

# FASTA: a new life for a former dinosaur adapting the modelling scope to new planning requirements

P. Tzieropoulos<sup>1,2</sup>, A. Curchod<sup>1,3</sup> & Y. Putallaz<sup>1,2</sup>

<sup>1</sup>EPFL: École polytechnique fédérale de Lausanne, Switzerland

<sup>2</sup>LITEP: Laboratory for Intermodality and Transport Planning,  
Switzerland

<sup>3</sup>Switzerland TRANSP-OR: Transport and Mobility Laboratory,  
Switzerland

## Abstract

FASTA has been developed for assessing timetable structures in strongly connected networks. It is built for “sketch simulation”, as it does not aim to produce detailed results but rather to assist in identifying major discrepancies between network and timetable, at early planning stages.

In the new planning paradigm (operations first and, then, infrastructure), long term vision starts by defining a timetable in only its basic structure, leaving spare capacity for allocation in a later stage (a year, a month, even a day before operation). Moreover, in long term timetable planning, many alternatives are drawn. To assess reliability of timetables designed 20 or 30 years ahead, there is no need (and no possibility) to use detailed simulation, which requires detailed definition of infrastructure and signalling. To apprehend the basic behaviour of a timetable under multiple assumptions, planners need at this stage to run a huge number of simulations fast and efficiently. Here comes the utility of sketch simulations and tools, such as FASTA, able to perform them efficiently.

The main question during development is to set the right balance between modelling simplicity (unavoidable due to uncertainties on far-future infrastructure and timetable) and accuracy of the results. The trade-off problem between simplicity and accuracy bears no single solution. Therefore, the initial options that have been taken for FASTA some twenty years ago had to be revisited. Developers had to reassess the consequences of simplification and to extend the concept, in order to build a new package, based on a sensibly modified balance, by trading off simplicity for better accuracy.

The paper presents the basic concepts of the simulation package, the initial trade-off between simplicity and accuracy and its justification, the reasons for changing the initial solution, and the latest developments of the FASTA package based on a changed scope in planning needs.

*Keywords: railways, railway network simulation, sketch planning, long term planning, fixed-interval timetable.*



## 1 Introduction

FASTA is a comprehensive software package for the simulation of trains running on a highly interconnected network. Its main purpose is to provide information needed to assess the stability of alternative timetables. The first FASTA version was already operational in 1992 and has been presented in several conferences and seminars. The purpose of this paper is not to provide still another presentation of the package, but to place the last development efforts within the context of the evolving planning philosophy. FASTA foundation principles will be explained in the next two sections, and then the focus will be brought to the evolution of railway planning philosophy and the peculiarities of the Swiss regular-interval timetable, before trying to roughly draw what has been learned through the development process.

Since its inception, FASTA has been designed with cyclic, fixed-interval timetables in mind. This periodic timetable structure offers several advantages to both railway operations and passengers, with the price – however – of increased vulnerability of the timetable for, at least, three reasons:

- periodicity of the operations does not leave any place for unresolved conflicts during timetable design, as any left over problem becomes recurrent, by repeating itself at the pace of the period;
- during operation, increased rigidity of the timetable results in greater risk of snowball effects that may be produced after an initial disturbance;
- further, in highly interconnected, dense networks, the need to guarantee connections in major stations significantly increases the risk of propagating disturbances among trains and spreading them through the whole network.

## 2 Simulating the train operation on a full-scale network

The initial development of FASTA has been the result of a PhD thesis [1] in the Laboratory for Intermodality and Transport Planning (LITEP) of the Swiss Federal Institute of Technology Lausanne (EPFL), with partial financial help of the Swiss Federal Railways (SBB). The very purpose of FASTA is to analyse the behaviour of various timetable structures and patterns on the network scale. By simulating the operations of all trains during a whole day, FASTA shows when, where, and how delays are generated, and how they propagate with time over the whole network. More specifically, FASTA has been designed to [2]:

- identify and analyse delays on the network level,
- detect causes of delays,
- show the geographical distribution of delays.

