

Technical report 06-013

Traffic adaptive control of a single intersection: A taxonomy of approaches*

R.T. van Katwijk, B. De Schutter, and J. Hellendoorn

If you want to cite this report, please use the following reference instead:

R.T. van Katwijk, B. De Schutter, and J. Hellendoorn, "Traffic adaptive control of a single intersection: A taxonomy of approaches," *Proceedings of the 11th IFAC Symposium on Control in Transportation Systems*, Delft, The Netherlands, pp. 227–232, Aug. 2006.

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: <https://www.dcsc.tudelft.nl>

*This report can also be downloaded via https://pub.deschutter.info/abs/06_013.html

TRAFFIC ADAPTIVE CONTROL OF A SINGLE INTERSECTION: A TAXONOMY OF APPROACHES

R.T. van Katwijk* B. De Schutter**
J. Hellendoorn**

* *Netherlands Organization for Applied Scientific Research,
TNO, ronald.vankatwijk@tno.nl*

** *Delft University of Technology, Delft Center for Systems
and Control {r.vankatwijk, b.deschutter,
j.hellendoorn}@dcsc.tudelft.nl*

Abstract: The design of a traffic adaptive control system is as well a science as an art. Along the way compromises have to be made in order to end up with a workable system that is not only able to come up with good signal timings, but is also able to deliver them on time. In this paper we propose a taxonomy of the various traffic adaptive control algorithms based both on their underlying principles and the compromises that were made to come up with a workable, albeit less optimal system.

Keywords: Traffic, Adaptive, Dynamic Programming, Branch-and-bound, Optimization, Simulation

1. INTRODUCTION

A common function of a traffic controller is to seek to minimize the delay experienced by vehicles through manipulation of the traffic signal timings. There are various levels of sophistication in traffic signal control system applications.

Basically, the modes of operation can be divided into three primary categories:

pre-timed Under pre-timed operation, the timing plan is based on predetermined rates. These predetermined rates are determined from historical data. Pre-timed control frequently results in inefficient use of intersection capacity because of the inability to adjust to variations in traffic flow and actual traffic demand; this inefficiency is pronounced when flows are substantially below capacity.

responsive In traffic responsive mode, signals receive inputs that reflect current traffic conditions, and use this data to choose an appropriate timing plan from a library of different plans. In traffic responsive mode, the signal timing plan responds to current traffic conditions measured by a detection system. The general traffic responsive strategies in use are either selection of a background signal timing plan based on detector data, or online computation of a background timing plan. The computation time interval may range from one cycle length to several minutes.

actuated Vehicle-actuated controllers operate in real-time by applying a control in response to the current traffic state. An actuated controller operates based on traffic demands as registered by the actuation of vehicle and/or pedestrian detectors. There are several types of actuated

controllers, but their main feature is the ability to adjust the length of the currently active phase in response to traffic flow. The green time for a phase is a function of the traffic flow, and can be varied between pre-timed minimum and maximum lengths depending on flows. Although vehicle-actuated controllers operate in real-time, they attempt no systematic optimization.

adaptive Traffic adaptive systems are currently the most advanced and complex control systems available. Traffic adaptive systems apply an optimization algorithm in real-time to create optimal signal timings. They differ from vehicle-actuated controllers because they incorporate decision making. That is, the system evaluates a set of feasible control actions and chooses an action that is optimal with respect to its current objectives. As traffic adaptive systems incorporate information from further upstream in their decision making (i.e. predicted arrivals) traffic actuated controllers are considered to be myopic (i.e. shortsighted) with respect to their control actions.

With recent advances in communication network, computer, and sensor technologies, there is increasing interest in the development of traffic adaptive signal control systems. Numerous systems have been proposed including PRO-DYN (Henry *et al.*, 1983), (Barriere *et al.*, 1986), (Henry and Farges, 1989), UTOPIA-SPOT (Mauro, 1990), OPAC (Gartner, 1983), (Gartner *et al.*, 1995), (Gartner *et al.*, 1999), RHODES (Sen and Head, 1997), (Mirchandani and Head, 2001), SPPOINT (Dion and Hellinga, 2001), (Dion and Hellinga, 2002) and ALLONS-D (Porche *et al.*, 1996), (Porche, 1998), (Porche and Lafortune, 2005). This overview is based on these references and the references contained therein.

