type). 424

Surprisingly, one does not need to access the extended 425 journeys of trains in s in order to construct G_s : one only 426 needs the trains' current and next resources. For this reason, 427 we may refer to G_s as the "next-stop graph" of s. Notice that the number of edges in G_s is at most the number of trains in 428 429 s: that is, $|E| \leq |T|$. Clearly G_s does not exceed the size 430 of s or of the journey data D beyond a constant factor, as 431 both s and D use $\Omega(|T|)$ space. Even so, G_s provides all the 432 information required for optimal deadlock detection. 433

435 **Optimal rule.** Our rule R_0 has an intuitive form.

436 Definition 2. For $s \in S$, $R_0(s)$ is the proposition: "In G_s , 437 every black vertex has a directed path to some red vertex."

438 We formally show that under the multi-track assumption, 439 R_0 is an optimal deadlock detection rule.

Theorem 3. If the problem instance satisfies Assumption 1, then for $s \in S$, $R_0(s) \iff SAFE(s)$.

Proof. Let $G_s = (V, E, C)$. We prove the theorem by induc-442 tion on $\phi(s)$. As base case, consider arbitrary $s \in S$ for which 443 $\phi(s) = 1$. If so, there is exactly one train t in the network, in a 444 resource that connects to t's terminus. Clearly s is safe since 445 t can be moved into its terminus (thus s transitions into s_{\top}). 446 Also notice that in this case, G_s comprises exactly two ver-447 tices $r_1, r_2 \in V$, with an outgoing edge from r_1 to r_2 . Since 448 each resource has at least two tracks, r_1 and r_2 must both be 449 red, making $R_0(s)$ trivially true. In short, when $\phi(s) = 1$, 450 $R_0(s)$ and SAFE(s) are both true, and thereby equivalent. 451

Our induction hypothesis is that for some integer $m \ge 1$, 452 $R_0(s) \iff \text{SAFE}(s)$ for all $s \in S$ having $\phi(s) = m$. 453 Now consider a state $s \in S$ for which $\phi(s) = m + 1$. We 454 separately prove the two implications in the theorem. 455

456

460

1. Proof of $R_0(s) \implies$ SAFE(s). Suppose that $R_0(s)$ is 457 true: that is, in G_s , every black vertex has a directed path to 458 some red vertex. We consider two complementary subcases. 458

1.1. Suppose G_s contains some red vertex $r \in V$ with no 461 outgoing edges (Figure 3a). As in our base case, r must 462 contain an incoming edge from some vertex $v \in V$. Notice 463 that r has two or more empty tracks. Hence, we can move 464 a train t from v to r to obtain a state s' with $\phi(s') = m$. 465 If there are any incoming edges into v in s, or v has trains 466 other than t in s, then v would also be a vertex in s', but 467 now coloured red. Otherwise v would not be a vertex in $G_{s'}$. 468 Depending on t's next stop from s', v could get a new edge 469 to an existing vertex or a new red vertex in $G_{s'}$. Regardless, 470 notice that if any black vertex had a path to a red vertex in 471 G_s , it would continue to have a path to a red vertex in $G_{s'}$. 472 Hence, if $R_0(s)$ is true, then $R_0(s')$ is also true. By the 473 induction hypothesis, s' is safe, and hence s is also safe. 474 475

1.2. The second subcase is that every red vertex in G_s has 476 an outgoing edge (Figure 3b); every black vertex in G_s will 477 anyway have at least one outgoing edge. In this case, we 478 consider a subgraph G' of G_s (Figure 3c), which differs only 479 in the set of edges. Indeed let G' = (V, E', C) so that (i) 480 each vertex $v \in V$ has exactly one outgoing edge in E', and 481 (ii) each black vertex in G' has a directed path to some red 482 vertex in G'. A natural way to construct E' would be to first 483 fix some outgoing edge to a red vertex from all black vertices 484 having such an edge in E, then to fix an outgoing edge to one 485 of those black vertices from all other black vertices having 486 such an edge in E, and proceeding similarly until all black 487 vertices have an outgoing edge. Thereafter each red vertex 488 can be given an arbitrary outgoing edge from E. 489