The software provides the user with, among others:

- delay diagrams along lines, for major nodes (stations) and by train category,
- geographical distribution of delays at any moment of operation.

FASTA makes it possible to assess the stability of a timetable, which may be defined as the ability of a timetable to recover normal planned operations within a given time after the occurrence of an incident (involving small disturbances).

The basic description of FASTA's operations may be defined as follows:



- event controlled, discrete simulation modelling, events being arrival and departure of trains in nodes (stations);
- two working modes; in deterministic mode, the running times of trains are fixed; in stochastic mode, the running times of trains are generated according a distribution function;
- input data include, besides the timetable whose stability is to be assessed, a broad description of the network (in terms of lines and nodes) and of the trains,
- the simulation takes into account specific constraints, such as minimal headways and the need to guarantee connections between trains at stations;
- inputs and results are stored in specific databases and may be exported for further analyses.

The rail network is modelled using classical graph theory concepts. Nodes are stations in the network, and arcs are sections between stations. Each section has a specified number of physical tracks on which direction movement is specified (mono- or bi-directional). Routes of trains are modelled as paths in the graph that are built as sequence of sections. The exact routing of the trains is determined dynamically during simulation of train-runs.

The model includes the following node types:

- connection stations, where passenger connections between trains are defined and should be ensured,
- stop stations, where trains stop, but there is no need to ensure connections,
- crossings, that are points in the network where train routes intersect,
- virtual nodes, used to model specific network details and exceptions.

Trains providing identical services are grouped in train families.

The timetable is defined by the arrival, departure and passing-through times of all trains. Using this information and a basic set of expert rules, the model generates a list of trains that should ensure passenger connections in major stations. This list can be modified and completed by the user. In the same way, the model generates also the rolling stock roster, although rules to do it are simpler than for train connections.

FASTA offers two distinct models for simulating the operation of the railway network:

- deterministic; train travel time are fixed, as defined in the timetable, and possibly shortened for delayed trains (as allowed by predefined margins);
- stochastic; travel times are drawn from pre-defined travel time distributions.

In deterministic mode, while trains are on time, they use the travel time defined in the timetable. If they are late, they can exploit their margin to reduce or to recover their delay. Stochastic simulation reflects slight variations of travel time, such as those found in daily operation. In stochastic mode, two travel time distribution curves are used, one for trains running on schedule, and another one, more optimistic, for delayed trains, assuming thus that delayed trains will try to run faster, in order to recover their delay.

In both modes, local disturbances may be introduced in connection nodes as well as on sections, to model disruptions (incidence of track works, extended dwell times in a station, and so on) that appear during normal operation.



Disturbances can be applied on one or several successive periods of time, one period being defined as 60 minutes. However, with the first version, there were fairly strict limits on the kind of disturbances that could be applied.

### 3 The balance between simplicity and accuracy

Modelling of traffic control and of several operational constraints was fairly rudimentary. This was so by design in order to keep the model within stringent requirements of usability and to ensure that computing times remain affordable. To understand the rationale of this choice, one has to keep in mind that FASTA has been built to assist planning of the operations at a very early stage, where assessing timetable structures is critical, and should precede the detailed design of full-scale timetables. At this stage, planners need to get a first idea of the big picture, before being submerged in a sea of details, which is unavoidable during accurate timetable design. To do it efficiently, the modelling effort should be kept below an affordable level; if ever modelling the network and the operations' alternatives requires too huge an investment in time, planners might as well step directly into designing the detailed timetable.

On the other hand, loss of accuracy (resulting from simplification) should also be kept to a minimum. If oversimplification leads to unacceptably inaccurate results, then the model defeats its own purpose, and planners would need to resort to detailed studies to get robust insights for their possible choices.

In this sense, usability may be defined as the need to achieve a very delicate trade-off between the effort required to make the model working, and the accuracy of its results. Both those requirements bear, obviously, antinomy.

