System convergence in transport models: algorithms efficiency and output uncertainty

Jeppe Rich[1]
Department of Transport, Technical University of Denmark

Otto Anker Nielsen[2]
Department of Transport, Technical University of Denmark

Transport models most often involve separate models for traffic assignment and demand. As a result, two different equilibrium mechanisms are involved, (i) the internal traffic assignment equilibrium, and (ii) the external equilibrium between the assignment model and the demand model. The objective of this paper is to analyse convergence performance for the external loop and to illustrate how an improper linkage between the converging parts can lead to substantial uncertainty in the final output. Although this loop is crucial for the performance of large-scale transport models it has not been analysed much in the literature. The paper first investigates several variants of the Method of Successive Averages (MSA) by simulation experiments on a toy-network. It is found that the simulation experiments produce support for a weighted MSA approach. The weighted MSA approach is then analysed on large-scale in the Danish National Transport Model (DNTM). It is revealed that system convergence requires that either demand or supply is without random noise but not both. In that case, if MSA is applied to the model output with random noise, it will converge effectively as the random effects are gradually dampened in the MSA process. In connection to DNTM it is shown that MSA works well when applied to travel-time averaging, whereas trip averaging is generally infected by random noise resulting from the assignment model. The latter implies that the minimum uncertainty in the final model output is dictated by the random noise in the assignment model.

Keywords: system convergence, transport models, assignment, transport demand, method of successive average.

1. Introduction

In transport modelling, one of the most fundamental equilibrium principles is the internal assignment equilibrium, where a route choice model (demand) iterates with a congestion model (supply). The complexity of this iteration scheme arises because increasing demand causes a disproportional increase in travel time, which in turn will reduce the demand. Stable convergence can be achieved by applying various averaging techniques, including the Method of Successive Averages (MSA) and variants of this method (see e.g. Sheffi, 1985). The analysis of system convergence for large scale models is closely related to model uncertainty. Firstly, for large-scale model systems, convergence is a very time-consuming process. In reality, it means that no system can be expected to be fully converged and it is therefore relevant to consider the magnitude of model uncertainty after a realistic number of iterations. In particular it is relevant to consider methods that may reduce the uncertainty through a more efficient model convergence.

---

[1] Bygningstorvet 116 Vest, 2800 Kgs. Lyngby, Denmark, E-mail: jr@transport.dtu.dk
[2] Bygningstorvet 116 Vest, 2800 Kgs. Lyngby, Denmark, E-mail: oan@transport.dtu.dk

Secondly, as state-of-the-art models often involve various degrees of randomness it is relevant to consider under which averaging strategies this randomness may propagate to the final output and, in particularly, how this problem can be minimized.

The iteration between demand and assignment, here referred to as the external equilibrium, is closely related to the mechanism of the internal equilibrium, but also different. It is similar in the sense that supply as represented by the assignment model interacts with a demand model. However, it is different in that the demand model is quite different and represents an aggregated measure of demand, typically at the origin-destination level. In any case, the importance of a unique and stable equilibrium at the system level, which can be achieved in a manageable number of iterations, is evident. If this was not possible, it would severely jeopardise the benchmark of different transport policies, as real differences in project efficiency (William and Lam; 1989) could not be separated from the uncertainty of achieving the equilibrium state. Fortunately, given that random noise in the model is (partly) controllable, convergence at the level of the system can be achieved.

Typically the feed-back between demand and supply is less volatile in the external equilibrium (due to aggregation) and the convergence will tend to exhibit a more "well-behaved" path. The fact that convergence in the external equilibrium is smoother has caused many applied models to use a simple Method of Repeated Approximations (MRA) principle[3], which in quite a few cases will perform quite well. However, the current paper shows that MRA will occasionally fail to converge or at least converge very slowly in a number of cases. This is seen in Section 4.1 (*Figure 6*) in a simulation exercise on a toy-network. It has also been experienced on a large-scale in the Danish national transport model and in the OTM model for Copenhagen (Vuk and Overgaard, 2010). In the Dutch National Model (Grol et al., 2010) it was concluded that, when using the MRA "*…it may take a long time to converge to the optimum, and possibly it may not converge at all*". The inefficiency of the MRA was also found in an activity-based model for New York (Vovsha et al., 2008).

The Dutch National Transport model represents one of the most developed transport systems in Europe (Fox et al., 2003). The convergence of the "old" Dutch model system applied a combination of MRA and "clever jump" based on approximations of the demand and supply curve, which occasionally failed to converge. Due to these convergence problems variants of the MSA were tested. Firstly, Grol et al. (2010) tested MSA with constant weight, which did not converge as well. This is not surprising as the regularity conditions in Blum's theorem (1954) are violated using this specification. In subsequent tests, reported in Grol et al., a simple MSA[4] was tested. It was shown to be difficult to find an optimal solution even when a large number of iterations were used (60 iterations). We believe Grol et al. had issues in their convergence criteria as they benchmarked against an MSA, which did not seem fully converged. Convergence for activity-based models was analysed in Vovsha et al. (2008) in an application for New York and in Bekhor et al. (2013) for the city of Tel-Aviv. In the paper by Vovsha et al., methods for obtaining convergence are classified into; i) averaging methods, and ii) enforcement methods. The latter includes re-use of random numbers in the microsimulation as well as gradual freezing of proportions of households. One of the problems in complicated activity-based microsimulation models is that the re-use of random seeds for the simulation can be difficult. This is due to the modelling of time-of-day choice and the possibility of changing trip chains dynamically. As the properties of other enforcement methods need further research, Vovsha et al. dealt with the introduced random noise by averaging on the trip tables using MSA. It was shown that the performance by averaging on trip tables was much better than averaging on LoS. Clearly, this is because they apply a static assignment. In other words, by averaging on the trip tables, which are infected by random noise, the noise is subsequently averaged out. This complies well with the

---

[3] This is also sometimes referred to as the "naive direct feedback".

[4] Based on averaging on trips which we in the current paper recommend not to use.

finding for the Danish national model as presented in this paper. We apply a simpler demand model, which although based on a micro-simulated population, is based on fixed a random seed and hence can be made deterministic between runs as we re-use the seed. The assignment model, on the contrary, is a stochastic user equilibrium (SUE) model with random effects for links and routes, which cannot re-use the random seed. We demonstrated that averaging on LoS is a powerful tool as it facilitates a smoothly and consistent convergence.

Hence, an important lesson is that if you have uncontrolled random effects in both the demand and the assignment model then the system convergence will never be better than the noise-floor of the system. The noise-floor is dictated by the model with the lowest degree of random noise. It is therefore quite important to be able to either re-use the random seed for the demand model or the assignment model so that random noise does not occur at both ends. Whether it is better to have uncontrolled noise in demand or assignment will depend on the modelling objective and possibly other characteristics such as the network layout and the zone system.

Another strong motivation for being concerned with the convergence of the external equilibrium relates to the computational effort. As the external equilibrium loop involves running a complete assignment model combined with a complete demand model, iterations are much more costly than for the inner loop. This does not justify a simple iteration scheme for the sake of simplicity. As only 5 to 10 iterations may be possible in practice within a reasonable calculation time for many large-scale models[5], it is important that these are spent wisely.

The paper thus considers the convergence of the external loop and discusses variants of the MSA algorithm. Section 2 firstly introduces the fix point problem and algorithms in general in a theoretical perspective. Section 3 considers system convergence in practise and lines up different methods to be tested. Section 4 describes simulation experiments where solution methods are first tested on a "toy-network" with a stylized demand model and subsequently on a large scale by investigating convergence performance in the Danish National Model. Summary and conclusions are offered in Section 5.

2. Introduction to fix point algorithms

The main issue in large-scale transport models is to find the equilibrium between a supply (assignment) model and a demand model. This equilibrium solution is a fixed point and the problem falls within the more general family of mathematical fixed-point problems. Results on existence and uniqueness for a broad class of assignment models were presented in Fisk (1980) and for assignment with elastic demand in Cantarella (1997). In the current paper, where the demand models are of the nested logit type and with a fixed demand profile between iterations due to the re-use of random seeds in the model, the uniqueness and existence is well covered by these papers although demand and supply are iterated externally.

Let demand be represented as a continuous vector function $f(x) : S \subseteq R^N \rightarrow S$ and S a non-empty, compact and convex set. According to Brouwer theorem (Agarwal et al., 2001) it has at least one fixed-point $x^* = f(x^*)$ (existence). If it can also be proved that a fixed-point exists and is unique, then this unique fixed-point may be found through many algorithms, whose general specification can be written as (Agarwal et al., 2001);

$$x_k = x_{k-1} + M_k \left(f(x_{k-1}) - x_{k-1} \right) \tag{1}$$

The algorithm is based on a starting solution $x_0 \in S$, and a properly defined matrix M_k (see among many others Kelley, 1995).

---

[5] The 60 iterations applied in the Dutch experiment as referred in Grol et al. (2010) would imply a runtime of 25 days in the DNTM.

The method of repeated approximations (MRA), which resembles the direct feedback approach, is given by $M_k = I$. However, it can be proven to converge only for contractions or for strictly non-expansive functions.

The Newton method is based on $M_k = [I - Jac(f(x_{k-1}))]^{-1}$ and will usually converge fast provided that the starting solution is close enough to the searched fixed-point.

The Broyden method (refer to, e.g. Agarwal et al., 2001) is a kind of secant approximations, where matrix M_k is updated at each iteration, from $M_k = [I - Jac(f(x_0))]^{-1}$. It gives some computation advantages with respect to the Newton method since derivatives need to be computed only at the first iteration. However, from a practical large-scale perspective neither the Newton method nor the Broyden method are feasible due to the dimensionality of the underlying models and matrices. The methods require matrix inversion of very large matrices, which is computationally cumbersome.