By its definition, G' cannot have a directed cycle with only black vertices (since that would imply that those vertices do 491



Figure 3: Graphs illustrating cases in proof of Theorem 3.

not have a directed path to some red vertex in G, hence invalidating $R_0(s)$). Also, since it has an outgoing edge for each vertex, G' must contain a directed cycle. In summary, we infer that G' must contain a directed cycle with at least one red vertex. Indeed let such a cycle C contain the sequence of vertices r, v_1, v_2, \ldots, v_m for some $m \ge 1$, where r is a red vertex. We are indifferent to the colours of v_1, v_2, \ldots, v_m .

Now consider the action of moving a train t from v_m to r, 499 leading to next state s'. The set of vertices in $G_{s'}$ remains 500 identical to that of s, except that $G_{s'}$ may get a new red 501 vertex that is the next stop for t in s'. The set of edges in 502 $G_{s'}$ is also identical to that of G_s , except for (i) the possible 503 loss of the single edge from v_m to r in case t was the only 504 train having that transition in s, and (ii) the possible gain of 505 a new outgoing edge from r. Regardless, v_m is necessarily 506 a red vertex in $G_{s'}$, and moreover, any black vertex that had 507 a directed path to r in s must have a path to v_m in s' (since 508 r has a path through the sequence of vertices in C to v_m). 509 Directed paths not involving C in G_s remain the same in 510 $G_{s'}$. Hence, we conclude that if $R_0(s)$ is true, then $R_0(s')$ 511 is also true. Since $\phi(s') = m$, we obtain from the induction 512 hypothesis that s' is safe. Since we can go from s to s' by 513 moving t, we observe that s must also be safe. 514

2. Proof of SAFE(s) $\implies R_0(s)$. Suppose $R_0(s)$ is not true: 516 that is, there exists a black vertex $b \in V$ with no directed path 517 to any red vertex in V (Figure 3d). b is fully occupied in s, 518 and so it must have outgoing edges in G_s . Since these edges 519 to do not lead to a red vertex in G_s , we surmise that there 520 exists a finite subset of black vertices $B \subseteq V$ of size at least 521 two such that for each vertex in B, all outgoing edges lead 522 to other vertices in B. Since every vertex $b' \in B$ is fully-523 occupied, no train can be moved out from or moved into any 524 vertex in B. If s is already in deadlock, by definition it is 525 unsafe. On the other hand, after any possible move of any 526 train outside of B in s to reach s', B remains a finite set of 527 black vertices in $G_{s'}$, with no directed path to any red vertex. 528 Since $\phi(s') = m$, the induction hypothesis gives us that s' is 529 unsafe, and hence s is unsafe. \square 530

Although R_0 is essentially a rephrasing of the rule given 531 by Reveliotis et al. [1997] for SURAS, its interpretation in 532 terms of a next-stop graph G_s is novel. The algorithm given 533 by Reveliotis et al. [1997] eliminates one node in each pass 534 through the set of resources, hence takes time that is quadratic 535 in the number of resources. On the other hand, it is easy 536 to see that $R_0(s)$ can be computed in time that is only lin-537 ear in the number of edges in G_s , which is generally much 538 smaller than $|U|^2$. The pseudocode in Figure 4 is for an im-539 plementation of R_0 by a standard search algorithm [Russell 540 and Norvig, 2022, see Chapter 3], taking O(|E|) operations 541 for input $\tilde{G}_s = (V, E, C)$. 542

543 **5** Experimental validation

515

We compare our proposed rule R_0 against other deadlock avoidance algorithms on real railway schedules as well as synthetically-generated ones. Data and code to reproduce the results reported in this section are available at https: //github.com/Hastyn/Linear-Time-Deadlock-Detection/.



Figure 4: $R_0(s)$ implementation with input $G_s = (V, E, C)$. The frontier may be implemented either as a stack or as a queue.