To contain the modelling effort within acceptable boundaries, simplifications had to be adopted during FASTA development, and they involve several issues on modelling the network and the operations. As an example, some of them are raised in the following part of this section [3].

Initially FASTA accepted only 3 different train categories, and that limits the number of train types that may be modelled; letting for instance model Intercity, fast and regional trains, but excluding freight trains from the simulation.

The list of connecting trains is used during simulation by forcing trains to wait for their connections. If the arriving train is late, its delay is transmitted to the waiting train, unless the incurred delay becomes bigger than a defined maximum threshold, in which case the connection between the two trains is broken. This threshold, i.e. the maximum waiting time for ensuring connections, is defined in each station and depends on the train categories.

The model ensures that train services from terminal stations start only if there is an available train-set as required by the roster plan. A threshold ensures that if the turn-round of the train-sets is broken (i.e. the delay of the incoming train-set is too high), an extra train-set is introduced to take over the service. There is no limit in availability of extra train-sets or drivers at the terminal stations.

Some operational parameters, such as minimal dwell times or minimal transfer times for passenger connections can be vary between peak and off-peak hours.



The capacity of the stations is only partially verified. The model ensures that the number of platform tracks is sufficient. However, delays due to lack of available routes to enter or leave a station are not modelled and are considered to be included in the running time of trains.

Sequence of trains on the same track section is controlled by the headway value. A basic headway is associated with a section and a running direction; it was valid for any train sequence. As the headway was only verified at the departure and arrival nodes, there was a potential risk that trains heap on a section when the destination node is occupied.

#### **4 Planning the services and, then, the infrastructure**

During the time the initial version of FASTA was developed, Switzerland underwent what we may now call a major shift in the philosophy of national transport policy. This shift was produced by a parallel evolution along two separate tracks.

The first track was launched in the mid 70s. Aware that infrastructure development in road and rail networks had already reached a satisfactory level, the Swiss authorities set up a “wise men” committee, involving both top-level academic personalities and policy makers, with the mandate to optimise the use of the multimodal transport system as a whole. Their goal was to lay the foundation of a “Global Swiss Concept for Transport” and the rationale behind it was that further improvement of level of service would be produced not by extending the networks’ capacity, but by making better use of existing facilities. This effort produced eventually (i.e. almost 10 years later) a comprehensive plan, whose key element was a major upgrading of the backbone of the rail network and, namely, the construction of two high speed lines crossing diagonally the country (SW-NE, and NW-SE). This plan was submitted to a large public consultation, which brought back very few positive comments and a huge wave of fundamental, negative reactions. The latter raised the regional equity issue, as the upgrade of the backbone of the rail net would benefit mainly the already strongly developed regions and, by widening the gap between them and the more peripheral regions, would be detrimental to balanced development.

The Swiss decision making system being heavily reluctant to persistent oppositions [4], the reaction of the authorities was to require the Swiss Federal Railways to set up a rail development policy that generates significant improvement of the level of service for all regions. This decision launched the second evolution track, and one year later, in 1985, the Swiss Railways came out with the so-called Rail 2000 project. The idea behind it was to take advantage of the periodic nature of a fixed-interval timetable, which makes that any particular event repeats itself at every period. Thus, by making all trains entering at a given moment in a station, one could ensure passenger connections from any origin to any destination; with an hourly fixed-time timetable, this “train meeting” occurs every hour. Now, to extend that “connecting to all destinations” principle over the whole network, one has to define pivot points on it that are distant at modulo the period between them. That means in practice that, with an hourly fixed



interval timetable, if we find major stations distant between them with a travel time of 1, 2 3 ... hours, we may ensure recurrent meetings of all entering trains, in all those stations, every hour.

This Rail 2000 principle put upside down the approach in railway planning. To make it work, one needs to ensure that trains run with the requested travel times (i.e. modulo 1-hour) between pivot stations. Instead of building new infrastructure, high speed lines for instant, and then design the services we may provide with it, the new planning paradigm was to start by designing the service (recurrent hourly many-to-many connections of all trains in every major station) and then upgrade the infrastructure and the rolling stock to meet the service objective.

