

# Towards Robust Admission Control in Delay-Constrained Routing Problems

Antonio Frangioni<sup>a</sup> Laura Galli<sup>a</sup> Giovanni Stea<sup>b</sup>

<sup>a</sup> *Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy. E-mail: frangio@di.unipi.it, laura.galli@unipi.it*

<sup>b</sup> *Dipartimento di Ingegneria dell'Informazione, Università di Pisa, Largo Lucio Lazzarino 1, 56122 Pisa, Italy. E-mail: giovanni.stea@unipi.it*

---

## Abstract

Delay-Constrained Routing (DCR) problems require to route a new flow in a computer network subject to worst-case end-to-end delay guarantees. The delay of a packet flow has three components, one of which is the “queueing delay”, that depends on the scheduling algorithm implemented by the routers of the network. When flows are not independent of each other, i.e., admitting a new flow changes the delay of the existing ones, *admission control* policies are necessary to ensure that existing flows do not become latency-unfeasible. It has been recently shown that admission control runs contrary to the usual objective function employed in these models, i.e., minimization of the reserved rates, significantly worsening network performance. In this paper we investigate the phenomenon and propose a heuristic way to overcome the problem.

*Keywords:* Routing problems, maximum delay constraints, scheduling algorithms, admission control, Mixed-Integer NonLinear Programming, Second-Order Cone Programs, Robustness

---

# 1 Admissible Delay-Constrained Routing

Nowadays, many Internet applications (e.g., industrial IoT, virtual/augmented reality) require stringent *Quality of Service* (QoS) guarantees, notably a maximum end-to-end *Worst-case Delay* (WCD). The WCD depends on the packet scheduling algorithm used at each link and can be bounded by enforcing a minimum rate at each link traversed by the application flow. This gives rise to a family of Delay-Constrained Routing (DCR) problems [2], which seek to simultaneously define the *route* and the *rates* on each link for a new flow, so that its WCD stays below a predefined deadline.

In DCR problems, a packet-switched network is modeled as a directed graph  $G = (N, A)$ ; nodes are routers, and arcs are the links interconnecting them. Packets passing through each node  $i \in N$ /arc  $(i, j) \in A$  experience a fixed *node/link delay*  $n_i/l_{ij}$  and a variable *queuing delay* at the nodes, which depends on all the flows  $Q$  simultaneously active in the network, as well as on the scheduling algorithm that the routers employ. Each  $q \in Q$  is characterized by a path in  $G$  (which, for notational simplicity, we will denote by  $q$  as well) and *reserved rates*  $r_{ij}^q$  for all  $(i, j) \in q$ ; the latter are the main driver of the schedulers' algorithms, as described in details below. An upper bound on the WCD of each packet can be computed if the arrivals  $F^q(t)$  of the flow over the time  $t$  are bounded. The standard assumption is that  $F^q(t + \tau) - F^q(t) \leq \sigma^q + \rho^q \tau$  has to hold for all  $t$  and  $\tau \geq 0$ , where  $\sigma^q$  and  $\rho^q$  are the known *burst* and the *rate* of the flow, respectively. One can then ensure that a proper choice of the paths and the reserved rates guarantees that each packet of a flow will be received within the given *deadline*  $\delta^q$ . The *Admissible Delay-Constrained Routing* (ADCR) problem then considers a “new” flow to be routed in  $G$  (from a source  $s$  to a destination  $d$ , with given burst  $\sigma$ , arrival rate  $\rho$ , and deadline  $\delta$ ) along with the existing ones, and seeks one feasible  $s$ - $d$  path  $p$  and the corresponding reserved rates at each of its arcs, so that *both the new flow and all the existing ones* meet their deadlines. Because the total amount of reserved rates on an arc is limited, it is natural to aim for the smallest possible reserved rates in order to “leave more room” available.