549

5.1 Data description

We use published timetables and infrastructure information 550 from three portions of the Indian Railway network. These 551 portions are from Ajmer (northwest India), Kanpur (north In-552 dia) and Konkan (west coast). Of these, Ajmer and Kanpur 553 are predominantly multi-track at stations as well as the con-554 necting sections between stations, while Konkan has multi-555 track stations but mostly single-track connecting sections (we 556 handle this scenario by moving trains from one station to the 557 next without stopping in the connecting sections, as outlined 558 earlier). Details from these networks are summarised in Table 559 1. In addition to the real data sets, we also generate 8 hypo-560 thetical lines and branching networks to have better control on 561 the characteristics. HYP-1 is a small toy dataset, HYP-2 and 562 HYP-3 share the same infrastructure, but HYP-3 has twice 563 the number of trains of HYP-2. HYP-4 and HYP-5 simulate 564

Sn.	Tns.	Time	Events	Con.	Density
		Span		Sec.	
52	444	5.5 day	13129	51	0.016
27	190	1.4 day	3858	26	0.036
59	85	2.2 day	2709	58	0.007
5	8	1.8 hrs	40	4	0.041
11	60	2.0 day	660	10	0.011
11	120	2.1 day	1320	10	0.021
4	350	0.8 hrs	1290	3	3.839
3	200	0.5 hrs	4766	2	31.773
6	6	0.8 hrs	23	5	0.044
26	500	3.0 hrs	4537	27	0.476
22	100	4.8 hrs	1046	21	0.084
	Sn. 52 27 59 5 11 14 3 6 26 22	Sn. Tns. 52 444 27 190 59 85 5 8 11 60 11 120 4 350 3 200 6 6 26 500 22 100	Sn. Tns. Time Span 52 444 5.5 day 27 190 1.4 day 59 85 2.2 day 5 8 1.8 hrs 11 60 2.0 day 11 120 2.1 day 4 350 0.8 hrs 3 200 0.5 hrs 6 6 0.8 hrs 26 500 3.0 hrs 22 100 4.8 hrs	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

Table 1: Data set description, giving number of stations (Sn.), trains (Tns.), span of the reference timetable, total number of departure events, number of connecting sections (Con. Sec.), and the traffic density in events per resource per minute. Note that the number of resources is the sum of stations and connecting sections.

very high traffic networks with only four stations but a large 565 number of trains. HYP-6, HYP-7, and HYP-8 are branch-566 ing networks with HYP-7 having high traffic. In Table 1, the 567 density reported is the total number of events occurring per 568 minute per resource in the network. 569

5.2 **Comparison with baselines** 570

The goal of our experiments is to test the reduction in ADD 571 (defined in (1)) as a result of dividing the action set in any 572 given state optimally into safe and unsafe labels. Considering 573 this to be a binary classification problem, we pick one base-574 line rule R_q which allows false negatives (marking an actu-575 ally unsafe state as safe) and one rule R_c which allows false 576 positives (marking an actually safe state as unsafe). We note 577 that the rules R_q and R_c correspond to 'greedy' [Prasad et 578 al., 2021] and 'critical first' [Khadilkar, 2017b] in their orig-579 inal forms. Briefly, the greedy algorithm marks a train move-580 ment as safe if it has at least two feasible forward moves. 581 Critical-first is a sufficient condition for deadlock free move-582 ment, which allows a train to move ahead as long as it only 583 stops in a resource where at least one additional track is free. 584 For every compounded set of primitive moves, this results in 585 a state where every resource has at least one free track. 586