The method of successive averages (MSA) is given by $M_k = a_k I$ with $a_k = 1/k$. This form was suggested in the seminal work by Robbins and Monro (1951), which proved convergence under two regularity conditions. These conditions were later softened by the Blum theorem (Blum, 1954) and by an extension in Cantarella (1997) with weakened regularity conditions. If the function is computed by means of Monte Carlo simulation which only provides an unbiased estimation of the value, only almost sure convergence can be considered. Since all intermediate solutions in the sequence are feasible, algorithms based on MSA are often called *feasible*, and the current solution may be considered an approximation to the searched equilibrium flows. These algorithms are also called *simple* since they only require computation of all the involved functions and will not require computation of derivatives and need no matrix algebra during iterations (Cantarella and Cascetta, 2009).

The "success" of the MSA (and variants of the MSA) is due to the "successive averaging" forming a contraction principle. Although it is not a contraction in a strict mathematical sense, it is a contraction principle which will lead to convergence as the number of iterations goes to infinity. The principle of MSA is illustrated in *Figure 1* below. As seen MSA will move towards the equilibrium, whereas MRA is unstable and moves away from the equilibrium. Depending on the shape of the curves, MRA may move away from the solution, may be cyclic unstable, or may move towards the equilibrium.
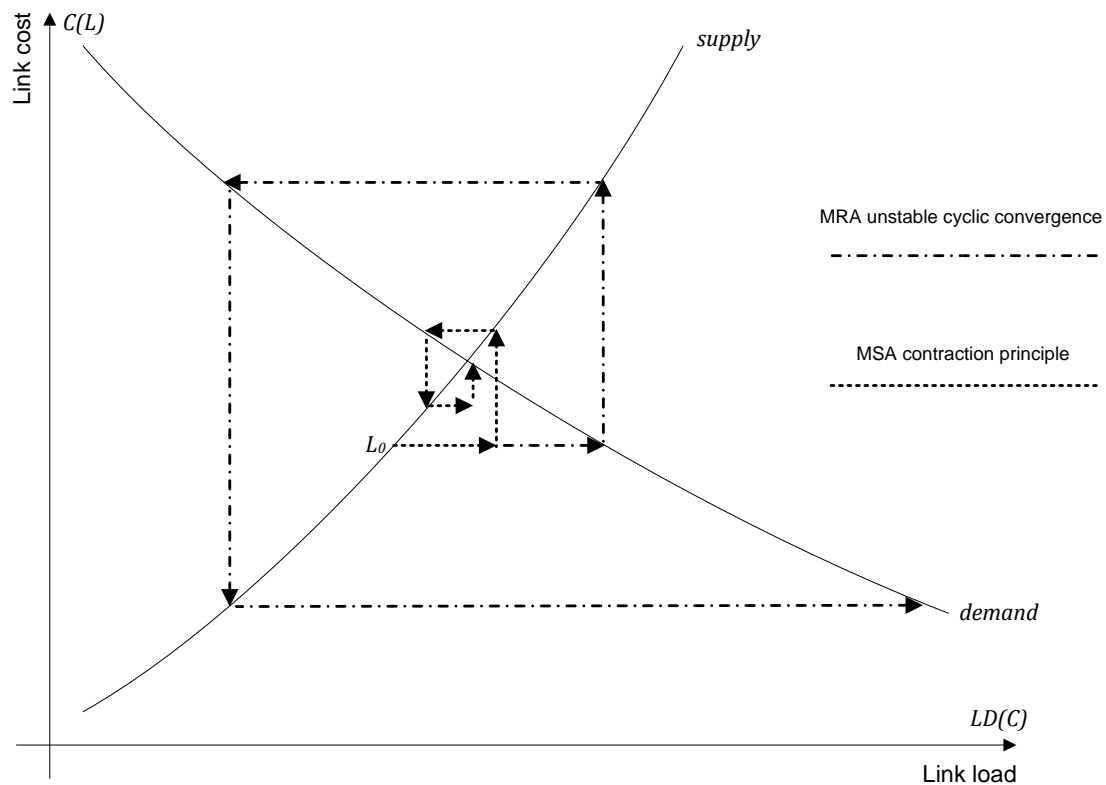
Figure 1. Illustration of unstable solution based on the MRA iteration scheme and the MSA contraction principle

A general finding in many practical applications is that the MSA principle is particularly efficient in the start of the sequence, whereas the speed of convergence tends to be relatively slow as the algorithm moves towards the equilibrium. This is particularly true if the starting solution is very different from the equilibrium as the "noise" from the first iterations will be inherited in the entire path towards the equilibrium. This is especially the case when transport models are used for long-term forecasts and/or project appraisals, where the input variables are changed compared to the data in the calibration year.

3. System convergence in practise

The equilibrium between the demand model and the assignment model is typically solved by an outer loop as illustrated below;

> **Step 0:** Define base-line OD matrices.
>
> **Step 1:** Calculation of Initial Level-of-service (traffic assignment based on base-line OD matrices).
>
> **Step 2:** Demand model, based on Level-of-service from Step 1.
>
> **Step 3:** Generation of OD matrices based on the demand model output (Step 2).
>
> **Step 4:** Calculation of Level-of-service (equilibrium inner-loop level-of-service based on demand matrices from Step 3).

Step 5: Iterate Step 2-Step 4 until convergence is obtained.

The MSA works by either averaging OD matrices (Step 3) or by averaging level-of-service matrices (Step 4). If averaging is applied to OD matrices it basically averages matrices generated in Step 3 from iteration to iteration. If averaging applies to level-of-service, it is based on averaging of the outcome of Step 4 from iteration to iteration.

An important issue in obtaining system convergence is that the use of different exogenous level-of-service attributes in the demand and route choice could influence the convergence. For instance, it is common to use fixed value-of-time estimates in both demand and assignment and these could well be different as they represent different choice situations. If these values are different it could introduce a flip-flop between the demand and supply and cause convergence to slow down. An example hereof is if the "mark-up" for congestion compared to free-flow time is 2 in the route choice model and 1 in the demand model. In this case, the route choice will tend to react more aggressively with respect to congestion than the demand model and produce trips which will be longer (de-tours) in order to circumvent congestion. The demand model will on the other hand counteract this process at the level of the matrices, as congestion is less costly. So it may lead to an assignment that seeks to avoid congestion and a demand model that "seeks" congestion which will slow down the convergence.

3.1 Averaging methods

Applying the MSA in the external loop is straightforward and implies that demand x_k (or supply) at iteration k is represented as a moving average. Hence, the final estimate of $\tilde{x}_k$, is represented as

$$\tilde{x}_k = a_k x_k + (1 - a_k) x_{k-1} \tag{2}$$

In the original paper by Robbins and Monro (1951) it was suggested that $a_k = 1/k$. With this parameterisation of a_k, the MSA puts most weight on the "history" and less weight on the current iteration. As discussed in Section 2 this generally works well if the starting solutions are not too noisy. However, if the starting solution is bad the inherited noise from the first iteration will mess up the performance of the MSA and the convergence will be slow. To deal with this issue several methods have been proposed. The underlying idea of most of these methods is to define a sequence a_k that conforms to the regularity conditions (Blum, 1954) stating that $\sum_k a_k = \infty$ and $\sum_k a_k^2 < \infty$ and where the weight of the first part of the iteration process is gradually damped.

A simple idea was put forward by Cascetta and Postorino (2001) who suggested resetting the iteration history at some points in the iterative process. If φ is the number of MSA iterations before reset, the resetting can then be accomplished by defining a new iteration index $\tilde{k}$ equal to;

$$\tilde{k} = mod(k, \varphi) \tag{3}$$

$$if\ \tilde{k} = 0\ then\ \tilde{k} = \varphi \tag{4}$$

The MSA is then simply defined with step-size $a_{\tilde{k}} = 1/\tilde{k}$, which will produce a repeated sequence $1, 2, \ldots, \varphi, 1, 2, \ldots \varphi, \ldots$

Cascetta and Postorino (2001) suggested resetting the history for every 5 steps, however, the optimal choice is strongly network-dependent and will also depend on the acceptable precision level for the final iteration and the number of iterations that are run. It should also be said, however, that in order for the resetting approach to be consistent with the regularity conditions (Blum's theorem), there should always be a point from where the reset is no longer used. Moreover, choosing a value of φ which is too small can be dangerous and it is generally not recommended to use values below 5.

Polyak and Juditsky (1990) introduced an alternative step-size for the MSA equal to $a_k = \frac{1}{k^{2/3}}$. Compared to the original suggestion by Robbins and Monro (1951), this specification put more weight on the newest iterations. It has also been common to use $a_k = \frac{1}{k^{1/2}}$ although it is does not satisfy the Blum theorem.

A weighted MSA approach was considered in Liu et al. (2009) where $a_k(d)$ is given by

$$a_k(d) = \frac{k^d}{1^d + 2^d + 3^d + \cdots + k^d} \tag{5}$$

where $d \geq 0$. Clearly, for $d = 0$ the ordinary MSA emerges, however, as d increases more emphasis is put on the latest iterations. It is easy to see that $a_k(1) = \frac{2}{k+1}$ and $a_k(2) = \frac{6k}{(k+1)(2k+1)}$.

The weighted MSA series satisfy the regularity conditions as

$$\frac{\pi^2}{6} = \sum_k^\infty a_k(0)^2 > \sum_k^\infty a_k(0.5)^2 > \sum_k^\infty a_k(1)^2 > \ldots > \sum_k^\infty a_k(d)^2 \tag{6}$$

and

$$\infty = \sum_k^\infty a_k(0) < \sum_k^\infty a_k(0.5) < \sum_k^\infty a_k(1) < \cdots \sum_k^\infty a_k(d) <, \forall d > 0 \tag{7}$$

3.2 Intersection approach

In relation to the work with the Danish National Model (Rich, 2010b) a slightly different approach was tested, which we will refer to as the intersection method (Paulsen, 2013). This method approximates supply and demand by linear functions and these functions are then solved for their intersection points. Each OD pair is considered separately, and cross-effects between the different OD pairs are ignored. Hence, a demand response does not include level-of-service changes from adjacent OD pairs.