A crucial part of developing ADCR models, discussed at length in [4] and only succinctly recalled here, is the description of the delay that the flows experience. Other than on the (fixed) flows  $q \in Q$  and their (fixed) reserved rates  $r_{ij}^q$ , and on the scheduling algorithm implemented by the routers, the delay depends also on the choices made for the new flow, i.e., the chosen  $s$ - $d$  path  $p$  in  $G$  and its reserved rates  $r_{ij}$  for each arc  $(i, j) \in p$ . To describe the *delay formulæ* we will use the following notation:  $P(i, j) = \{q : (i, j) \in q\} \subseteq Q$

is the set of existing paths traversing arc  $(i, j)$ , which leads to partitioning  $A = A' \cup A''$ , where  $A'$  contains the arcs  $(i, j)$  that are “empty” ( $P(i, j) = \emptyset$ ) and  $A''$  those that contain at least one flow. For the scheduling algorithms to provide guarantees it is necessary to assume that each arc is not oversubscribed, i.e.,  $r_{ij} \leq w_{ij} - \bar{r}_{ij}$ , where  $w_{ij}$  is the *link capacity* and  $\bar{r}_{ij} = \sum_{q \in P(i, j)} r_{ij}^q$  ( $\geq 0$ ) is the total reserved rate of all the existing flows traversing link  $(i, j)$ . Most schedulers share the link capacity proportionally among the flows, based on their reserved rate. Hence, a flow’s *guaranteed rate* on each  $(i, j) \in p$  is given by  $g_{ij} = (w_{ij}r_{ij})/(\bar{r}_{ij} + r_{ij})$ . It is easy to see that  $r_{ij} \leq g_{ij} \leq w_{ij}$ : the upper bound is achieved when  $\bar{r}_{ij} = 0$  ( $\equiv (i, j) \in A'$ ), i.e.,  $p$  is the only path traversing  $(i, j)$ , whereas the lower bound when  $\bar{r}_{ij} + r_{ij} = w_{ij}$ , i.e., the arc is fully subscribed. Under the assumption that the minimum guaranteed rate among all links of  $p$  is at least as large as the traffic injection rate, i.e.,  $g_{ij} \geq \rho$  for all  $(i, j) \in p$ , the WCD is finite; ensuring that each packet meets the deadline  $\delta$  can then be expressed by the *WCD constraint*

$$(1) \quad \frac{\sigma}{\min\{g_{ij} : (i, j) \in p\}} + \sum_{(i, j) \in p} (\theta_{ij} + l_{ij} + n_i) \leq \delta ,$$

where  $\theta_{ij}$  is the *link latency* experienced by the flow when traversing the arc  $(i, j)$ , whose exact form depends on the scheduling algorithm employed. Up until [4],  $\theta_{ij}$  has always been computed under the “pessimistic” assumption  $g_{ij} = r_{ij} \forall (i, j) \in p$ , which logically implies that  $(i, j) \in A''$ . Delay formulæ obtained under the “pessimistic” assumption are referred to as *bound estimates*, as opposed to the more accurate *worst-case estimates* employing the exact formulæ based on the guaranteed rates. The latter are exact, but make the models somewhat more complex, while the former still compute safe upper bounds on the WCD. The formulæ for  $\theta_{ij}$  for four main classes of schedulers are summarized in the following table, where  $P^{ij} = |P(i, j)|$ ,  $r_{ij}^{min} = \min\{r_{ij}^q : q \in P(i, j)\}$ , and  $L$  is the maximum packet length in the network:

	bound	worst-case
SRP	$\frac{L}{r_{ij}} + \frac{L}{w_{ij}}$	$\frac{L}{w_{ij}} + \begin{cases} L/g_{ij} & \text{if } (i, j) \in A'' \\ 0 & \text{otherwise} \end{cases}$
GB	$3\frac{L}{r_{ij}} + 2\frac{L}{w_{ij}} \leq \theta_{ij} \leq 6\frac{L}{r_{ij}} + 2\frac{L}{w_{ij}}$	
WRP	$P^{ij}\frac{L}{w_{ij}} + \frac{L}{r_{ij}}$	$P^{ij}\frac{L}{w_{ij}} + \frac{L}{g_{ij}}$
FB	$\frac{L}{w_{ij}} \frac{w_{ij} - r_{ij}}{\min\{r_{ij}, r_{ij}^{min}\}} + P^{ij}\frac{L}{w_{ij}} + \frac{L}{r_{ij}}$	$\frac{L}{w_{ij}} \frac{\bar{r}_{ij}}{\min\{r_{ij}, r_{ij}^{min}\}} + P^{ij}\frac{L}{w_{ij}} + \frac{L}{g_{ij}}$

The key fact about Strictly, Weakly Rate-proportional (SRP, WRP), Group- and Frame-based (GB, FB) schedulers of the above table are:

- (i) SRP has the smallest latency, and WRP outperforms FB. GB can be expected to have a large latency, too, due to the multiplying constants. However, SRP schedulers are more complex than WRP, GB or FB, which allows a network engineer to select different trade-offs between complexity and latency [4].
- (ii) Except for the bound SRP formula and the GB one (that, to the best of our knowledge, does not have a worst-case version), all formulæ introduce *dependences* among flows, for several reasons: if  $\bar{r}_{ij}$  increases, the guaranteed rate decreases, moreover the term  $P^{ij}$  depends on all the flows, and so does  $r_{ij}^{min}$ . For bound-SRP and GB instead, a flow's latency only depends on its own reserved rate.
- (iii) All the above latency expressions are non-linear, but can be expressed by Second-Order Cone constraints.

Due to point (ii) above, in general routing the new flow  $p$  in the network increases the delay of any existing flow  $q$  that shares with  $p$  at least one arc. To ensure that the existing flows remain delay-feasible without changing their routing or reserved rates, *admission control constraints* need be put in place. Those constraints guarantee that the increase in delay due to the new flow does not exceed the *delay slack* of each existing flow  $q$ , which has the general expression

$$(2) \quad \bar{\delta}^q = \delta^q - \frac{\sigma^q}{g_{min}^q} - \sum_{(i,j) \in q} (l_{ij} + n_i + \theta(r_{ij}^q)) \quad ,$$

where  $\theta(r_{ij}^q)$  depends on the scheduler as illustrated in the table. Denoting by  $\Delta\theta(r_{ij})$  the increase in delay due to the new flow passing through  $(i, j)$ , whose form again depends on the type of scheduler, the general expression of the *Admission Control Constraint* (ACC) for an existing flow  $q$  has the form

$$(3) \quad \sum_{(i,j) \in q} \Delta\theta(r_{ij}) \leq \bar{\delta}^q \quad .$$

In particular,  $\Delta\theta(r_{ij})$  does not depend on the existing flow  $q$  (whose contribution is entirely counted into  $\bar{\delta}^q$ ), but only on the reserved rate of the new flow. Constraints (3) may be either linear or nonlinear depending on the scheduler class. However, they can always be expressed within a MI-SOCP. What will be relevant for our arguments is that the choice between bound and worst-case formulæ can be made *independently* between the WCD constraint (1) for the *new* flow and the ACC (3) for the *existing* ones. That is, while intuitively one would either use bound formulæ in both, or worst-case ones in both, nothing prevents mixing the two. As we shall see, this is actually beneficial.

## 2 Admission Control Performance Analysis

In [4] it has been shown that all models solution times are compatible with real-time use in a realistic environment; thus, in [5] the impact on network performance of using different scheduler classes was evaluated. Here we complete the study in two respects: a) we also test worst-case formulæ, which was not done therein, and b) we propose a way to improve the “robustness” of the approach, and therefore the network performance, based on properly mixing the two sets of latency formulæ.