Promoted with the slogan “Not as fast as possible, but just as fast as needed”, Rail 2000 was easily and largely confirmed in a popular vote in 1986 (57% approval rate), and shapes since then the development of the Swiss railways.

Meanwhile, along the first track, a last trial to produce something usable met a dead-end, with the massive rejection of a “Coordinated Transport Policy” constitutional amendment in 1987 (55% of the rejection rate among the voters, and also rejected by 19 out of 23 cantons). That put a gravestone to any whim to explicitly define a national transport policy. Now, the amusing paradox is that today there is one, well defined and easily recognisable Swiss transport policy; it was just built steadily, by adding one element after another, by patiently matching every single piece of the puzzle to the already existing ones.

## 5 Fixed-interval timetables and their specificities

Fixed-interval or regular timetables are periodic; this is a tautology. Now, as already stated, periodicity involves both advantages and some less desired properties. Advantages for the rail customer are maximised when the period of the timetable is a “round” number, preferably 60 minutes. With a 60-minutes periodical timetable, regular users need only to remember the minutes of departure and arrival of their frequently used services; they can use the system without previously consulting the timetable. This property, combined to the fact that connections are also regular and recurring, greatly enhances the perceived, seamless usability of the system. Train operators as well, with regular, repetitive operational patterns, have a really easier job in planning and organising the operations: conflict solving during timetable construction extends to all recurring conflicts through the day, design of personnel and rolling stock rosters becomes also easier, and so on. All this comes at the price of increased operational costs, as the introduction of regular timetables results in increasing train mileage due to the unavoidable increase of supply in terms of volume of provided services. It should be stated however that the growth rate of operational costs is less than the one of the train mileage, thanks to better productivity for the overall system; moreover, actual experiences show that if the introduction of regular timetable is carried out and marketed successfully, increased revenue due to higher demand may sometimes completely compensate the increased costs.



On the other hand, public transport services are based on discrete flows, and this is especially true for railways; operation of the system is heterogeneous. Periodicity of the operations promotes even greater heterogeneity, by concentrating events at specific moments and at specific locations. It was this concentration that proved instrumental in designing the Rail 2000 concepts. It is also the same concentration that creates bottlenecks when the volume of services comes close to the infrastructure capacity. That makes both for an easier and a more complex response to the need of extending current capacities. In a network based on (more or less) continuous flow, such as the road network, upgrading capacity calls for an action on the nodes and in enlarging entire sections (or links). Generally, this action provides sustainable better capacity, i.e. capacity that will be useful even when (and if) the structure of flow changes. In a rail network operated with a regular timetable, bottlenecks appear at the stations and on specific points in the network, where are located the train crossings or the overtaking of trains. To increase the capacity within the given timetable structure, one needs not enlarging a full section; it is sufficient to provide extra capacity at the very spots that constitute bottlenecks. This is the pleasant part: action is more concentrated and, thus, less costly. At the same time, this is also the concept's fragility: changing the timetable shifts the location of the bottlenecks and makes useless many previous investments in extra capacity.

What comes out of this is the need to carefully analyse multiple timetable structures, in order to orient future investment in extra capacity and make them as less dependent as possible on a given timetable structure. And here comes in FASTA, which makes it possible to analyse the largest possible spectrum of timetable structures and to pinpoint structural vulnerabilities of the network.

## 6 Shifting the trade-off

With those prerequisites in mind, the first version of FASTA has been developed to let timetable planners test in due time the largest possible spectrum of timetable structures, leaving the detailed design of a given timetable for a later stage. FASTA also assists network managers in assessing sustainability of their investment for increasing capacity, i.e. in choosing the options that are the less dependent on a particular service structure. To that purpose, timetable specification had to be alleviated [5], as had to be simplified the modelling of several system components, such as signalling, for instance.