In the following two mappings are considered;

- $D(t_k)$: Mapping of demand as a function of LoS t_k at iteration k. Here demand is represented as the accumulated car demand OD matrices.

- $t(D_k)$: Mapping of LoS as a function of demand at iteration k. Here LoS is represented at the car LoS matrix.

The process applies the following steps:

Step 0: Set the iteration number $k = 0$. Start out with an initial level-of-service matrix; t_{k-1} (this matrix may define trips between zones $\{i, j\}$ as well as possible trip purpose categories).

Step 1: Evaluate the following point estimates $D_k = D(t_{k-1}), t_k = t(D_k), D_{k+1} = D(t_k)$ and $t_{k+1} = t(D_{k+1})$. This requires two evaluations of the demand model and two evaluations of the assignment, which leads to four data points $\{D_k, t_{k-1}\}, \{D_k, t_k\}, \{D_{k+1}, t_k\}$ and $\{D_{k+1}, t_{k+1}\}$. In the following, the demand curve is approximated from $\{D_k, t_{k-1}\}, \{D_{k+1}, t_k\}$ and the supply curve from $\{D_k, t_k\}, \{D_{k+1}, t_{k+1}\}$.

Step 2: Define (for each OD pair) the linear system where demand and supply is approximated by the four most recent data points.

$$D = q_1 + b_1 t \tag{8}$$

$$t = q_2 + b_2 D \tag{9}$$

where

$$q_1 = \left(1 - \left(\frac{D_{k+1}-D_k}{t_k-t_{k-1}}\right)\right)D_{k+1}, \; b_1 = \left(\frac{D_{k+1}-D_k}{t_k-t_{k-1}}\right) \tag{10}$$

$$q_2 = \left(1 - \left(\frac{t_{k+1}-t_k}{D_{k+1}-D_k}\right)\right)t_{k+1}, \; b_2 = \left(\frac{t_{k+1}-t_k}{D_{k+1}-D_k}\right) \tag{11}$$

Step 3: Solve the linear system in Step 2. The solution is given by:

$$t^* = \frac{q_2 - b_2 q_1}{1 - b_1 b_2} \tag{12}$$

$$D^* = q_1 + b_1 \frac{q_2 - b_2 q_1}{1 - b_1 b_2} \tag{13}$$

Step 4: Let $k = k + 1$. Calculate $D_{k+1} = D(t^*)$ and subsequently $t_{k+1} = t(D_{k+1})$.

If an MSA function with sequence $a_k(d)$ is applied on the point sequence $t_{k+1}, t_k, \dots, t_1$ we get

$$\tilde{t}_{k+1} = a_k(d)t_{k+1} + (1 - a_k(d))t^* \tag{14}$$

Use the MSA point for $\tilde{t}_{k+1}$ to calculate the corresponding point of demand $\widetilde{D}_{k+1}$.

$$\widetilde{D}_{k+1} = D(\tilde{t}_{k+1})$$

Step 5: If $\|\tilde{t}_{k+1} - \tilde{t}_k\| < \varepsilon$ for all OD pairs, stop. Otherwise let $k = k + 1$ and define $D_{k+1} = D(t^*)$ and $t_{k+1} = t(D^*)$. Go to Step 2.

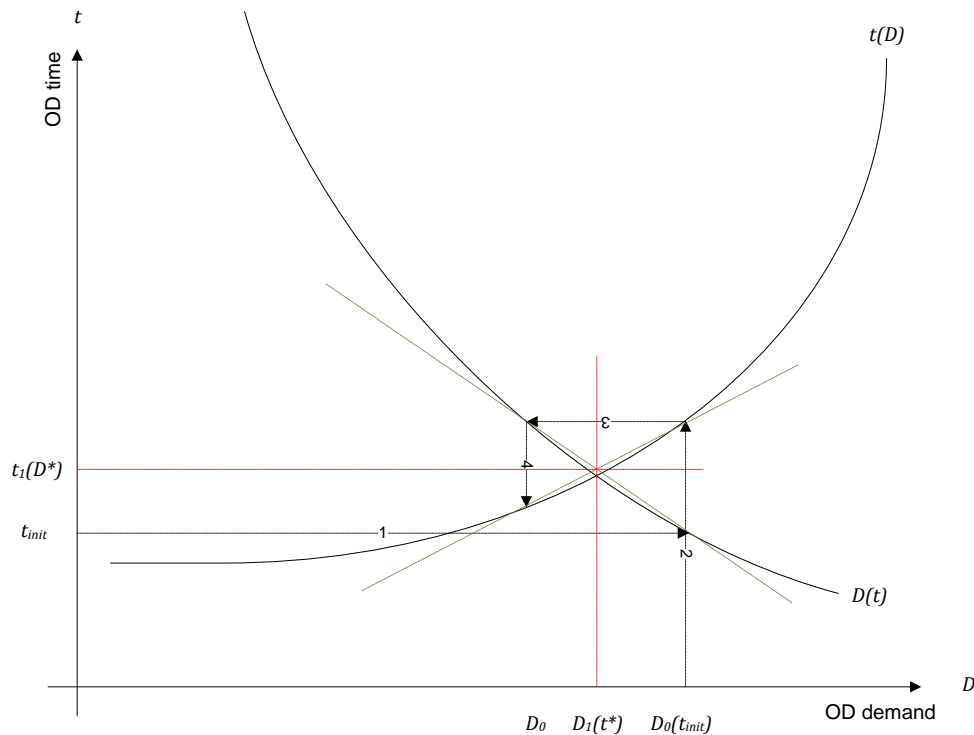A graphical illustration of the process is illustrated in Figure 2 below.

Figure 2. Illustration of the Intersection Method

As seen from Figure 2 the intersection method performs very well when applied to a one-dimensional case. Only 2-3 iterations are needed to get very precise approximations.

A problem when looking at a single cell is however that the $D(t)$ and $t(D)$ functions may not be unique due to substitution effects. Changes of nearby cells may have a second-order effect on the cell under consideration. If in one iteration, cell $\{i, j\}$ attracts much traffic, $\{i, j - 1\}$ may lose

traffic correspondingly. To build some robustness into the algorithm we introduce the MSA averaging in equation (14).

Generally the intersection method seems to be somewhat similar to what was used in the Dutch model according to the description provided in Grol et al. (2010) and the final conclusion seems to be similar as well.

Although the MSA averaging in equation (14) helps, it does not fully resolve the problem of correlation and we have experienced positive failure rates for the intersection method for most settings in the sense that it failed to converge to an equilibrium in 1000 iterations. The paper does not allow a detailed elaboration on the rigorous simulation exercise that was carried out (Paulsen, 2013). However, below we present the most important conclusions and invite the reader to consider Appendix A where a small subset of the results is presented. For additional information readers should refer to (Paulsen, 2013).

For the intersection method, the following was concluded:

- Excellent performance in one-dimension.

- For N-dimensional systems.

    o When used without MSA a significant share of the test cases did not converge and when convergence was experienced the rate of convergence was slow compared to a weighted MSA.

    o When used with MSA and proper weighting the failure rate could be reduced significantly and may even be zero for reasonably soft convergence criteria. Still the rate of convergence was slow compared to a weighted MSA.

Based on these findings it was concluded that the intersection method, although tempting at first sight, could not outperform the weighted MSA and (more importantly) could not guarantee stable convergence. Due to this, the idea will not be considered further[6].

4. Simulation experiments

In the following section, we test the various MSA techniques mentioned in Section 3.1 in two different model settings

- A small/medium sized "toy-network" consisting of 9 zones and a total of 3,804 routes

- A large network for the whole of Denmark consisting of 907 zones and more than 30,000 road links.

The strategy was to use the toy-network to "screen" for good solution candidates and then test the most successful ones in the Danish National model.

4.1 Analysis of "toy-network"

The lay-out of the "toy-network" is illustrated in Figure 3 below.

---

[6] We are aware of slightly similar ideas by Lv et al. (2010) on a two-directional approach which is shown to converge. Whether this approach is applicable to large-scale implementation and can compete with the efficient weighted MSA is an interesting future research topic.
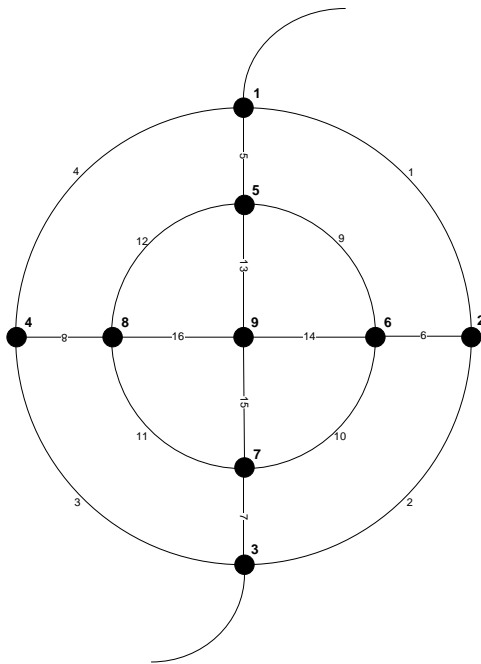
Figure 3. Layout of toy-network

The network is represented by 9 nodes (or zones), which by elimination of inter-zone traffic gives 72 OD pairs connected by 16 arcs. There is no assumption of limitations in the network (e.g., closed links or one-way links) and the total combination of routes is 3,804 representing 22,404 connected arcs. As a result, the average route consists of 5.88 arcs.