The critical first algorithm further provides a ranking order 587 when multiple train movements are marked as safe, based on 588 the number of presently free tracks in the resource the train 589 is currently occupying (*criticality* of the node). We use this 590 logic to rank train moves among the set marked *safe* by each 591 rule. The resulting ADD for all problem instances and algo-592 rithms is summarised in Table 2. We emphasise that only the 593 safe action masking (by using R_0 , R_q , or R_c) differs among 594 the algorithms, and the remaining scheduling/rescheduling 595 policy is the same. In order to generate statistical results, 596 we generate perturbed versions of each instance by moving 597 598 the entire journey of each train in the reference timetable by an amount picked uniformly at random in [-30, 30] minutes. 599 Refer to Prasad et al. [2021] for details of the perturbations. 600

From Table 2, we first confirm that both R_c (sufficient) and 601 R_0 (necessary and sufficient) conditions result in deadlock-602 free schedules for all instances. Second, R_0 performs signif-603 icantly better than R_c in all instances, in terms of schedule 604 efficiency. This demonstrates the advantage of employing an 605 optimal deadlock detection condition. R_g outperforms R_0 in 606

Instance	Critical First (R_c)	Greedy (R_g)	R_0
Ajmer	4.76 ± 0.07	$4.19 {\pm} 0.08$	4.12 ±0.09
Kanpur	1.35 ± 0.07	3.57 ± 0.09	1.29 ±0.05
Konkan	60.33±0.69	42.22 ±0.59	42.60±0.51
HYP-1	19.30±1.74	16.22 ±1.29	16.49±1.30
HYP-2	6.33±0.33	4.29 ±0.21	4.31±0.21
HYP-3	7.02 ± 0.91	5.01 ± 0.16	0.83 ±0.07
HYP-4	1773.73±12.79	deadlock	1170.11 ±2.46
HYP-5	654.9±17.26	568.07±2.71	524.23 ±2.39
HYP-6	12.07±1.43	deadlock	6.42 ±1.20
HYP-7	1487.88 ± 22.11	deadlock	1228.30 ±1.61
HYP-8	776.96 ± 27.74	$215.74{\pm}2.04$	169.48±1.90

Table 2: ADD values in minutes with their standard error averaged over 10 perturbed versions of the reference timetables.

three instances, all of which have low traffic density (see Ta-607 ble 1). However, the ADD for R_0 is competitive even in these 608 instances. On the other hand, R_q deadlocks in instances with 609 high traffic density, and hence in general would not be a suit-610 able choice for real-time rescheduling. 611

Policy Improvement 5.3

As a second experiment, we consider the effect of optimal 613 deadlock detection on the efficiency of resulting schedules, 614 by performing policy improvement using roll-outs [Tesauro 615 and Galperin, 1996; Agarwal, 2022]. Under policy improve-616 ment, a base schedule is progressively improved by updating 617 each action to one that minimises delay when followed by a 618 roll-out policy. In Table 3, we start with the schedule pro-619 duced by R_0 in Table 2 for all three algorithms (for a fair 620 comparison). For every decision taken in the sequence, we 621 roll out the individual trajectories for all alternative actions 622 which are also marked as safe by the relevant rule. We choose 623 the schedule with the least ADD out of these results, move to 624 the next decision in the sequence, and repeat. The results in 625 Table 3 show that the rollouts using R_0 (which provides the 626 maximal set of feasible actions) are predominantly more ef-627 fective than those using R_c and R_q . In some cases, R_c and R_q 628 are unable to improve on the baseline schedule given by R_0 , 629 while R_0 results in improvement over Table 2 in all instances. 630

Conclusion 6

In this paper, we show that in contrast to the accepted 632 characterisation of railway scheduling in the literature, a 633 polynomial-time deadlock detection method from the re-634 source allocation literature applies to a large class of 635 (re)scheduling problems. Our version of the implementation 636 is in fact linear-time for arbitrary network topology, so long as 637 each resource (station or connecting section) has at least two 638 tracks. Further, we show that under a mild assumption (avail-639 ability of a sequence of moves to bring all trains in single-640 track resources to a multi-track resource), we can also handle 641 scheduling in the single-track scenario. Our empirical results 642 show that using an *optimal* deadlock detection strategy sig-643 nificantly improves scheduling efficiency, in addition to pro-644 viding feasibility guarantees. One important open question 645 for the future is to evaluate the usefulness of optimal action 646 masking while training data-driven scheduling policies. 647