Meanwhile, with time passing and with added experience by using the first version of the software, FASTA eventually proved to be an extremely versatile tool, useful not only to the first purpose for which it has been built (timetable structures' assessment), but also for all other kind of analyses where the overall vulnerability of the network was involved. One such case was the attempt to identify a global optimum between holding up the tracks for maintenance works and the use of infrastructure's capacity: maintenance uses time resources and, thus, reduces the system's output and revenue; but faulty or insufficient maintenance leads to degraded substance of the infrastructure and, then, reduces



system output and revenue [7]. Trying to deal with this issue highlighted many limits of the first FASTA version, which precluded using it satisfactorily. That led to the development of an almost new software, called for historical reasons FASTA II, where - compared to the initial version - the trade-off between simplicity and accuracy has been deliberately shifted towards more accuracy.

Among a multitude of changes, the most notorious are:

1. Extending the model limits in terms of maximum number of nodes and sections, and in terms of maximum number of trains, as well.
2. Enhancing the modelling of incompatible routes through junctions, to bring the modelling as close as possible to real conditions.
3. Complete redesign of the service typology, setting at 20 the upper limit of user-defined train services. The initial limit of only 3 train categories made it impractical to model freight trains, and - to take into account the impact of un-modelled trains - the user had to resort into tricking the stochastic mode.
4. Minimum headway in train sequence can be modulated according to the categories of the trains that follow each other.
5. A capacity value is automatically assigned to each section, corresponding to the maximum number of possible trains on the track; entering a track is precluded not only when a train already runs in the opposite direction, but also when the number of trains on the track reaches the track's capacity. Doing so avoids the risk that trains heap on a section when the destination node is occupied, due to the way FASTA checks headways, as already mentioned.
6. Greater flexibility in controlling the roster of rolling stock. Contrary to the initial version, where rostering of trains between train categories was fairly restricted (although not impossible), rostering of train-sets is now possible across train families and may vary through time periods. The price to pay for it, however, is increased user involvement, as train-set roster has now to be defined by the user.
7. Full upgrade of the specification of disturbances. Disturbances can now occur in sections and in any type of node (not only in connection stations), and may also include the reduction of available platform tracks in stations, speed restrictions in sections, and increasing of minimum headway. Disturbances can be applied on one or several successive periods of time, one period being defined as 15 minutes. Several disturbances may combine into scenarios.
8. Possibility to launch a cycle of consecutive simulations, each successive one using a pre-defined list of disturbance scenarios.

This shift towards greater accuracy is detrimental both to computing time needed for a simulation and to the effort a user should invest before being able to generate results. The computing time proved to be a non-issue, thanks to ever improving computer performances. To compensate for the greater user involvement, several specific actions have been taken to provide even better user-friendliness and to assist the user in exploiting and analyzing the results (Figure 1), especially as the latter became more voluminous and more complex.





FASTA II uses databases for the input data and the results (instead of binary files), user controlled functions have been seamlessly integrated; the user is assisted in managing the disturbances' scenarios, and the presentation of the results has been greatly enhanced. All those actions concur to a significantly better user experience with FASTA.

