

SUMO's Lane-Changing Model

Jakob Erdmann
German Aerospace Center, Berlin, Germany

1 Abstract

SUMO is an open source microscopic traffic simulation. A major component of modelling microscopic vehicle behavior is the lane-changing behavior on multi-lane roads. We describe a new model which uses a 4-layered hierarchy of motivations to determine the vehicle behavior during every simulation step and motivate in which ways it improves the current lane-changing model.

Keywords: microscopic simulation, lane changing

2 Introduction

The SUMO application suite [1, 2] provides tools for the *Simulation Of Urban MObility*. It consists of a microscopic simulator for multimodal road traffic and a host of applications for preparing simulation input data (network import and modification, traffic import, routing) and for working with simulation outputs. The microscopic driving dynamics of road vehicles are determined by the interplay of several models briefly listed below:

- **Car-following model:** determines the speed of a vehicle in relation to the vehicle ahead of it.
- **Intersection model:** determines the behavior of vehicles at different types of intersections in regard to right-of-way rules, gap acceptance and avoiding junction blockage.
- **Lane-changing model:** determines lane choice on multi-lane roads and speed adjustments related to lane changing.

When simulating traffic on complex road networks with multi-lane roads, most routes which a vehicle might use require changing lanes. Even where there are no such hard necessities, lane-changing behavior is often a major determinant for traffic efficiency which underscores the importance of the respective model.

The lane-changing model in SUMO has been under continuous development since the start of the project in 2001 and will certainly undergo changes in the future. Due to a large number of improvements in 2013 we see the need to report on the current state of the model. These changes were prompted by problems and visibly implausible behavior in some of our simulation scenarios.

- Motorway traffic which requires many vehicles to change lanes at a point where the motorway splits exhibited heavy jamming contrary to real-world measurements (*A92 scenario*).
- Heavy jamming where motorway traffic in the main direction came to a stop because of vehicles merging at on-ramps (*Braunschweig scenario*).
- Jamming because vehicles did not change to their respective turn lanes in time and thus blocked the flow (*Braunschweig scenario*).

- Jamming because vehicles only used the outer lanes of a two-lane roundabout (*ACOSTA* scenario)

The model changes which were undertaken to alleviate these problems are tightly interwoven with the previous model which makes it impractical to discuss them in isolation. Instead we will describe the new model fully in the following sections and then describe areas of improvement relating to the above scenarios in section 1.9.

The lane-changing model described herein fulfills two main purposes: It computes the change decision of a vehicle for a single simulation step based on the route of the vehicle and the current and historical traffic conditions in the vehicles surroundings. Furthermore, it computes changes in the velocity for the vehicle itself and for obstructing vehicles which promote the successful execution of the desired lane change maneuver.

In comparison to other microscopic lane-changing models, this model explicitly discriminates between four different motivations for lane-changing:

- 1) Strategic change
- 2) Cooperative change
- 3) Tactical change
- 4) Regulatory change

After discussing the general architecture of lane-changing within the simulation in Section 3, the handling of these four motivations will be discussed in detail in Sections 4-7. The complete formulas and decision trees used in the implementation cannot be given due to lack of space. For those wishing to re-implement or modify these models, this paper should serve as a useful guide when reading the source files of the implementation in SUMO [3]. In Section 8 external control of the lane-changing model via the TraCI interface [1] is discussed. Section 9 gives a detailed account of the way conflicting lane-change motivations are resolved. Section 10 documents simulation results of the lane changing model in comparison to older models and section 11 gives an outlook on further developments.

This paper is an extension of [6]. It gives a more detailed account of the decision trees for effecting speed adjustments (section 4.3) and covers changes which were made since that publication. One major change is the handling of regulatory lane changes discussed in section 7. Furthermore, it contains additional evaluation results.

3 Architecture

Road traffic simulation in SUMO represents the road network in terms of **edges** which are unidirectional street segments between intersections and remain constant in their number of lanes, and their maximum speed (among other attributes). An edge consists of one or more parallel **lanes** which correspond to the (mostly marked) lanes found in European road networks. These lanes are **indexed** from right to left starting at 0. The route of a vehicle is stated in terms of the edges it needs to follow but during the simulation it moves along the lanes with mostly free choice of lane usage (except where lane usage restrictions are explicitly defined). Connectivity in the road network is defined on the level of lanes, with each lane having 0 or more successor lanes. If the lane on which a vehicle drives does not have a successor lane which belongs to the next edge along this vehicles route, the vehicles must change its lane in order to continue.