The need for point b) is illustrated by Figure 1 for a number of real-world IP network topologies. On each, traffic simulations were performed following the guidelines of [3,5]. In the figure, the horizontal axis is the *load* of the system (average number of new flow requests per second, considering that each flow lasts one second in average), whereas the vertical axis is the *blocking probability*, i.e., the ratio of rejected flows over the total number of requests. In each simulation a different model is used to choose the feasible routing and rates (or prove that there is none); the model is a combination of scheduling algorithm (among SRP, GB, WRP and FB), latency formula in the WCD constraint (1) for the new flow (either “B”, for bound, or “W”, for worst-case), and latency formula in the ACC (3) (again, either “B” or “W”). All the models were solved by `Cplex 12.6`, ran with default parameters, in short time [4], hence the only relevant fact is their effect on network performance.

The results for “BB” models confirm the expectations: SRP outperforms WRP which outperforms FB, with GB being largely the worst. The differences between the three are significant. However, the results of “WW” models are somewhat counter-intuitive. In particular SRP-WW performs much worse than SRP-BB, even at low loads, despite guaranteed rates being, in general, larger than reserved rates, especially if the network is lightly loaded. Furthermore, the \*-WW models perform very close to each other, with SRP-WW not being any better (and sometimes visibly worse) than the others.

The latter phenomena (for “BB” models) have already been reported in [5], where they are found to be due to a “conceptual mismatch” between the objective function and the ACC. Indeed, when a flow is admitted, its rates are computed based on the *current state* of the network. Because the objective function minimizes the cost of rate allocation, the *smallest* possible rates are selected. Given that rates and delays are inversely proportional, the selected (smallest) rates are those that produce the *largest possible feasible delay*: in other words, when a flow is admitted, its delay slack (2) is always zero. Therefore, unless some flow disappears later on, *any* new flow that

attempts to use the same links will increase the WCD of any existing flow: but because the WCD is already at its maximum, even a fractional improvement is impossible. Thus, new flows may be found to be impossible to route not because there is not enough rate to support them, but because they would disrupt the current flows (i.e., because of the ACC). This indicates that, when ACC is in place, it is necessary to *over-allocate* rates, to introduce some *extra slack* which will allow later flows to use the same arcs. Of course a trade-off shows here, where too much slack takes away precious rate and therefore may again increase the blocking probability, but too little a slack has a similar effect. In [5] it was proposed to simply decrease the deadline  $\delta$  by a small percentage in the WCD constraint (1), thus forcing the model to over-allocate the rates and thereby protect the flow from subsequent increases of the WCD. However, getting the “right” percentage was found to be nontrivial, and in particular it was highly dependent on the specific network topology.

In this paper we explore a different way to protect flows from the latency increase due to the new flows. In a nutshell, the basic observation is that bound formulæ already over-estimate the WCD experienced by a flow, producing a more conservative resource allocation than strictly necessary, i.e., a nonzero delay slack *if the WCD is evaluated with the worst-case ones*. Hence, one may use bound formulæ in (1) for the new flow, but worst-case ones in the ACC of the existing ones. This forces the models to (over-)allocate the rates when admitting a flow, but uses the exact WCD when evaluating whether or not a flow is, later on, disrupted by a new flow request. The corresponding “BW” models do not need any parameter tuning. As shown in Figure 1, the performance of WRP-BW and FB-BW is very close to that of SRP-BB, despite using less complex schedulers, that in general require larger rates to achieve the same WCD, and much better than that of the corresponding “BB” or “WW” models. We remark that there is no SRP-BW model because when a flow is admitted using SRP-B formulæ, its WCD is independent from the other flows, and therefore there is no ACC. The running time of the “BW” models is comparable with the others’ (which is hardly surprising since they are somewhat simpler than “WW” ones [4]), and therefore not an issue.