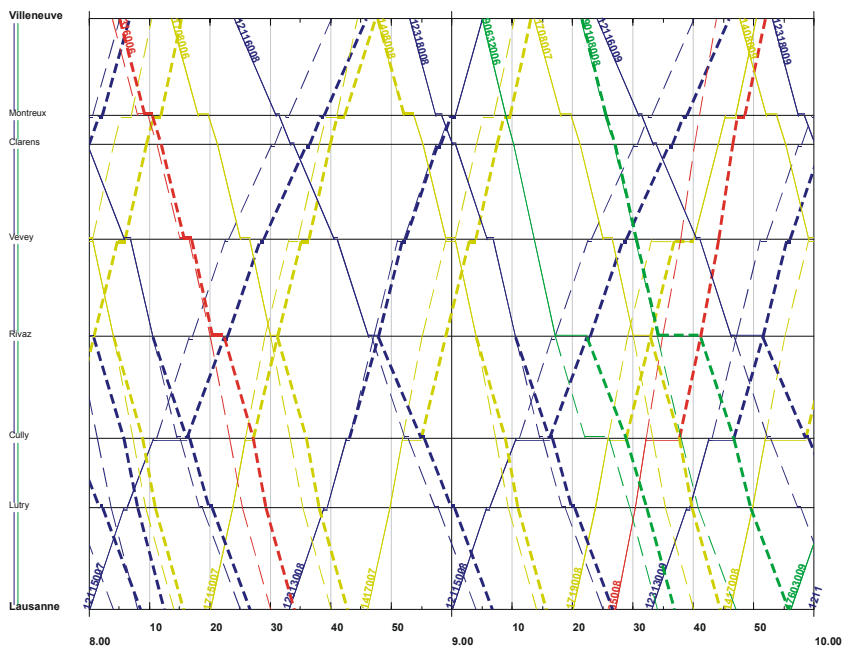


Figure 1: Simulated vs. planned.

## 7 Lessons learnt, the way ahead and conclusions

In the almost 20-years period from the initial conception of the model's objectives to the current development of FASTA II, the railway planning environment underwent a near revolution. A software, that was mainly meant to be a tool for the sketch appraisal of intrinsic stability of timetable structures, has gradually seen its use extended, and was requested to respond to significantly more complex requirements. Evolution of the planning philosophy and the scope of the model are strongly interlinked, and that may lead to completely redesign existing tools, even if they successfully achieve their initial objectives.

There is also a strong interdependence between the purposes of a software and the balance it has to provide between simplicity and accuracy. The more the objectives become complex, the more the break-even point moves away from simplicity. Although this seems intellectually obvious, this evolution, supported by both programmers' and practitioners' natural tendency to be perfectionists, may be so detrimental that eventually the tool becomes unusable for any practical purpose. The only way to control this drift is to exploit advances in

computing power, and to resist to possible but not necessarily crucial extensions and options. Improving user-friendliness also helps; a lot. Finally, developers (and users, as well) must accept the basic fact that increased versatility comes with a price: higher sophistication.

FASTA II is currently in a stable, operational state, although further developments are already foreseen [6]:

- More accurate modelling of large stations, with more sophisticated management of routes, related to platform tracks and to subsequent track assignment in sections.
- Extensive modelling of traffic control and related strategies. Current modelling of traffic control is still quite rustic, and that leads to suboptimal behaviour when complex disturbance patterns are dealt with.
- Inclusion of economical indicators, to assess economic consequences of delays and the resulting reduction of capacity [7].

## References

- [1] Noordeen M., *Stability analysis of cyclic timetables for a highly interconnected rail network*, PhD Thesis Nr 1435, École Polytechnique Fédérale de Lausanne, 1995
- [2] Curchod A., Noordeen M., Rivier R.E., FASTA: Modèle d'étude de la stabilité de l'horaire Rail 2000, *6<sup>th</sup> World Conference on Transportation Research (WCTR)*, Lyon (France), July 1992
- [3] Noordeen M., Curchod A., Hachemane P., FASTA - Rapport technique, Doc ITEP 303/3, August 1993, (revised in January 1995)
- [4] Tzieropoulos P., Les référendums suisses: une chance et une entrave, *Colloque RATP "Enjeux du transport public dans les villes européennes: quelle recherche prospective?"*, Paris, April 1998
- [5] Curchod A., FASTA (Version 3.0) - Manuel de l'utilisateur, Doc. 303/4, October 2003
- [6] Curchod A., Analysis of timetable stability and delays resulting of track maintenance works by means of simulation, *7<sup>th</sup> Swiss Transport Research Conference (STRC)*, Monte Verità / Ascona, September 2007
- [7] Putallaz Y., *Gestion stratégique de la maintenance et de la capacité d'un réseau ferré*, PhD Thesis Nr 3852, École Polytechnique Fédérale de Lausanne, 2007