Different OD matrix structures were generated randomly to reflect different OD flow patterns and congestion levels.

The inner loop in the simulation experiment used a full stochastic loading on all routes for each loop in the assignment (approximately 50 routes per OD pair). This is in contrast to a normal assignment model, where routes are sampled sequentially in the MSA loop.

The simulation experiment did not apply link-based random effects. This is usually applied to improve performance and robustness of the inner loop; however, in this case with a small network and very fast convergence (due to the full stochastic loading) it was not needed.

In all of the simulation schemes on the toy-network we applied a simple demand and supply setup. Demand for car transport for a given OD pair $\{i, j\}$ at iteration k is for simplicity notated as $D(t_k)$. Mathematically it was represented as a logit model[7] for the choice of mode m scaled by a fixed OD matrix $T_{i,j}^0$ as illustrated below in (15)

$$D(t_k) = D(t_k | m = car, i, j) = T_{i,j}^0 \frac{exp(\vartheta_{m=car} + \theta t_{m=car}(i,j))}{\sum_{m\prime} exp(\vartheta_{m\prime} + \theta t_{m\prime})} \tag{15}$$

The supply $t(L_k)$ function was represented as a traditional BPR formula (Highway Research Board, 1965) as seen below in (16) where L_k represent the link-load at iteration k.

$$t(L_k) = t_0 \left(1 + \alpha \left(\frac{L_k}{c}\right)^{\beta}\right) \tag{16}$$

These values of these functions were then updated iteratively to investigate convergence failure and rate of convergence. In the following section we consider the simulation results.

---

[7] A similar setup was used in the toy-network experiments presented in Section 4.1.

Non-convergence of the outer loop

The first and very obvious issue is whether it is necessary to consider convergence problems in the outer loop at all? As many models have applied MRA iteration schemes with success it is a question whether it is possible to construct a counter example? However, it turned out to be straightforward to generate many examples where the external loop diverged from the equilibrium solution. If the slopes of the demand and supply curves were moderately flat, usually the performance of the MRA was quite good and sometimes even better than the MSA. However, as the slope increased the divergence was amplified and at some point it became cyclically unstable.

Rather than illustrating this in a separate figure we refer to Figure 6, which indicates a nearly cyclically unstable MRA convergence and, in all cases, very slow convergence.

Benchmark of averaging methods

Figure 4-6 presents a series of convergence benchmarks for the MRA and five averaging methods. Unfortunately, even for a small-scale problem as represented by the toy network it is very difficult to find the exact solution. As a result, we have approximated the system equilibrium by iterating the system until stability. We used more than 300 iterations and tested the performance using different averaging methods. The convergence criterion is below the double-precision floating-points in SAS. For each method, we evaluated the mean % deviation $\frac{1}{N}\sum_{ij}\left(\frac{T_{ij}-T_{ij}^*}{T_{ij}^*}\right)$ from the approximated equilibrium trip matrix T_{ij}^* to facilitate a visual inspection of the convergence scheme of the different methods as presented in *Figure 4-Figure 6*.

A first and interesting observation is that although the MRA failed to converge in certain cases it was quite efficient compared to many averaging methods in other cases. The main explanation for this is that many averaging methods will tend to converge slowly when the slope of the supply curve up to the point where it crosses the demand curve is relatively flat.
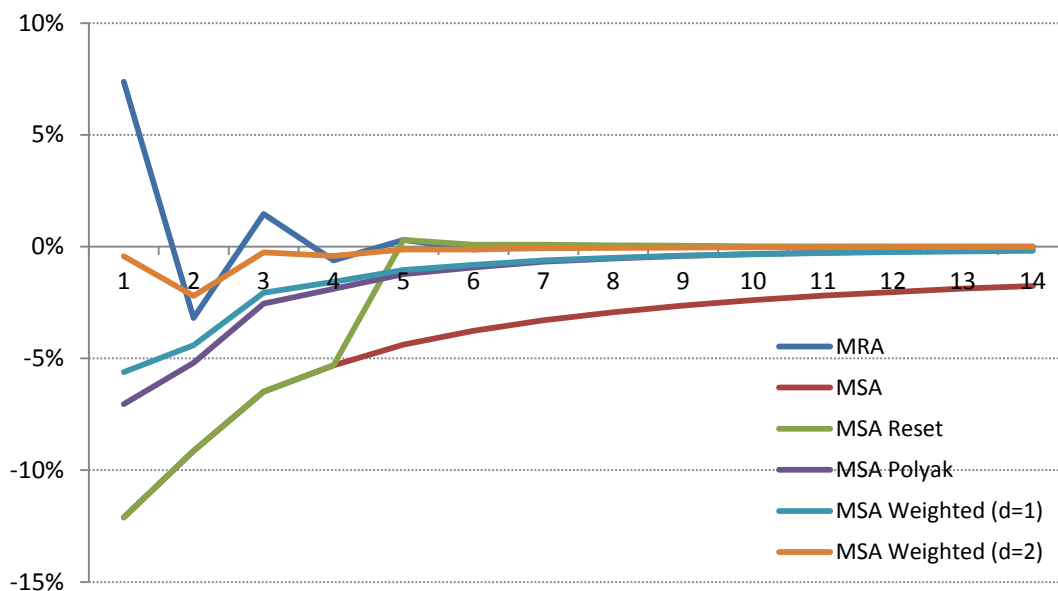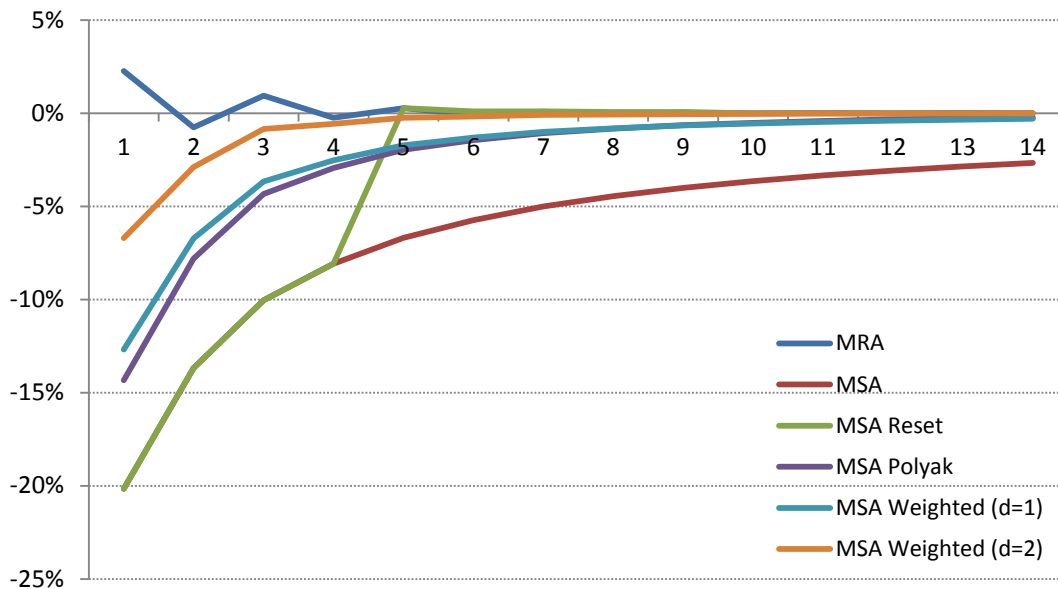

Figure 4. Convergence performance measured in terms of mean % deviation of trips from equilibrium for six methods (iterations in the horizontal axis); "normal" congestion

Figure 4 represents a situation with normal congestion. The BPR parameters are given by $\{\alpha, \beta\} = \{0.5, 2\}$, logit demand parameter (route choice) -0.5 and logit parameter for demand model -0.1.

In terms of absolute deviation, the weighted MSA with $d = 2$ performs quite well for only 3 iterations. However, the practical convergence rate is found to be close to 1 meaning that the relative convergence improved linearly. The MSA with reset is also efficient; however, it requires that the reset point is passed (in this case 5 iterations). In fact, from an infinitesimal point of view, the MSA with reset is the best performing method in this example for 10+1 iterations.

Figure 5 below illustrates a situation where the congestion level is higher. For the BPR $\{\alpha, \beta\} = \{0.5, 2.5\}$, the logit demand parameter (route choice) is -0.5 and the logit parameter for demand model is -0.13. As can be seen, the performance declines for all of the methods.

Figure 5. Convergence performance measured in terms of mean % deviation of trips from equilibrium for six methods (iterations in the horizontal axis) "aggressive" congestion

It is however interesting to see that the weighted MSA with ($d = 2$) is generally relatively efficient for even a small number of iterations. The MSA with reset tends to be efficient after the reset point of 5 iterations.