The design of a traffic adaptive control system is as well a science as an art. Along the way compromises have to be made in order to end up with a workable system that is not only able to come up with good signal timings, but is also able to deliver them on time. Each adaptive system mentioned above has a different approach in dealing with the computational complexity of determining the best set of signal timings.

All adaptive systems therefore have their own specific strengths and weaknesses that make that system more - or - less suited for particular networks and traffic demand patterns. In this paper we propose a taxonomy of the various traffic adaptive systems based both on their underlying principles and the compromises that were made




Fig. 1. decision tree

to come up with a workable, albeit less optimal system.

2. A TAXONOMY OF APPROACHES

Traffic signal control essentially comes down to making the right decisions at the right time. As such the traffic signal control problem solved by all traffic adaptive systems can be formulated in the form of a general decision problem. This general decision problem in turn can be represented as a simple decision tree such as that shown in figure 1.

The root of a decision tree represents the current state s_i , $s_i \in S$, where i is the current time index and S is the set of all states. States in the decision tree of figure 1 are represented by open circles and actions by solid circles. The cost involved in order to transition to the subsequent state, s_{i+1} when deciding for an action u_i is denoted by c_i . In figure 1 there are two possible successor states depicted for each action which permits the specification of stochastic outcomes.

In general, the nodes of a search tree represent choices. These choices are mutually exclusive and therefore partition the search space into two or more simpler sub-problems. At each time step, the controller observes the system's current state s_i , and selects a control action, $u \in U_i$, where u is the action and U_i is the finite set of actions available to a controller in state s_i . When the controller chooses an action $u \in U_i$, the cost incurred by taking that action and subsequently transition to state s_j with probability $p_{i,j}(u)$, is denoted by $c(i)$. The objective of a traffic adaptive system is to find an optimal sequence of actions.

Looking at the various traffic adaptive systems we can discern the following features on which they differ:

- the *optimization method*. Is the optimal sequence of actions found by searching the decision tree using a rule-based method, or an

approach based on dynamic programming or branch-and-bound?

- the possible *actions* (u_i) considered in the optimization. Is the order in which phases can be given green to predetermined or can this be determined (and optimized) on-line?
- the length and resolution of the *planning horizon* over which an optimal sequence of actions is sought (i.e. the depth of the decision tree). Is the length of the horizon fixed (e.g. 2 minutes) or dependent on current traffic conditions? Is the resolution static (e.g. is the horizon divided into 5 seconds intervals) or is it dynamic (e.g. dependent on projected arrival times)?
- the *update frequency*. How often can the optimization be done (i.e every 0.5 seconds or every 5 seconds)?
- the *delay model*. How is the performance (c_i) of each evaluated action (u_i) evaluated? How accurate is the model used in optimizing the signal timings? Is a fast vertical queuing model used instead of a slow but possibly more accurate simulation model?

The following sections elaborate on each of these features and how each traffic adaptive system differs in how these features are filled in.

3. OPTIMIZATION METHOD

The objective of the system is to operate such that the total cost over the entire planning horizon is minimized. Thus, the task of the controller is to obtain a sequence of control actions $[u_0, u_1, \dots, u_T]$, also referred to as a policy or control trajectory, such that the expected cost is minimized. In the case of an infinite planning horizon, a discount factor, $\gamma < 1$, is typically applied to future costs to obtain a finite estimate of the *cost-to-go* from the current state i , denoted by $f(i)$. The optimal cost-to-go value, denoted by $f^*(i)$, is a function of the immediate cost of applying the control plus the expected cost-to-go from the subsequent state, a relationship encapsulated in the following recursive expression which is also known as Bellman's Equation.

$$f^*(i) = \min_{u \in U(i)} \left\{ c_i(u) + \gamma \sum_{j \in S} p_{i,j}(u) f^*(j) \right\} \quad (1)$$

As the decision space has a tree-like structure, the search for the optimal sequence of decisions corresponds to building the tree. An exhaustive search of the entire decision space results in a full tree being built. Since search space size grows exponentially with problem size, it is not possible to explore all assignments except for the smallest problems. The only way out is to not look at