SUMO simulates the movement of vehicles along the aforementioned lanes. In the context of this work the term **vehicle** refers to the model of a real-world vehicle and its driver (sometimes called vehicle-driver unit). The speed of a vehicle is mainly determined by the next vehicle in front of it called the **leader**, which may be on the same lane or on the preferred successor lane after the current lane. This preference is discussed in section 1.4.1. The speed for following the leader is defined by the car-following model which is not discussed in this paper [5]. A vehicle may only change its lane if there is enough physical space on the **target lane** and if it neither comes too close to the leader on the target lane nor to its immediate **follower** on the target lane (*too close* being defined by the car-following model). If either of these conditions is not met, the vehicle is said to have a **blocking leader** or a **blocking follower**. To distinguish the vehicle currently under consideration from its leaders and followers we will refer to it as the **ego** vehicle. A vehicle that advances to a lane on the next edge is said to **advance** the lane, whereas a vehicle that changes to a parallel lane on the same edge is said to **change** lane. By default, lane changes are instant¹. A vehicle is situated completely on the original lane in one simulation step and in the next simulation step it is situated completely on the target lane.

During each simulation step, the following sub-steps are executed in order for every vehicle:

- 1) Computation of preferred successor lanes (called **bestLanes**)
- 2) Computation of safe velocities under the assumption of staying on the current lane and integration with lane-changing related speed requests from the previous simulation step
- 3) Lane-changing model computes change request (left, right, stay)
- 4) Either execute lane-changing maneuver or compute speed request for the next simulation step (involves planning ahead for multiple steps). Whether speed changes are requested depends on the urgency of the lane-changing request.

The sub-steps 3 and 4 are handled by a customizable software component the *laneChangingModel*. This gives a high amount of configurability within the bounds of the architecture. The *laneChangingModel* described in this paper is can be swapped against the previous model by setting user-configurable parameters. In the following, the four motivations for lane-changing are discussed in the order of their priority beginning with the most important. In section 9 we explain how conflicts between these motivations are resolved.

4 Strategic lane changing

Whenever a vehicle must change its lane in order to be able to reach the next edge on its route, we call this type of lane changing *strategic*. This happens whenever the current lane of the vehicle has no connection to the next edge of the route. In this case we say that the vehicle is on a **dead** lane. Note that such a lane does not have to be a dead-end in the common sense. A left-only turn lane is dead from the perspective of a vehicle that wants to go straight. A vehicle may perform a strategically motivated lane change well in advance before reaching the dead lane if no other motivation prevents it. This topic is discussed in the next two sections.

¹ There exists the simulation option `-lanechange.duration` which enables continuous lane change maneuvers. The vehicle occupies both lanes for a part of the given duration depending on its width. This functionality is less mature than the default.

4.1 Evaluating subsequent lanes

Vehicles (or rather their assumed drivers) need to decide a sequence of lanes to follow along their route of edges. In this they have some degree of restriction (because some lanes are dead-ends) and they have some degree of freedom because there are multiple lane sequences available. In SUMO a data structure is computed which allows retrieving the following information necessary for subsequent computations:

- For every lane on the current edge, a sequence of lanes that can be followed without lane changing up to the next dead-end or to a maximum distance (**bestLanes**).
- For every lane on the current edge, the traffic density along the bestLanes (**occupation**)
- For every lane on the current edge, the offset in lane index to the lane which is strategically advisable (**bestLaneOffset**)

Note, that multiple lanes may have a bestLaneOffset of zero. In this case, the bestLaneOffset of other lanes points to the closest best lane. Most parts of this data-structure are only updated whenever the vehicle advances to the next lane. The algorithm for computing this data structure is discussed in [3]. The strategically advisable direction is not part of the customizable architecture because it is rather unambiguous (being based on maximizing the drivable distance without changing lanes and minimizing the number of necessary lane changes). Nevertheless multiple bugs were resolved in this part of the code base during the work on improving the lane-changing model in SUMO. Figure 1 illustrates the bestLaneOffset at a motorway off-ramp.