Instance	R_g Rollout	R_c Rollout	R_0 Rollout
Ajmer	3.43 ± 0.04	3.79 ± 0.05	$\textbf{3.43}\pm0.07$
Kanpur	1.29 ± 0.05	1.13 ± 0.05	$\textbf{1.09}\pm0.05$
Konkan	39.67 ± 0.57	42.59 ± 0.51	40.24 ± 0.61
HYP-1	15.75 ± 1.26	16.18 ± 1.28	15.74 ± 1.26
HYP-2	3.99 ± 0.24	4.19 ± 0.18	4.08 ± 0.20
HYP-3	0.83 ± 0.07	0.82 ± 0.07	$\textbf{0.74} \pm 0.06$
HYP-4	1170.11 ± 2.46	1170.11 ± 2.46	1166.26 ± 2.49
HYP-5	524.23 ± 2.39	524.23 ± 2.39	517.23 ± 2.09
HYP-6	6.35 ± 1.22	5.90 ± 1.01	4.20 ± 0.62
HYP-7	1215.73 ± 1.57	1228.3 ± 1.61	1225.99 ± 1.36
HYP-8	169.48 ± 1.90	169.48 ± 1.90	165.79 ± 2.42

Table 3: Policy improvement starting with the baseline schedule produced by R_0 in (the last column of) Table 2.

631

612

648 Acknowledgements

- We thank Spyros Reveliotis for sharing full versions of relevant literature that was otherwise not available. We thank Ab-
- hiram Ranade from IIT Bombay for informative discussions,
- and Shripad Salsingikar from TCS for providing railway data.

653 **References**

- Keshav Agarwal. Real-time railway scheduling. Master'sthesis, Indian Institute of Technology Bombay, 2022.
- ⁶⁵⁶ Toshiro Araki, Yuji Sugiyama, Tadao Kasami, and Jun Okui.
- 657 Complexity of the deadlock avoidance problem. In 2nd
 658 IBM Symp. on Mathematical Foundations of Computer Sci 659 ence, pages 229–257, 1977.
- ⁶⁶⁰ X Cai and C Goh. A fast heuristic for the train scheduling ⁶⁶¹ problem. *Computers & Op. Res.*, 21(5):499–510, 1994.
- L Chen, C Roberts, F Schmid, and E Stewart. Modeling and solving real-time train rescheduling problems in railway
- bottleneck sections. *IEEE Trans. on Intelligent Transporta- tion Systems*, 16(4):1896–1904, 2015.
- Simona Cocco and Rémi Monasson. Trajectories in phase diagrams, growth processes, and computational complexity:
 How search algorithms solve the 3-satisfiability problem. *Physical review letters*, 86(8):1654, 2001.
- 669 *Fhysical Teview tellers*, 80(8).1034, 2001.
- Maria Pia Fanti, Bruno Maione, Saverio Mascolo, and
 A Turchiano. Event-based feedback control for deadlock
 avoidance in flexible production systems. *IEEE Transac- tions on Robotics and Automation*, 13(3):347–363, 1997.
- E Mark Gold. Deadlock prediction: Easy and difficult cases.
 SIAM Journal on Computing, 7(3):320–336, 1978.
- Rob MP Goverde. A delay propagation algorithm for large scale railway traffic networks. *Transportation Research Part C: Emerging Technologies*, 18(3):269–287, 2010.
- A Higgins, E Kozan, and L Ferreira. Optimal scheduling of
 trains on a single line track. *Transportation Research Part*B, 30(2):147–161, 1996.
- A Higgins, E Kozan, and L Ferreira. Heuristic techniques for
 single line train scheduling. *Journal of Heuristics*, 3(1):43–
 62, 1997.
- Harshad Khadilkar. Data-enabled stochastic modeling for
 evaluating schedule robustness of railway networks. *Trans- portation Science*, 51(4):1161–1176, 2017.
- Harshad Khadilkar. Scheduling of vehicle movement
 in resource-constrained transportation networks using a
 capacity-aware heuristic. In 2017 American Control Conference (ACC), pages 5617–5622. IEEE, 2017.
- Harshad Khadilkar. A scalable reinforcement learning algorithm for scheduling railway lines. *IEEE Transactions on Intelligent Transportation Systems*, 20(2):727–736, 2018.
- S Liu and E Kozan. Scheduling trains as a blocking parallel machine job shop scheduling problem. *Computers & Op- erations Research*, 36(10):2840–2852, 2009.
- Alexander H Lovett, C Tyler Dick, and Christopher PL
 Barkan. Determining freight train delay costs on railroad
 lines in north america. *Proceedings of Rail Tokyo*, 2015.