Figure 6 illustrates a situation where the congestion level is "extensive". The BPR parameters are now given by $\{\alpha, \beta\} = \{0.5, 3\}$, and the logit demand parameter (route choice) is -0.5 and logit parameter for demand model is -0.2.
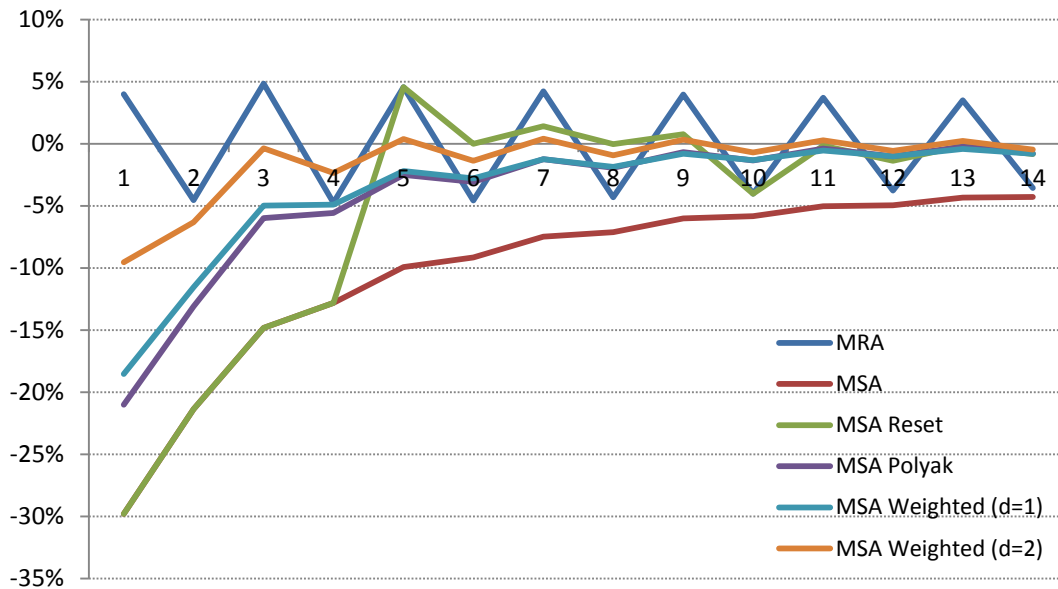
Figure 6. Convergence performance measured in terms of mean % deviation of trips from equilibrium for six methods (iterations in the horizontal axis); "hyper" congestion

Generally, the performance of the MSA as well as the MSA with Polyak step size and weighted MSA ($d = 1$) declines and tends to be unacceptable for a small number of iterations. However, the weighted MSA ($d = 2$) and MSA with reset perform well. The best performing method is again the weighted MSA with $d = 2$. Values beyond 2 were shown to be too aggressive. The MSA with reset is again quite good, although "jumps" can occur, especially around reset points. Also, note the oscillating pattern for the MRA and the extreme slow convergence of the MSA. To supplement the visual inspection we present the Root Square Error (RSE) between the approximated equilibrium point D^* and the solution at iteration k given by D_k. This is shown in Table 1 below.

Table 1. Convergence performance measured as norm difference for different methods and congestion profiles

	RSE_1	RSE_2	RSE_3	RSE_4	RSE_5	…	RSE_{15}
Normal congestion							
MRA	38.454	13.166	5.204	2.220	0.983	…	0.002
MSA	38.454	25.223	18.360	14.238	11.552	…	4.170
MSA Reset	38.454	25.223	18.360	14.238	11.552	…	0.006
MSA Polyak	38.454	21.895	13.658	9.020	6.229	…	0.661
MSA Weighted (d=1)	38.454	20.970	12.882	8.525	5.973	…	0.898
MSA Weighted (d=2)	38.454	17.690	9.491	5.545	3.443	…	0.222
Aggressive congestion							
MRA	48.194	13.928	5.826	2.942	1.719	…	0.026
MSA	48.194	28.503	20.100	15.157	12.296	…	4.405
MSA Reset	48.194	28.503	20.100	15.157	12.296	…	0.015
MSA Polyak	48.194	23.789	13.976	8.694	6.019	…	0.627
MSA Weighted (d=1)	48.194	22.521	12.949	8.072	5.669	…	0.839
MSA Weighted (d=2)	48.194	18.285	8.823	4.774	3.012		0.199
Hyper congestion							

MRA	54.362	16.568	13.149	9.686	9.966	…	4.378
MSA	54.362	26.460	20.554	14.276	12.668	…	4.445
MSA Reset	54.362	26.460	20.554	14.276	12.668	…	0.470
MSA Polyak	54.362	20.336	14.656	7.553	6.926	…	0.732
MSA Weighted (d=1)	54.362	18.841	13.567	6.987	6.465	…	0.957
MSA Weighted (d=2)	54.362	15.007	10.496	4.723	4.650	…	0.463

MSA with reset

As indicated in the simulation experiments, the MSA with reset generally performs very well. However, it requires as a minimum that the number of iterations exceed the first reset point and at best 2 or 3 reset points as "jumps" can occur as shown in *Figure 6*.

An obvious question is which reset point is optimal? Although the question cannot be answered in any precise way as it is network-dependent, simulation can give indications on how choice of reset point affects the iteration history.
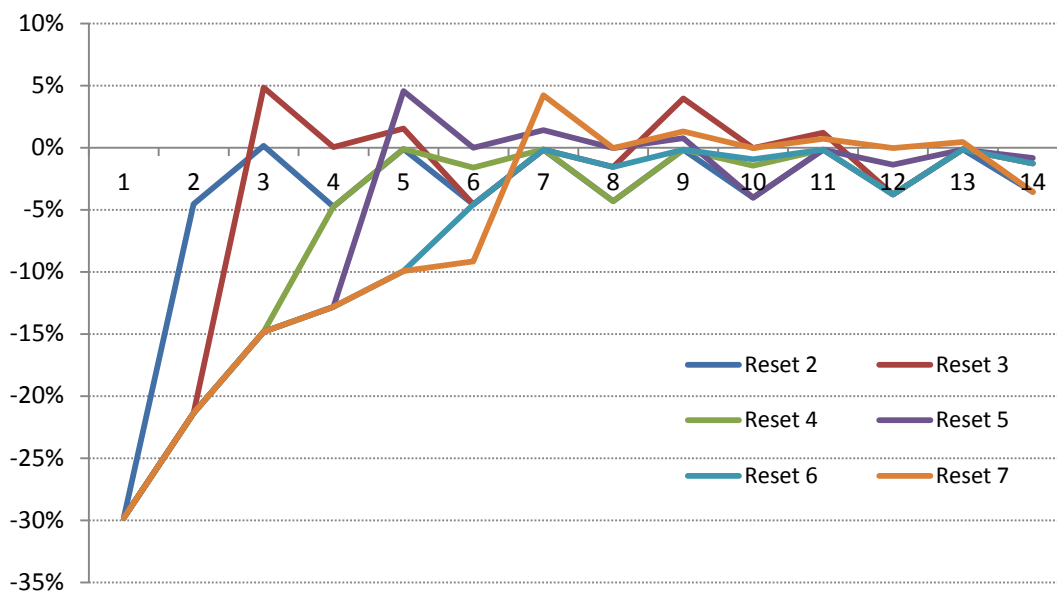
Figure 7. Convergence performance measured in terms of mean % deviation of trips from equilibrium of MSA reset with different reset points; "hyper" congestion

Figure 7 suggests an interesting finding, namely that the choice of reset point, whether it is 2, 3, 4, 5, 6 or 7 are not so important even in a "hyper" congestion regime. It is thus important to choose the correct stoppage point after the reset point. Given that iterations are costly, the optimal point seems to always be 1 plus the reset point. Hence, for a reset point of 2 the algorithm should stop after the third iteration, for a reset point of 5, it should stop at 6, etc. This finding was consistent for a range of different random OD matrices. Based on the various Monte-Carlo runs (as also indicated in *Figure 7*) the best reset point was 5. However, to fully investigate whether this is due to "peculiarities" in the data or points to a more general finding needs to be assessed in more detail.

Weighted MSA

Another interesting question is whether it may be efficient to further decrease the emphasis on "history" in the weighted sequence and increase the value of d to 3 or 4. Again the "hyper" congestion regime from Figure 9 was used, since it provided the greatest challenge to the algorithms.
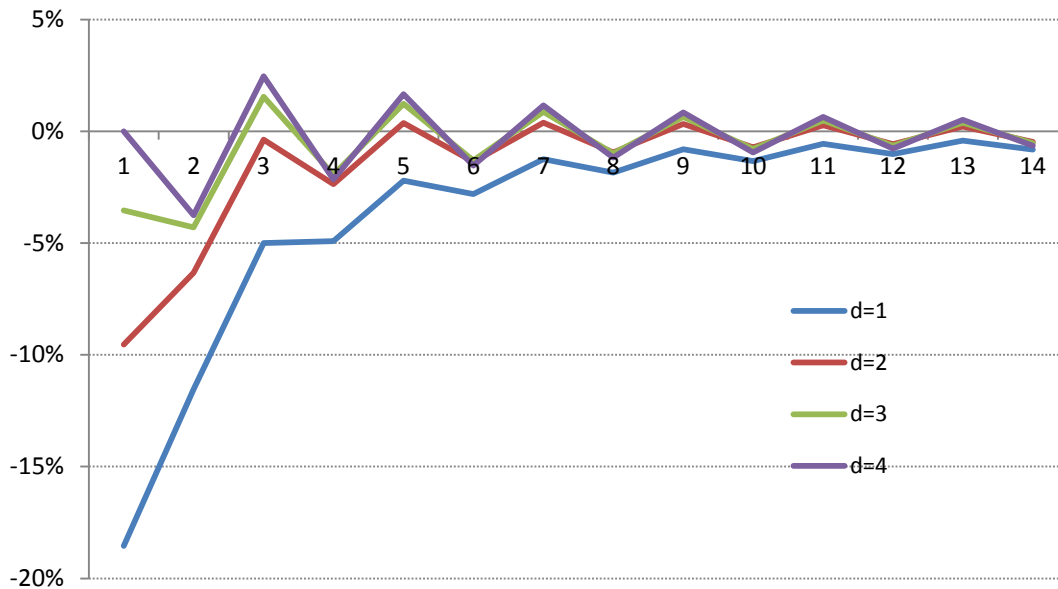
*Figure 8. Convergence performance measured in terms of mean % deviation of trips from equilibrium of weighted MSA with different **d**-values; "hyper" congestion*

As can be seen, there is an optimal value somewhere between 1 and 4. However, $d = 2$ is the optimal choice in this case.