### 3 Conclusions and Future Research

To conclude, this paper improves on [5] showing that one can actually use “cheaper” schedulers, e.g. WRP and FB ones, and obtain network performance similar to (or even better than) those of the more costly SRP ones. However, doing so requires using the exact WCD formulæ instead of the

bounds employed so far in the literature. Furthermore, it requires doing so “carefully”; in particular, using the exact WCD for both rate allocation for a new flow, and ACC for the existing ones, leads to too small (zero) delay slacks, and therefore to non-robust solutions. A simple, yet effective way to avoid that is to construct models properly mixing the two types of formulæ, which (at least on all the instances in our test set) attains good network performance without any need for hard-to-tune robustness parameters. The topic deserves future investigations. It would be useful to define a proper concept of *admission control robustness* and possibly to remove the assumption that all existing flows maintain the same path and rate allocation, thus introducing the possibility of *recourse actions*, as described e.g. in [1,6] for different applications. This, however, requires that the recourse actions be actually implemented in the context of packed-based computer networks, where some flows are sensitive to even minor QoS disruptions (e.g., real-time applications like equipment control). Finding the right compromise between flexibility and network stability may lead to a significant improvement in the network performance.

## References

- [1] V. Cacchiani, A. Caprara, L. Galli, L. G. Kroon, G. Maróti, and P. Toth. Recoverable robustness for railway rolling stock planning. In *Proceedings of ATMOS 2008, Karlsruhe, Germany, September 18, 2008*.
- [2] A. Frangioni, L. Galli, and M.G. Scutellà. Delay-constrained shortest paths: Approximation algorithms and second-order cone models. *Journal of Optimization Theory and Applications*, 164(3):1051–1077, 2015.
- [3] A. Frangioni, L. Galli, and G. Stea. Optimal joint path computation and rate allocation for real-time traffic. *The Computer Journal*, 58(6):1416–1430, 2015.
- [4] A. Frangioni, L. Galli, and G. Stea. Delay-constrained routing problems: Accurate scheduling models and admission control. *Computer and Operations Research*, 81:67–77, 2017.
- [5] A. Frangioni, L. Galli, and G. Stea. Qos routing with worst-case delay constraints: Models, algorithms and performance analysis. *Computer Communications*, 103:104–115, 2017.
- [6] L. Galli. Combinatorial and robust optimisation models and algorithms for railway applications. *4OR*, 9(2), 2011.