Figure 1: The ego vehicle (green) needs to move to the bottom lane (with index 0) in order continue on its desired route (green). This lane has a bestLaneOffset of 0. The yellow lane (with index 1) has a bestLaneOffset of -1 indicating a necessary lane change to the right. The red lane (with index 2) has a bestLaneOffset of -2 and is strategically unadvisable.

4.2 Determining urgency

While approaching a dead-end lane, a vehicle has some amount of freedom to pursue the strategically advisable lane (which may involve changing or staying) or to follow conflicting motivations. The urgency for following the strategic necessities (i.e. changing to the left if $\text{bestLaneOffset} < 0$ and changing right if $\text{bestLaneOffset} > 0$) correlates with the following factors:

- remaining distance to the dead-end (negative correlation)
- the presumed speed while approaching the end of the dead lane (**lookAheadSpeed**)
- magnitude of the bestLaneOffset
- occupation on the ultimate target lane (lane with bestLaneOffset = 0)
- occupation of the intermediate target lane (next target in direction of the bestLaneOffset)

A strategic change is deemed urgent if the following relation holds true:

$$d - o < lookAheadSpeed * abs(bestLaneOffset) * f$$

Where the d is the distance to the end of the dead lane, o is a discount due to occupation and f is a factor that encodes the time typically needed to perform a successful change maneuver set to 10 for changing to the left and 20 for changing to the right.

Notably, if there are multiple lanes in between the current lane and the ultimate target lane, all their occupations should also matter, but are not currently evaluated. The `lookAheadSpeed` depends on the current and historical speed of the vehicle. This is necessary to avoid vehicles which temporarily have to slow down from losing all sense of urgency. The expected number of seconds until reaching the end of the dead lane is divided by the number of necessary lane changes (`bestLaneOffset`) to obtain the available time for the current lane change (**remainingSeconds**). This value is used in subsequent computations. Currently, urgency is only considered for strategic lane changes but we discuss how it could apply to other motivations in section 11.

4.3 Speed adjustment to support lane-changing

Whenever a desired lane change cannot be executed due to blocking vehicles, a vehicle may adjust its speed to allow the lane change to succeed in later steps. Furthermore a vehicle may exert an influence on the speed of blocking vehicles (in reality this typically happens as a reaction to observing the turn signals of the ego vehicle). Due to the importance of completing strategic lane changes, it is assumed that the ego vehicle will take careful adjustments to enable the change. Basically, vehicles are assumed to drive at the maximum safe speed, so speeds can only ever be adjusted downwards. However, as a part of the car-following model, vehicles may have a stochastic component (called **dawdling**) which prevents them from using their maximum possible acceleration. Preventing this stochasticity is a way of increasing vehicle speeds somewhat.

To compute the desirable speed adjustments, the following hierarchy of situations is distinguished by comparing the **plannedSpeed** of the ego vehicle, blocker speed, gaps and `remainingSeconds`:

(1) Leader is blocking

- a. **able to overtake leader:** request leader to refrain from speeding up, prevent ego dawdling, (prevent overtaking on the right where forbidden by law)
- b. **unable to overtake leader**
 - i. slow down to stay behind the leader
 - ii. keep speed since the leader is faster anyway

(2) Leader is not blocking: set a maximum speed to ensure that the distance to the leader remains sufficiently high

(3) There is no leader: drive with the maximum safe speed

The decision whether a blocking leader may be overtaken (1a vs 1b) is illustrated in Figure 2. The choice is made by checking a list of necessary conditions for overtaking:

- a. The ego vehicle is faster than the blocker ($dv = plannedSpeed - blocker\ speed > 0$). The `plannedSpeed` incorporates speed requests by surrounding vehicles.

- b. The blocking vehicle is to the left of ego or overtaking on the right is allowed (in urban situations, on congested motorways or if the simulation option *-lanechange.overtake-right* is set)
- c. The remaining space to the end of the dead lane is sufficient for overtaking
- d. The time remainingSeconds is sufficient to overtake the leader at the current speed difference dv




Figure 2: The ego vehicle (green) needs to change to the right to continue with its route. It could either change behind the truck or it could overtake the truck and change in front of it.

The above decision tree results in an updated value of **plannedSpeed** with regard to a blocking leader. Another decision tree is used to compute the behavior in regard to a blocking follower. This decision tree considers plannedSpeed, blocker speed, gaps and remaining seconds.