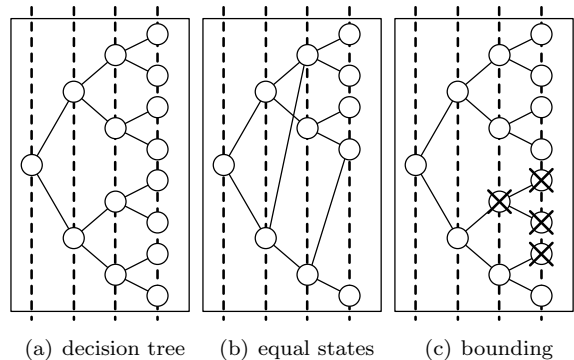


Fig. 2. Different approaches to not having to build the entire search tree

the whole search space. Efficiency in searching the decision space is considered by the degree to which the entire tree will not have to be built to find an optimal path. In (Shelby, 2004) several well-known algorithms are assessed based on computational speed and on the quality of the results (in terms of delay).

Although the resulting state of traffic is subject to probabilistic outcomes as depicted in figure 1, all adaptive systems reviewed have chosen to employ a deterministic model. Due to this simplification, decision tree diagrams corresponding to the traffic signal control problem will generally neglect the use of explicit actions nodes as each action results in only one possible next state.

Dynamic programming and branch-and-bound (and combinations thereof) are the techniques that are predominantly used in traffic adaptive systems.

3.1 Dynamic Programming

The applicability of the approach depends on the opportunities for state aggregation within the decision tree. The strength of dynamic programming is that it can prevent that optimal solutions to subproblems it has already solved are recomputed. In order to do this, the solutions to already solved problems are saved. This approach is called memoization (not memorization, although this term also fits). It is however only possible to reuse a previous solution when states and thus the corresponding subproblems can be considered equal. This is why approaches like RHODES and PROLYN use an approximate state equivalence relation to attain greater efficiency. Figure 2(b) shows how dynamic programming can assist in not having to build the entire search tree as depicted in 2(a). RHODES, PROLYN, OPAC and SPPORT all employ dynamic programming as their method of optimization.

3.2 Branch-and-Bound

Branch-and-bound is a general method for finding optimal solutions of various optimization problems, especially in discrete and combinatorial optimization. It belongs to the class of implicit enumeration methods. One way to do this is by proving that certain areas of the space contain no solutions. The core of the approach is the simple observation that (for a minimization task) if the lower bound for a sub-problem A from the search tree is greater than the upper bound for any other (previously examined) sub-problem B, then A may be safely discarded from the search. This is the bounding-part of the branch-and-bound approach. Figure 2(c) shows how branch-and-bound can assist in not having to build the entire search tree depicted in figure 2(a). Of the adaptive systems reviewed only ALLONS-D and SPOT employ the branch-and-bound method in its pure form. RHODES employs a hybrid system in which branch-and-bound techniques are applied within a dynamic programming framework.

In order to obtain a tight upper bound an initial path must be established through the search tree for which it is most likely to obtain a good solution. This involves that initially parts of the search space that are unlikely to contain good solutions are ignored. This is done by using *heuristics*. Heuristics are used to explore promising areas of the search tree first. This can be done by using problem specific knowledge (often borrowed from current practices in tuning traffic responsive and vehicle-actuated controllers) or by reusing information gained from previous optimizations.

4. ACTION SPACE

The width of the tree to be searched is dependent on the number of decisions that can be made at each point in time. In its simplest form the choice available is that between extending the current phase or switching to the next phase. This is the approach taken in OPAC, ALLONS-D, PROLYN and SPPORT. Although this approach significantly reduces the number of options to consider, it does not allow arbitrary phase sequencing. In its most elaborate form the choice available is that between phases. This approach allows the arbitrary sequencing of phases but comes at a cost in the width of the search tree. This is the approach chosen by UTOPIA-SPOT. Both of these approaches are shown in figure 3.

A compromise between these two extremes is found in allowing phase skipping. When the skipping of phases is allowed any phase sequence can be attained. This is shown in figure 4. This is the approach taken in the COP-system. The downside




Fig. 3. Different approaches regarding actions considered




Fig. 4. phase skipping vs. selection

of this approach is that when the initial phase sequence is chosen wrong the gain in width is counteracted with an increase in tree depth. This is not strictly necessary however. It is in principle possible to choose the next phase dynamically.