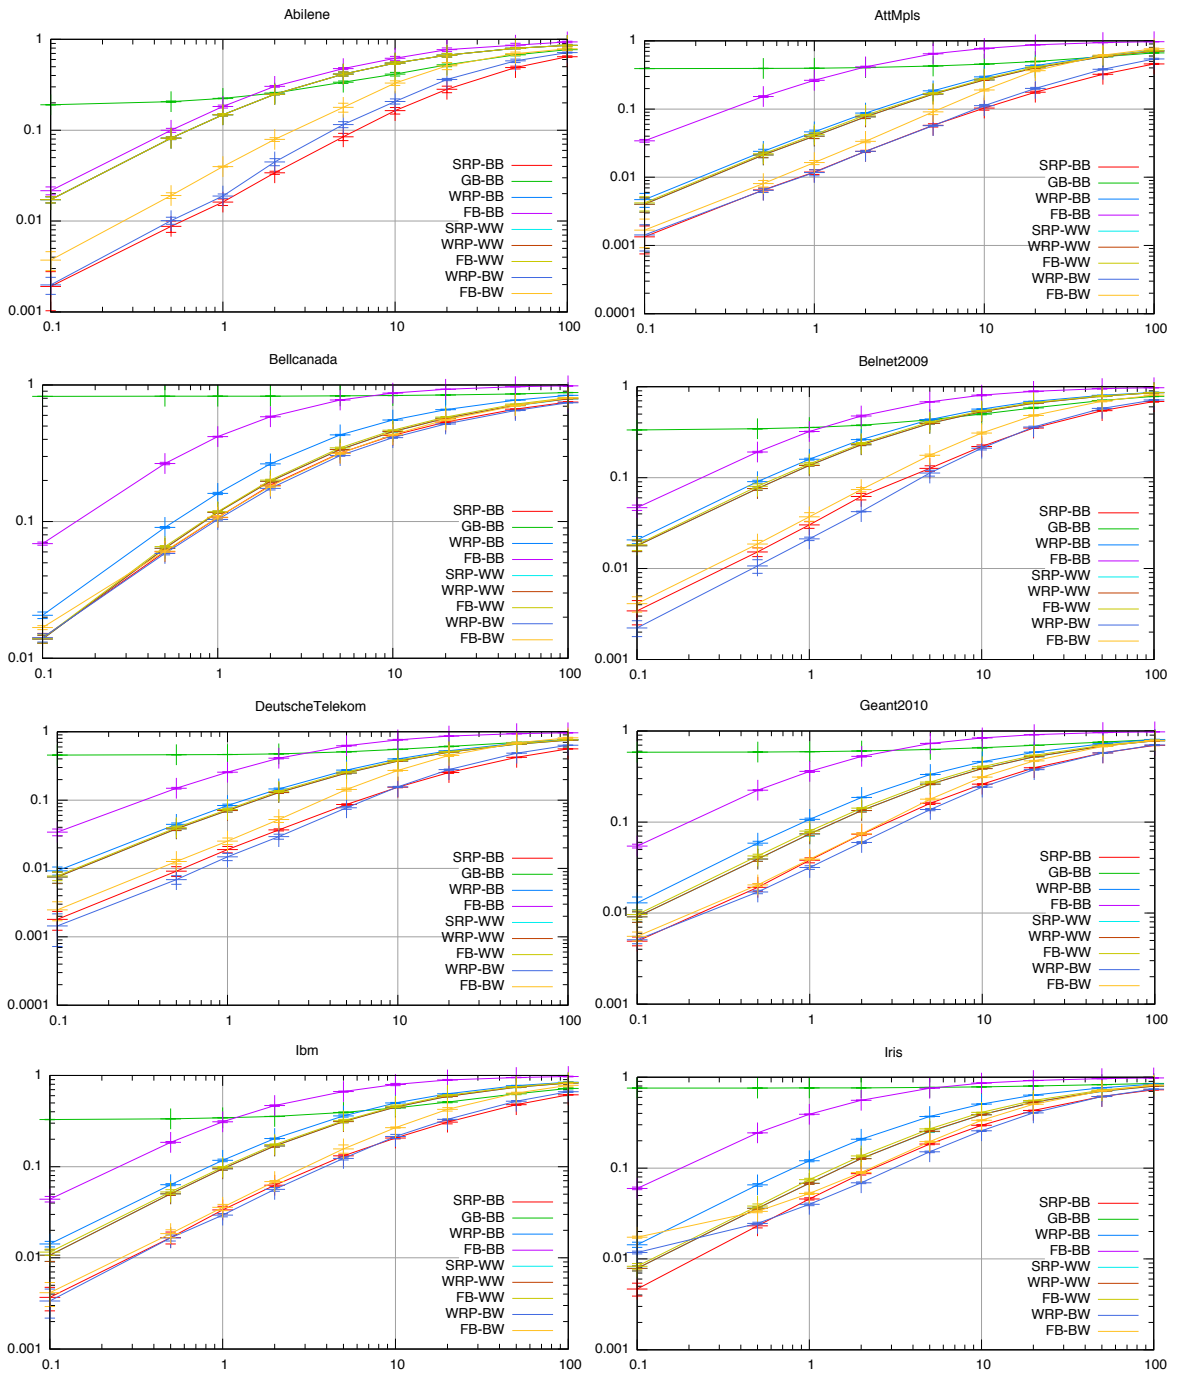


Fig. 1. Blocking probability for all topologies