(4) Follower is blocking

a. will be able to cut in before follower

- i. **fast enough to do so with current speeds:** request follower to refrain from speeding up, prevent ego dawdling
- ii. **follower decelerating once is sufficient to open a gap:** request follower to decelerate as much as needed, prevent ego dawdling

b. needs to be overtaken by follower

- i. **follower should slow down a bit to increase the chance that subsequent followers will be slow enough:** request follower to decelerate a bit, slow down to be overtaken fast enough
- ii. **follower should overtake quickly:** prevent follower from dawdling, slow down to be overtaken fast enough

c. follower cannot overtake on the right: request follower to slow down if above a minimum speed threshold

(5) Follower is not blocking: request follower to maintain speed so as to remain non-blocking

(6) There is no follower: drive with the maximum safe speed

Speeds, computed in this way are integrated with the maximum safe speed (**vSafe**) as computed by the car-following model by using the minimum of vSafe and all requested speeds.

The distinction between cases (4)b.i and (4)b.ii warrants further explanation. Whenever a vehicle tries to change from an on-ramp onto the motorway it has to yield to vehicles already on the motorway. These vehicles may slow down slightly to help merging vehicles, but they must not cause the flow on the motorway to break down. For this reason, vehicles that try to change to the left only cause blocking followers to slow down if their own speed exceeds a threshold value (currently 27m/s).

4.4 Preventing deadlock

If a vehicle needs to stop on a dead lane because changing to a continuing lane did not succeed it creates an undesirable impediment to traffic flow. The measures in the previous section help to prevent this situation from occurring too often (and it does occur in reality as well). However, if two vehicles on adjacent lanes both need to change to the lane occupied by the other vehicle (referred to as **counterLaneChange**) and both vehicles reach the end of a dead lane, a deadlock occurs. Neither vehicle has the option of driving any further nor can either vehicle get the space it needs to execute the strategic lane change (vehicles in SUMO cannot go backwards). This situation blocks the flow of traffic on both lanes and is highly undesirable. Currently, it can only be resolved by moving vehicles in a non-standard way (teleporting) after a time threshold is elapsed. Figure 3 illustrates situations which may lead to a deadlock.

To prevent deadlock, special care is taken whenever two vehicles are in a counterLaneChange relation. We refer to the vehicle which is closer to the end of the dead lane as the **blocking leader** and the other vehicle as **the blocking follower**. Note that this relation may change from one simulation step to the next. Generally, the blocking follower slows down when approaching the dead-end to ensure that the blocking leader has enough space to complete its lane change. In some cases the blocking follower is too fast or the blocking leader is too long. In this case the blocking leader must slow down to leave enough space for the follower before the dead-end.

Unfortunately, dead-lock situations can still arise if vehicles need to perform strategic lane changes across multiple lanes. In this case, a counterLaneChange situation can arise at a time where both vehicles have already reached the dead-end and are unable to move. To prevent this, vehicles reserve additional space in front of the dead-end whenever they have to change across more than one lane. Currently, additional space of 20m is reserved for vehicles which need to change to the right by more than one lane and 40m for vehicles which need to change to the left by more than one lane. The asymmetry is necessary to prevent yet another type of deadlock. The values were selected because they were found to perform well in preventing deadlock. Eventually they should be made configurable and be subject to rigorous calibration.

An important aspect of preventing stopping at a dead lane (and thus deadlocks) is avoiding detrimental lane changes. Generally speaking, the fewer lane change maneuvers vehicles have to perform, the less chance they have to become stuck. One change that was found to be quite beneficial was the avoidance of changing to a dead lane which continues elsewhere. This is most often the case for turn-lanes at an intersection. A vehicle which intends to go straight should not use the left-only turn lane to get ahead because it will find it difficult to go back onto the required lane. In reality these turn-lanes often have directional markings at their start and there are rules which prohibit their use by vehicles which follow another direction.





Figure 3: A deadlock may arise between the purple vehicle which needs to change left and the green vehicle which needs to change right to continue with its route. Another deadlock may arise between the purple vehicle and the yellow vehicle which needs to change to the right twice.