4.2 Analysis on the Danish National Transport Model

The following analysis considers convergence performance for the Danish National Transport Model[8] (DNTM). The model operates on Denmark as a whole and includes 907 Danish zones and more than 300 foreign zones to take account of traffic in and out of Denmark.

The passenger model[9] is based on forecasts of the entire Danish population, which is grouped into synthetically generated households fed into the demand model. In total the model operates with more than 5,000 different person classes divided by age, gender, labour market association, income and family status per zone. All individuals are then subsequently grouped into household entities in a micro-simulation process which includes "spouse-matching" and simulation of kids. This produces a long list of individuals, which is fed into the demand model in which car ownership, trip frequency, destination choice and choice of mode are calculated. In the discrete choice framework, which is of the nested logit type, more than 15 separate models for different transport segments and car ownership are processed. This in turn produces a list of trips which is then fed into the assignment model. In the demand model it is possible to re-use the random seed in the simulation, which implies that there are no random noise from iteration to iteration that can be attributed to the demand model.

The assignment model is a stochastic user equilibrium (SUE) model which operates on a network of more than 30,000 links and 15 passenger user classes and 6 freight vehicle classes. The stochastic part of the model is a Mixed Probit model, where the error term is simulated (as in Sheffi, 1985, Appendix A.2) but where the model in addition used a logarithmic normal distributions of the value of times which is solved by simulation. Freight demand is determined in a separate freight model, which provides demand matrices (vans and various truck types) as input to the assignment model as well.

---

[8] The tested version is version 1.07. This official version will be version 1.08, which will be made official during the autumn of 2014.
[9] Refer to Rich et al. (2010b) and (2010c) for an overall description of the Danish Transport Model.

The DNTM model applies a pivot-point principle where the model is pivoted relatively to a baseline matrix for 2010. Hence, if we apply assumptions very close to those for 2010 in terms of population and network characteristics we would reproduce the 2010 solution and the convergence would be good. Hence, in order to "challenge" the model we apply a 2040 population with a 25% reduction in travel cost for cars. This causes a rather large exogenous chock to the system with large re-bound effects.

Measurement of convergence

When testing different iterative root-finding algorithms such as MRA and weighted MSA with different weights, it is important to be able to benchmark the convergence speed of these different algorithms. The problem is obviously that the true equilibrium is unknown and cannot be approximated infinitesimally close as for the toy-network. The only thing we know, based on Blum's theorem, is that convergence will be achieved given that the regularity conditions in equation (6) and (7) are fulfilled (note that this is not the case for MRA in general).

A common measure of convergence for fixed point problems is $\|D(x) - x\|$. If averaging is carried out on trips (hence, x represents a trip matrix) it can be deducted from the output of the successively iterated OD matrices and the MSA weight. From equation (17) we have that $x_k = x_{k-1} + a_k (D(x_{k-1}) - x_{k-1})$. By rearranging, we get

$$\|D(x_{k-1}) - x_{k-1}\| = \frac{1}{a_k} \|x_k - x_{k-1}\| \qquad (17)$$

As can be seen, the MRA iteration scheme implies that $\|D(x_{k-1}) - x_{k-1}\| = \|x_k - x_{k-1}\|$. If a_k is different from one, however, we need to weight the right hand side to account for the memory. If averaging is carried out based on LoS as in Section 0 we can calculate RSE in LoS-space at each iteration in a similar way. However, if applying LoS averaging we cannot easily examine the RSE in trip space and because of this we will examine convergence performance in the LoS space.

As a different measure of convergence performance, we have tested whether it was possible to use an approximate convergence rate as suggested in Senning (2007). However, the experiments were not unambiguous and the variation in the convergence scheme from iteration to iteration influenced the rate. We believe the approach could be valuable if more iterations could be produced. In that case we suggest to calculating the average convergence rate. In all of our experiments, we experienced a convergence rate close to 1 indicating linear convergence.

MSA applied to trips

In the first experiment we applied MSA and MRA to OD trip matrices. That is, the successive averaging is applied to trip matrices from iteration to iteration. However, the problem with this strategy is the presence of random noise from the assignment model. It was found that the variation in the model was always above a certain noise threshold (amounting to ~ 1% deviation of trips between iterations). The problem is that when trips are used as basis for the MSA (or MRA), the stochastic element of the assignment is "renewed" for every iteration, which causes the model to converge, not to a unique equilibrium, but to any solution below the noise threshold of the model. If the random seed of the assignment model could be re-used this would eliminate the problem, however, this is complicated and generates a lot of computational overhead and requires methodologies for how to deal with possible new routes. A solution would be to choose a simpler deterministic assignment; however this would result in a much too simplified network loading. Due to these problems, it was decided to consider MSA on the basis of level-of-service rather than on trips. It was also decided not to use MRA due to possible convergence problems as illustrated for the toy-network and the poor performance when averaging on trips. The MSA averaging on LoS is considered below, and tests with different d values are carried out.

MSA applied to LoS

The averaging of LoS was carried out by weighting of the car LoS matrix. More specifically, LoS was defined as the perceived travel time between OD pairs which is the sum of free-flow and congested time (un-weighted).

We measured convergence by looking at $\|S(\boldsymbol{t}_k) - \boldsymbol{t}_k\|$, the RSE, between the LoS function $S(\boldsymbol{t}_k)$ and the LoS solution at iteration $\boldsymbol{t}_k$. This term can be calculated from the model output in a straightforward manner.
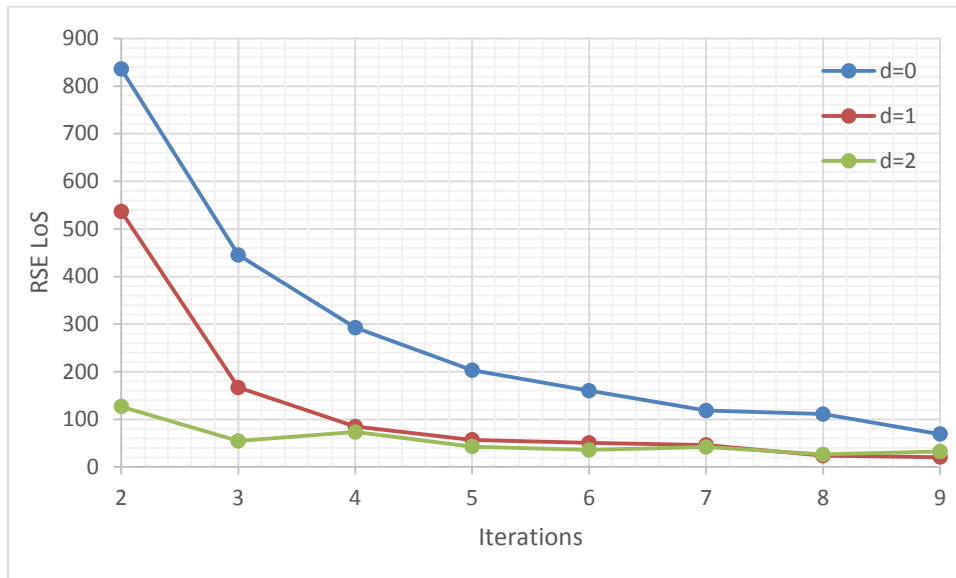
Figure 9. Convergence measured by $\|S(\boldsymbol{t}_k) - \boldsymbol{t}_k\|$ in LoS space. Based on averaging by LoS for different values of $\boldsymbol{d}$ in the weighted MSA.

As for the toy-network, the simple MSA is outperformed by the weighted MSA schemes. It is also notable that convergence was relatively steep in the first couple of iterations, whereas it became flatter as we moved on to the right.

Another way of benchmarking the performance of the different d-values is to look at the difference between $S(\boldsymbol{t}_k)$ and $\boldsymbol{t}_k$ relative to $\boldsymbol{t}_k$. That is, we consider $abs\left(\frac{S(t_k) - t_k}{t_k}\right)$ and divide the degree of deviation into seven categories from below 0.25% deviation to above 12.5% deviation. In doing so, it is possible to picture the distribution profile on the different deviation categories for different iterations. The resulting frequency table is weighted with the number of trips in order to give more weight to larger cells compared to smaller cells. The result is shown in *Figure 10 - Figure 12* below.
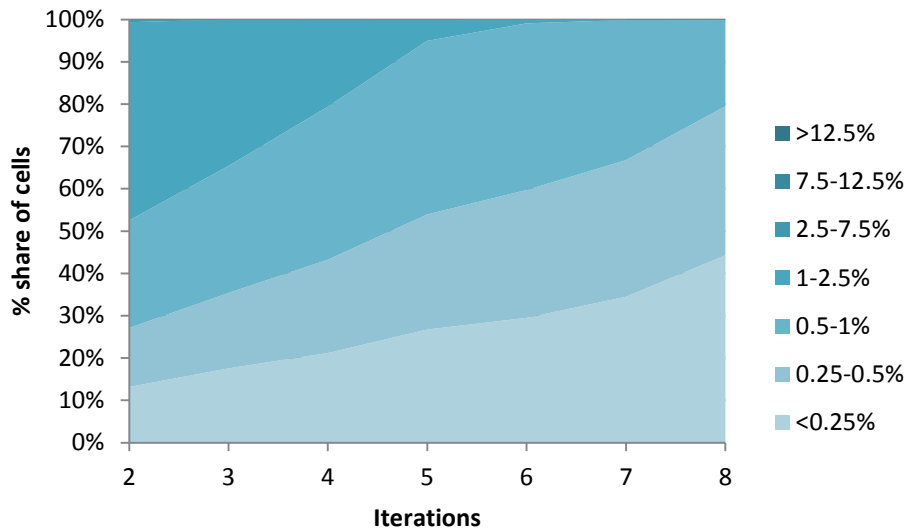
Figure 10. System convergence illustrated as a distribution of $\boldsymbol{abs}\left(\frac{S(t_k)-t_k}{t_k}\right)$ *for LoS and classified in seven deviation bands, with* $\boldsymbol{d} = \boldsymbol{0}$ *and MSA by LoS.*