5. PLANNING HORIZON

Traffic adaptive systems employ a traffic model to evaluate alternative traffic signal timings over a planning horizon. The length of the planning horizon as well as how the horizon is split up into successive intervals differs between each adaptive system. Typically however the horizon has a fixed length (of typically 1 to 2 minutes) and is subdivided into fixed intervals. From their descriptions we can deduce that OPAC, PROLYN, SPPORT, and ALLONS-D all use or have used 5-second time-steps.

If the horizon is chosen to short and the optimization algorithm is faced with the choice whether to a) completely serve a phase dispersing at a slow rate, or b) preempt that phase in order to switch to a phase with a higher dispersion rate, it would counter-intuitively chose for the latter. This is why many of the adaptive systems that employ shorter horizons have introduced terminal

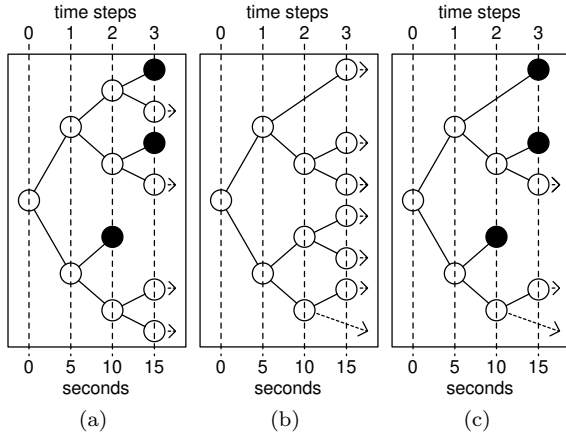


Fig. 5. Different approaches regarding the length of the planning horizon

costs in order to penalize residual queues at the end of the horizon (Newell, 1998), (Shelby, 2004).

The ALLONS-D algorithm takes a different approach wherein the length of the horizon depends on the current traffic situation. The ALLONS-D algorithm enlarges the horizon until it finds a solution in which all projected arrivals are cleared. Although the idea of a horizon that shrinks or grows dependent on the traffic situation sounds attractive, it might not turn out this way in the case of the ALLONS-D algorithm. In saturated conditions - with many projected arrivals - the length of the horizon might become so large that the optimization method used by ALLONS-D might be unable to come up with an answer in time.

The approach where both the length of the planning horizon and the length of the time intervals in which it is subdivided are variable is not applied by any of the algorithms reviewed.

6. UPDATE FREQUENCY

Traffic adaptive systems rely on predicted arrivals. As the distance over which these arrivals are predicted increases the reliability of these predictions often decreases. This is why a rolling horizon is often applied. The concept of a rolling horizon originated in operations research and is used to determine a short term policy based on a longer term analysis. All adaptive systems reviewed that depend on arrival predictions employ the concept of a rolling horizon. These algorithms implement only the first (few) action(s) of the control plan after which a new optimization is performed.

The rolling horizon concept is visualized in figure 6, where each horizontal bar denotes the calculated control plan for a decision horizon. The system commits to this control plan until the optimization process refreshes it. The period between updates is denoted by the commitment period.




Fig. 6. Rolling Horizon

The amount of time that passes between each subsequent optimization (the roll - or - commitment period) is, for all adaptive systems reviewed, equal to the length of the intervals which subdivide the planning horizon. For most adaptive systems reviewed the length of the interval is typically equal to 5 seconds. Waiting 5 seconds between decisions to switch or extend the current phase can however have a significant impact on delay.

Consider, for example, the case where a queue dissipates earlier than predicted. With a 5 second commitment period, an adaptive system may take up to 5 seconds to realize the error, resulting in the waste of green time. With a 1-second decision resolution, controllers could quickly terminate phases as queues clear out, reallocating this time or capacity to phases that do have traffic to serve.

Note that, as all adaptive systems choose their commitment period equal to the length of the interval in which the planning horizon is subdivided, switching from a 5 second to a 1-second decision resolution increases the number of time-steps in the planning horizon by a multiple of 5. This imposes too much of an increase in computational effort for many algorithms to solve in real-time. Thus, the typical trade-off is to also decrease the duration of the planning horizon.

7. DELAY MODEL

All adaptive systems reviewed consider individual vehicles in determining the control delay brought about by a chosen control plan. In that respect all adaptive models can be considered to use a microscopic model. However, because the delay model is applied many times when exploring the search tree, all models have to make some sacrifices with respect to the level of detail on the employed model.