5 Cooperative lane-changing

In some real-world situations vehicles (or rather their drivers) perform lane-changing maneuvers with the sole purpose of helping another vehicle with lane-changing towards their lane. In the current model, vehicles are informed by other vehicles about being a blocking follower (the reason being that the turn-signals of the vehicle being blocked are always visible to the follower, whereas being a blocking leader is less obvious). If there are no strategic reasons against changing the lane, the ego vehicle may change in either possible direction to clear a gap for the blocked vehicle. Contrary to expectation, this may have a beneficial impact on traffic flow even if the ego vehicle attempts to change towards the blocked vehicle. This effect is not yet understood and warrants further investigation.

Vehicles which cannot perform a cooperative lane change adjust their own speeds slightly to increase the success probability for subsequent simulation steps. However, they do not request speed changes if they are blocked.

A special case for cooperative behavior arises at multi-lane roundabouts. Typically, all vehicles enter the roundabout at the outermost lane and also need to leave again at the outermost lane. Due to the short distances involved, this means they should always remain on the outermost lane for strategic reasons. However, this effectively turns all multi-lane roundabouts into one-lane roundabouts and thus degrades throughput. For this reason, the lane-changing model compels vehicles which are not yet on their final roundabout edge to change towards the inner lane. While this ignorance of strategic motivations sometimes results in stranded vehicles it has a beneficial impact on roundabout performance (see results for the ACOSTA scenario).



6 Tactical lane-changing

Tactical lane-changing refers to maneuvers where a vehicle attempts to avoid following a slow leader. It requires balancing the expected speed gains from lane changing against the effort of lane-changing (which is arguably a very driver-subjective value). The expected speed gains must also be balanced against the obligation for keeping the overtaking lane free. Failure to do so results in situations where slow vehicles with minor speed differences become major impediments to traffic flow. Figure 4 shows a situation in which tactical lane changing may take place.

This part of the model is left unchanged from the old model [4]. Each vehicle maintains a signed variable **speedGainProbability** which by its sign indicates the beneficial change direction (-1 for right, 1 for left) and by its magnitude the expected benefit. If the magnitude exceeds a threshold value, a tactical lane change is attempted. If the lane change succeeds the value is reset to 0. The accumulation over multiple time steps prevents oscillations. During each simulation step and for each considered change direction d , the potential gain $g = (v - u)/v$ is computed. If g is positive, the variable speedGainProbability is incremented by $d \cdot g$. If g is not positive and speedGainProbability has the same sign as d , it is halved instead.

Under some circumstances, overtaking another vehicle on its right side is forbidden. To overtake in this case, the ego vehicle needs to change lanes to the left and wait for the slower leader to move to the right itself (effectively swapping lanes). If there are more than two lanes available, the ego vehicle may also change to the left twice in order to overtake the leader on its left side. The latter situation is illustrated in Figure 4. As one of the additions of the new lane-changing model, this behavior is now triggered if all of the following conditions are met:

- The option **--lanechange.overtake-right** is not set
- The ego vehicle is driving at a speed of 60km/h or above
- There is a leader vehicle on the adjacent lane to the left
- The leader is slower than the ego vehicle
- The ego vehicle would need to slow down if it were to follow the leader vehicle

In this case the ego vehicle will slow to the safe following speed and receive an impulse to change to the left by incrementing speedGainProbability by a fixed amount (enough to exceed the threshold if the conditions hold for three simulation steps). It should be noted that the 60km/h check in the above conditions stands as a rough approximation for a far more complex set of rules which are dependent on legal rulings in the applicable country to be modelled [8]. For Germany, the value of 60km/h corresponds not to a text of law but to a legal precedent which concretizes the wording of “slow speed”.




Figure 4 The ego vehicle (green) is faster than its leader vehicle (purple). To prevent slow down, the ego vehicle may change to the left in order to overtake its leader. Since overtaking on the right is forbidden, two lane changes are necessary for overtaking.