Figure 11. System convergence illustrated as a distribution of $abs\left(\frac{S(t_k)-t_k}{t_k}\right)$ *for LoS and classified in seven deviation bands, with* $d = 1$ *and MSA by LoS.*
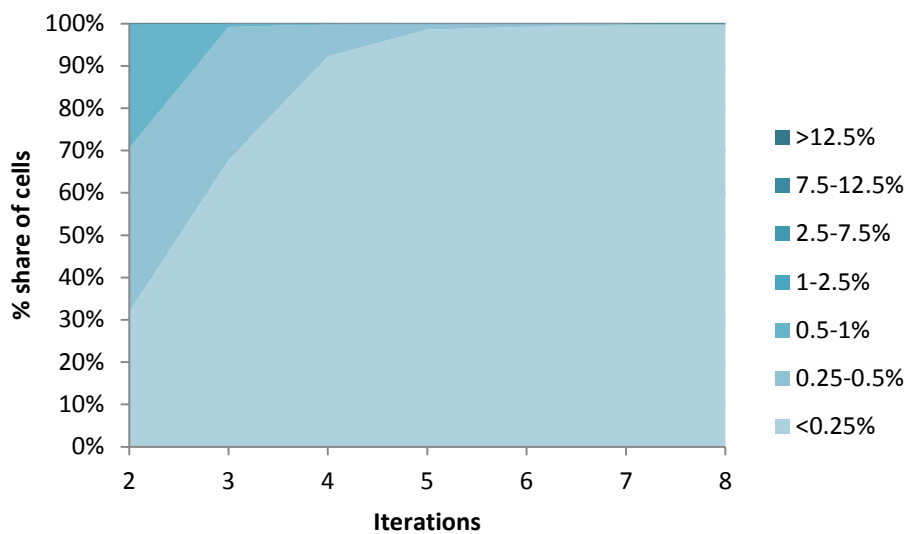
Figure 12. System convergence illustrated as a distribution of $abs\left(\frac{S(t_k)-t_k}{t_k}\right)$ *for LoS and classified in seven deviation bands, with* $d = 2$ *and MSA by LoS.*

Although *Figure 10-Figure 12* does not represent a formal test of convergence, the figures are still informative. In a fully converged scheme for a given iteration number, everything would be "light-shaded". The figures support the conclusion that $d = 2$ is the most efficient choice. In particular, all of the figure underline that $d = 2$ is a very efficient choice for the first round of iterations.

The figures can also be illustrated without weighting by trips. In that case the deviation represents "cell-base" deviation and does not account for the size of cells. These figures, which we do not include in the paper, are quite similar to *Figure 10-Figure 12*. However, they exhibit a more volatile pattern indicating that the main fluctuations are only present in very small cells with limited impact. For $d = 2$ we have tested convergence for 20 iterations and it converged approximately with the same pace as seen from iteration 3 to 10.

5. Summary and conclusions

The paper investigates the external convergence between a demand model and a separate assignment model. The performance of the convergence has direct consequences for the model output uncertainty. Firstly, as large-scale models in practise cannot be expected to be fully converged due to excessive computation time, it is relevant to consider magnitude of uncertainty and how this can be reduced by the use of proper averaging methods. Secondly, even in the best of all worlds where computation time is unlimited, the final model output can be affected by randomness from one or more model parts. In this case, it is crucial that the averaging principle is implemented in such a way that the contribution from the random noise is gradually averaged out. If not, the minimum uncertainty in the final model output is newer better than the level of the random noise.

Whereas much research has been invested in how the inner-loop converges, the convergence of the external loop has not been given much attention in the literature. A reason for this may be that the convergence in the external loop (due to aggregation) is typically well-behaved and can in many cases be iterated without applying averaging principles. However, evidence from both simulation experiments and practise indicates that convergence problems may occur in large-scale applications as well and depends on the volatility the underlying scenario will cause.

The little focus on the external loop also contradicts the fact that the iterations in the external loop are far more costly compared to the inner loop as they involve a complete assignment to be run for every iteration. Hence, it is important to spend these iterations wisely and a difference of 5-10 iterations may well determine whether a model can realistically be applied or not. As an example the Danish National Model currently has a run-time of 50 hours for 5 iterations.

The paper lists a range of averaging principles, which were tested in a small-scale simulation exercise with the following findings;

- The common MRA (Method of Repeated Approximations) algorithm will usually be relatively efficient in low-congested networks, however, convergence is not guaranteed and cyclically unstable behaviour may be present. Generally, as the congestion level increases, the performance of the MRA declines and the failure rate increases.

- The MSA (Method of Successive Average) in its original form with $a_k = 1/k$ will converge but at times very slowly if applied to "volatile" small-scale networks.

- A weighted MSA (d = 2) is particularly efficient, even in situations with few iterations and "hyper" congestion. Generally, the d-parameterisation seems to include plenty of freedom to parameterize most processes.

- A MSA with reset is also very efficient, although the reset point and the iteration stoppage point need to be specified. It appears to be efficient to stop the algorithm one iteration after the reset points irrespectively of the reset length. Although the reset method was better for low and medium congested toy- networks it was not better than the weighted MSA for heavily congested networks. Because of the uncertainty of the reset-point and stoppage point this was not implemented for the national model.

- MSA with Polyak step size and weighted MSA with (d = 1) is significantly better than the MSA but also significantly worse than the weighted MSA d = 2 and MSA with reset.

- Variants of a new intersection approach were presented and tested. However, it falls into the category of tempting ideas that do not work out in practice since it failed to converge in some cases, and since convergence was generally significantly slower than the weighted MSA in the cases where it did converge.

The second wave of tests examined the convergence performance in the Danish National Model. Focus was on the weighted MSA due to the findings from the toy-network simulation experiment. There were two main findings;

1. It was found that if the assignment model involves random noise such as link-based random components, MSA should not be applied to trips but to LoS instead. In case MSA is (wrongly) applied to trips, random components from the route-choice model will be renewed from iteration to iteration and the model will converge, not to a unique equilibrium, but to any solution below the noise-floor of the model. As a result, we recommend averaging over LoS attributes. This finding has a more general consequence, namely that model convergence requires that at least one of the models is free from uncontrolled random noise. This is an important design criterion for complex activity-based models based on micro simulation.

2. It was found that the most efficient weighting in the weighted MSA was d = 2 and hence similar to what was found for the toy-network experiment.

5.1 Future research

Only few publications have been concerned with system convergence in transport model systems and there is a number of interesting future research topics to consider.

- Further investigation into the MSA reset method and investigation into the robustness of large-scale systems.

- The use of series acceleration methods such as the Aitken's delta-squared process (Atkinson, 1989) and the Steffensen's method (Johnson and Scholz, 1968). The latter is a root-finding method that achieves quadratic convergence without using derivatives.

- Investigation into "enforcement methods" as suggested in Vovsha et al. (2008) and the impact on performance and model robustness.

- Assessment whether joint averaging of demand and LoS will further improve performance.

- Research on how to re-use random numbers in Stochastic User Equilibrium Assignment models as well as in the demand model of activity-based micro simulation models. The issue is specifically how to deal with new generated routes and trips between runs.

Further, it is worthwhile considering if the approximate convergence rate, as suggested in Senning (2007), can be used as an alternative measure of convergence speed.

Acknowledgements

References

Agarwal, R.P., Meehan, M., O'Regan, D. (2001). Fixed Point Theory and Applications. Cambridge University Press, ISBN 0-521-80250-4.

Atkinson, K. E. (1989). An Introduction to Numerical Analysis. John Wiley & Sons, Inc., ISBN 0-471-62489-6.

Bekhor, S.; Kheifits, L.; and Sorani, M. (2013). Stability analysis of Activity-Based models: Case study of the Tel Aviv Transportation Model. Unpublished working paper, Technion - Department for Transportation and Geo-Information Engineering, 2013.

Blum, J.R. (1954). Multidimensional Stochastic Approximation Methods. Ann. Math. Stat.,Vol. 25, pp. 737-744.

Cantarella, G.E. (1997). A General Fixed-Point Approach to Multi-Mode Multi-User Equilibrium Assignment with Elastic Demand. Trans. Sci., Vol. 31, pp. 107-128.

Cantarella, G. E., Cascetta, E. (2009). Basic Static Assignment to Transportation networks. In Cascetta, G.E. (eds). Transportation Systems Analysis. Springer.

Cascetta, E., Postorino, M. N. (2001). Fixed Point Approaches to the Estimation of O/D Matrices Using Traffic Counts on Congested Networks. Transport Science, Vol. 35, no. 2, pp. 134-147.

Fox, J., Daly, A., Gunn, H. (2003). RAND Europe Review of RAND Europe's Transport Demand Model Systems. Rand Europe, Report prepared for TRL, ISBN: 0-8330-3362-X.

Grol, R.V., Smit, R., Hofman, F. (2010). Convergence of the Dutch National Model System. Paper presented at the European Transport Conference (ETC), Glasgow Oct.11-13, 2010.

Highway Research Board (1965). Highway Capacity Manual 1965. Special Report 87, National Academy of Sciences, National Research Council, Publication 1328, Washington, D.C. 1965.

Johnson, L. W., Scholz, D. R. (1968). On Steffensen's Method. SIAM Journal on Numerical Analysis, Vol. 5, no. 2, pp. 296–302.

Liu H.X., He, X., He, B. (2009). Method of Successive Weighted Averages (MSWA) and self-Regulated Averaging Schemes for Solving Stochastic User Equilibrium Problem. Network Spatial Economy, Vol. 9, pp. 485-503.