- S. Mackenzie. *Train scheduling on long haul railway corridors*. PhD thesis, University of South Australia, 2010. 702
- Alan S Manne. On the job-shop scheduling problem. Operations research, 8(2):219–223, 1960. 704
- Alessandro Mascis and Dario Pacciarelli. Job-shop scheduling with blocking and no-wait constraints. *European Journal of Operational Research*, 143(3):498–517, 2002. 707
- Jörn Pachl. Deadlock avoidance in railroad operations simulations. In *PROMET Traffic & Transportation*, pages 359– 369, 2012. 710
- Mauro José Pappaterra, Francesco Flammini, Valeria Vittorini, and Nikola Bešinović. A systematic review of artificial intelligence public datasets for railway applications. *Infrastructures*, 6(10):136, 2021. 714
- Rohit Prasad, Harshad Khadilkar, and Shivaram Kalyanakr ishnan. Optimising a real-time scheduler for Indian rail way lines by policy search. In 2021 Seventh Indian Control
 Conference (ICC), pages 75–80. IEEE, 2021.
- John Preston, Graham Wall, Richard Batley, J Nicolás Ibáñez, and Jeremy Shires. Impact of delays on passenger train services: evidence from great britain. *Transportation research record*, 2117(1):14–23, 2009. 722
- Spiridon A Reveliotis, Mark A Lawley, and Placid M Ferreira. Polynomial-complexity deadlock avoidance policies for sequential resource allocation systems. *IEEE transactions on automatic control*, 42(10):1344–1357, 1997. 726
- Stuart Russell and Peter Norvig. *Artificial intelligence : a* 727 *Modern Approach.* Pearson Education, 4th edition, 2022. 728
- Bryan W Schlake, Christopher PL Barkan, and J Riley Edwards. Train delay and economic impact of in-service failures of railroad rolling stock. *Transportation research record*, 2261(1):124–133, 2011. 732
- Darja Šemrov, Rok Marsetič, Marijan Žura, Ljupčo Todor ovski, and Aleksander Srdic. Reinforcement learning approach for train rescheduling on a single-track railway.
 Trans. Res. Part B: Methodological, 86:250–267, 2016.
- Sudhir Kumar Sinha, Shripad Salsingikar, and Siddhartha
 SenGupta. An iterative bi-level hierarchical approach for
 train scheduling. *Journal of Rail Transport Planning* &
 Management, 6(3):183–199, 2016.
- Christian Strotmann. *Railway scheduling problems and their* 741 *decomposition*. PhD thesis, Univ. Osnabrück, 2007. 742
- Gerald Tesauro and Gregory Galperin. On-line policy improvement using monte-carlo search. *Advances in Neural Information Processing Systems*, 9, 1996. 743
- Johanna Törnquist and Jan A Persson. N-tracked railway traffic re-scheduling during disturbances. *Transportation Research Part B: Methodological*, 41(3):342–362, 2007. 748
- Robin Vujanic and Andrew J Hill. Computationally efficient dynamic traffic optimization of railway systems. 750 *IEEE Transactions on Intelligent Transportation Systems*, 751 23(5):4706–4719, 2022. 752