Known, commercially available, microscopic simulation models like Paramics, VISSIM, and AIM-SUN are unfit for use in real-time optimization. This is why simple event-based and cellular automaton models are predominantly used within adaptive systems. At first these models employed vertical queuing models, but many adaptive systems have since switched to using horizontal queuing models so that queue spill backs to upstream

intersections can explicitly be considered in the optimization.

8. CONCLUSIONS

The previous has shown that there are many different ways to configure a traffic adaptive system. Although the core of each of the traffic adaptive systems reviewed is based on the idea of finding a short term policy on the basis of a long term analysis, they differ with respect to the search algorithm applied, the length and resolution of the planning horizon, the update frequency and the delay model used.

Unfortunately computational boundaries still prevent the configuration of a traffic adaptive system in which no compromises have to be made in order to end up with a workable system that is a) able to come up with good signal timings and b) is able to deliver them on time.

As the base performance of an adaptive system is at least as good as that of an actuated controller there are considerable advantages to the deployment of an adaptive system. However, in order to gain the full advantage of traffic adaptive control, the system should be carefully tuned. Computationally complexity, geometry of an intersection, and demand patterns should be considered.

9. ACKNOWLEDGEMENTS

Research funded by the TU Delft spearhead program "Transport Research Centre Delft: Towards Reliable Mobility", the BSIK project "Towards Sustainable Mobility (TRANSUMO)" and the TNO spearhead program "Sustainable Mobility Intelligent Transport Systems (SUMMITS)"

REFERENCES

- Barriere, J.F., J.L. Farges and J.J. Henry (1986). Decentralization vs. hierarchy in optimal traffic control. In: *Proc. of the 5th IFAC/IFIP/IFORS Symposium on Control in Transportation Systems*. Vienna, Austria.
- Dion, F. and B. Hellinga (2001). A methodology for obtaining signal coordination within a distributed real-time network signal control system with transit priority. In: *Proc. of the 80th Annual Meeting of the Transportation Research Board (TRB'01)*. 80th Annual Meeting of TRB.
- Dion, F. and B. Hellinga (2002). A rule-based real-time traffic responsive signal control system with transit priority: Application to an isolated intersection. *Transportation Research Part B* **36**, 325–343.
- Gartner, N. H., C. Stamatiadis and P. J Tarnoff (1995). Development of advanced traffic signal control strategies for intelligent transportation systems: Multilevel design. *Transportation Research Record* (1494), 98–105.
- Gartner, Nathan H., P.J. Tarnoff and C.M. Andrews (1999). Evaluation of the optimized policies for adaptive signal control strategy. *Transportation Research Record* (1683), 105–114.
- Gartner, N.H. (1983). Opac: A demand-responsive strategy for traffic signal control. *Transportation Research Record* (906), 75–81.
- Henry, J.J. and J.L. Farges (1989). Prodyn. In: *Proc. of the 6th IFAC/IFIP/IFORS Symposium on Control, Computers, and Communications in Transportation*. Paris, France.
- Henry, J.J., J.L. Farges and J. Tuffal (1983). The prodyn real time traffic algorithm. In: *Proc. of the 4th IFAC/IFIP/IFORS Symposium on Control in Transportation Systems*. Baden-Baden, Germany. pp. 307–312.
- Mirchandani, P. and L. Head (2001). A real-time traffic signal control system: architecture, algorithms and analysis.. *Transportation Research Part C* **9**, 415–432.
- Newell, G.F. (1998). The rolling horizon scheme of traffic control. *Transportation Research Part A* **32**, 39–44.
- Porche, I., M. Sampath, R. Sengupta, Y.-L. Chen and S. Lafortune (1996). A decentralized scheme for real-time optimization of traffic signals. In: *Proc. of the 1996 IEEE International Conference on Control Applications*.
- Porche, I.R. (1998). Dynamic Traffic Control: Decentralized and Coordinated Methods. PhD thesis. Department of Electrical Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, Michigan.
- Porche, Isaac and Stephane Lafortune (2005). A game-theoretic approach to signal coordination. *Submitted to Transportation Research B*.
- Sen, S. and K.L. Head (1997). Controller optimization of phases at an intersection. *Transportation Science* **3**, 5–17.
- Shelby, Steven G. (2004). Single intersection evaluation of real-time adaptive traffic signal control algorithms. In: *Proc. of the 84rd Annual Meeting of the Transportation Research Board (TRB'04)*.