Kelley, C. T. (1995). Iterative Methods for Linear and Nonlinear Equations, Siam, Philadelphia.

Lv, Q., Skrobanski, G., Smith, M., Springham, J., Woods, A., Gordon, A. (2010). A two-direction method for variable demand traffic assignment. In Heydecker, B.G. (eds) Mathematics in Transport. Elsevier, London.

Paulsen, M. (2013). Determination of Transport Network Equilibria - Using the Method of Intersection of Straight Line Approximations of Demand and Cost Curves. B.Sc.-2013-33 Thesis, Technical University of Denmark, DTU Compute.

Polyak, B.T., Juditsky, A.B. (1992). Acceleration of Stochastic approximation by averaging. SICON, Vol. 30, no. 4, pp. 838-855.

Rich J., Nielsen, O.A., Cantarella, G. (2010a). System convergence in transport models. Paper presented at the European Transport Conference (ETC), Glasgow Oct.11-13, 2010.

Rich J., Nielsen, O.A., Brems, C., Hansen, C.O. (2010b). Overall design of the Danish national transport model. AUC Proceedings, Aalborg, August 23-25, 2010.

Rich J., Prato, C., Daly, A., (2010c). Activity-based demand modelling on a large scale: Experience from the new Danish National Model. Paper presented at the European Transport Conference (ETC), Glasgow Oct.11-13, 2010.

Robbins, H., Monro, S. (1951). A Stochastic Approximation Method. Annals of Mathematical Statistics, Vol. 22, no. 3, pp. 400-407.

Senning, J. R. (2007). Computing and estimating the rate of convergence. Department of Mathematics and Computer Science, Gordon College, 255 Grapevine Road, Wenham MA, 01984-1899, http://www.math-cs.gordon.edu/courses/ma342/handouts/rate.pdf.

Sheffi, Y. (1985). Urban Transportation Networks – Equilibrium Analysis with Mathematical Programming Methods. Prentice-Hall, Inc. Englewood Cliffs10.

Vovsha, P., Donnelly, R., Gupta, S. (2008). Network Equilbrium with Activity-Based Microsimulation models – The New York Experience. Transportation Research Record, No. 2054, pp. 102-109.

Vuk, G., Overgaard, C.H. (2010). Some Validation Tests of the OTM, ver. 5.0. AUC Proceedings, Aalborg, August 23-25, 2010.

Williams, H.C.W.L, Lam, W.M. (1991). Transport policy appraisal with equilibrium models I: Generated traffic and highway investment benefits. Transportation Research Part B, Vol. 25, Issue 5, pp. 253–279.

---

[10] Downloadable from http://sheffi.mit.edu/urban-transportation

Appendix A: The intersection approach

Below in *Figure 13-Figure 16* we present selected results where the Intersection method and the MRA are investigated in Monte Carlo experiments for the toy-network. These experiments are parameterised with a relative steep speed-flow curvature and as a result, the model is quite responsive to congestion. Also, demand is fixed at a level which will lead to substantial congestion. As seen the shock causes the algorithm to produce positive failure rates. The tests in *Figure 13 - Figure 16* is parameterised differently that the tests in *Figure 4 -Figure 8*.
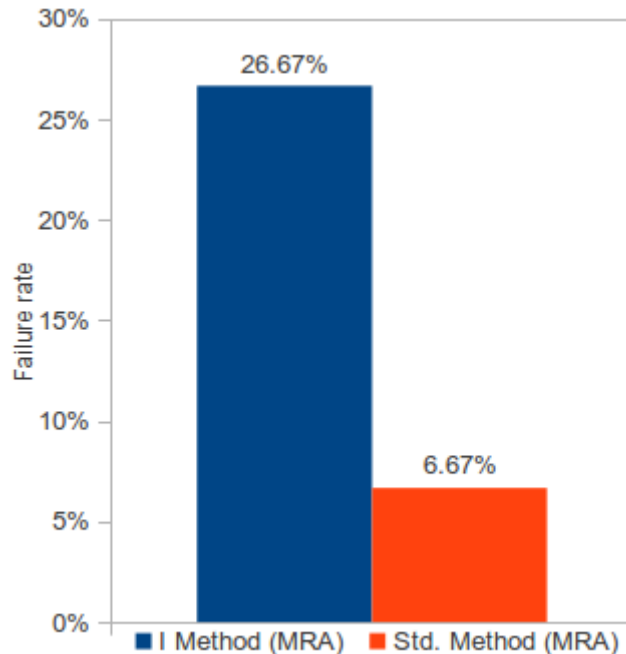

Figure 13. Failure rates of I-method (Intersection method) and MRA based on simulations on toy-network for 1000 iterations. Hard-convergence criteria (maximum absolute difference for OD pairs less than 0.001).
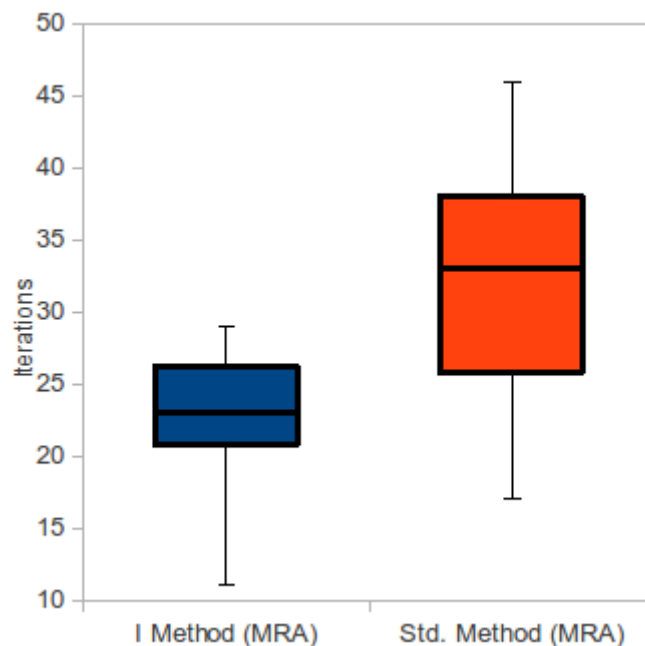

Figure 14. Convergence for the I-method and MRA for the 20 best converged schemes. Hard-convergence criteria (maximum absolute difference for OD pairs less than 0.001).

Figure 14 basically describes the performance (in terms of the number of iterations used to obtain hard-convergence) for the 20 best simulation experiments. This is included in order to describe the performance if everything goes well and failure does not occur. In these cases the I-method is significantly better than the standard MRA.
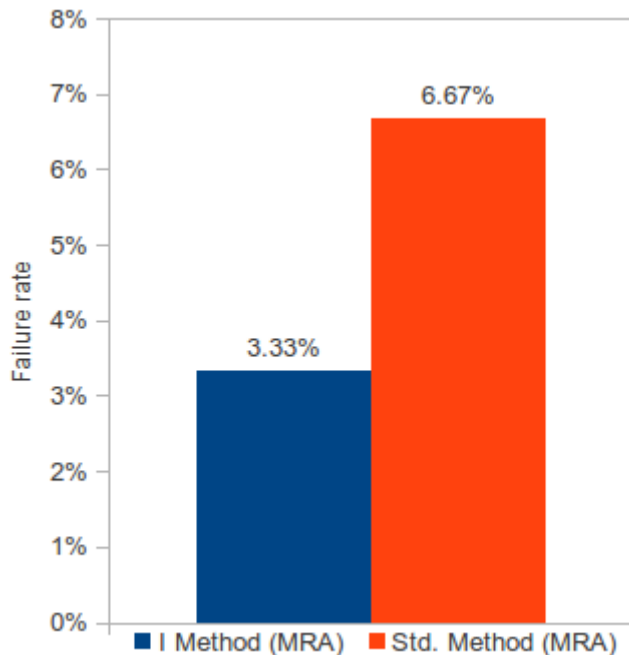

Figure 15. Failure rates of I (Intersection) method and MRA based on simulations on toy-network for 1000 iterations. Soft-convergence criteria (maximum absolute difference for OD pairs less than 0.01).
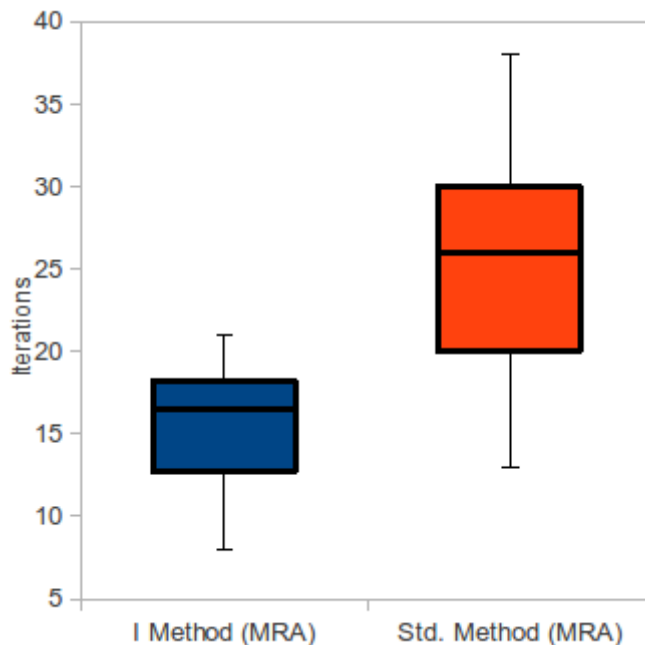

Figure 16. Convergence for the I-method and MRA for the 20 best converged schemes. Soft-convergence criteria (maximum absolute difference for OD pairs less than 0.001).

The interpretation of *Figure 16* is similar to that of *Figure 14* except that we are now looking at soft-convergence.