7 Regulatory lane-changing

In jurisdictions with right-handed driving, the left lane(s) are designated as overtaking lanes. Drivers are under the obligation to clear that lane whenever they do not use it for an overtaking maneuver. The obligation to clear the overtaking lane could be framed as cooperative behavior because it helps other faster moving vehicles. However, contrary to the cooperative lane-changing behavior described in Section 5 which is optional, the behavior described in this section is mandated by traffic laws [9]. Figure 5 illustrates a situation which calls for regulatory lane changing

In the current lane-changing model, each vehicle maintains a variable **keepRightProbability** which is decremented over time and triggers a lane change to the right once a lower threshold value is exceeded (negative values are used in allusion to the variable **speedGainProbability**). The formula for computing the new value of keepRightProbability q' from the old value q is designed for clearing the overtaking lane as soon as possible while at the same time avoiding oscillations from repeated cycles of overtaking on the left and changing back to the right lane.

$$q' = q - T \frac{t \cdot m}{d \cdot v}$$

Here, t denotes the expected time which for which the ego vehicle will be able to drive on the right lane with full speed (based on the distance to the leader vehicle and its speed). The value of m denotes the speed limit on the target lane while d denotes the desired speed of ego and v its current speed. T is a constant for tuning the urgency of regulatory changing. This formulate encapsulates the following concepts:

- t : vehicles that can expect to continue with high speed on the right lane are more likely to change to the right
- m / v : slower vehicles (w.r.t the speed limit) are more likely to change to the right
- $1 / d$: vehicles which desire high speed are more likely to stay on the overtaking lane




Figure 5 The ego vehicle (green) is about to overtake. After finishing the maneuver, it is required to move back onto the rightmost unobstructed lane

8 Remote controlled lane-changing (TraCI)

Running SUMO simulations can be controlled by external programs using an interface called TraCI (**Traffic Control Interface**). Among the things that can be controlled is the choice of lane among the available lanes for each vehicle. External requests to change to a target lane or keep the current lane must be integrated with the “internal” requests computed by the lane-changing model. This is accomplished by letting the user determine the urgency and the priority of remote requests by setting appropriate flags.

For each of four change motivations discussed above the following options can be independently configured:

- a) Ignore internal request
- b) Ignore internal request when in conflict by an external request
- c) Always follow internal request regardless of external request

Furthermore, the following options for configuring the urgency of external requests:

- a) Following request regardless of surrounding vehicles, perform urgent speed adaptations
- b) Follow request unless it would cause an immediate collision but ignore safety gaps to surrounding vehicles, perform urgent speed adaptations
- c) Only change if all safety constraints are met, perform urgent speed adaptations
- d) Only change if all safety constraints are met, perform no speed adaptations

As an example, the interface allows the remote program to specify that a given vehicle should try to change to the left lane with urgency (i.e. with speed adjustments to itself and to blockers), unless there are urgent strategic reason against changing to the left and that the vehicle should ignore all other requests by the lane-changing model.

9 A hierarchy of lane-changing motivations

The four motivations discussed above are considered in a hierarchical fashion as described by the following decision schema. The first statement which applies determines the vehicles change request. In every simulation step, each vehicle first considers changing to the right, and if no change to the right is performed, a change to the left is considered as well. Accordingly, the currently considered direction d is either right (-1) or left (1) according to the resulting change in lane index.

1. **Urgent strategic change to d needed:** change (strategic)
2. **Change to d would create an urgent situation:** stay (strategic)
3. **Vehicle is a blocking follower for another vehicle with urgent strategic change request:** change (cooperative)
4. **speedGainProbability above threshold and its sign matches d :** change (tactical)
5. **keepRightProbability above threshold and $d = -1$:** change (regulatory)
6. **non-urgent strategic change to d needed:** change (strategic)

10 Improvements over the earlier model

In the following we present measurements which document the effect of model changes on traffic flow and lane changing efficiency. Section 10.2 presents additional measurements which were undertaken since [6] when working on the model for regulatory lane changing.

10.1 Efficiency of lane-changing

For a quantitative evaluation of the improvements, the following metrics were computed for a selection of benchmarking scenarios.

- **avgWaitingTime:** the average time each vehicle spent with speed below 0.1m/s
- **wrongLaneTeleports:** the count of vehicles which had to be moved artificially (teleported) because they could not complete a strategic lane change (after a threshold time t)
- **jamTeleports:** the count of vehicles which had to be moved artificially (teleported) because the successor lane was occupied (after a threshold time t)

The scenario *Braunschweig* contains the urban area of the German city of Braunschweig (Brunswick) and the surrounding area with sections of motorway. The scenario spans one day and contains 650000 vehicle movements. The threshold time for teleporting was set to 120 seconds.

The scenario *A92* consists of a motorway section in southern Germany with a length of 20km. It contains 63000 vehicle movements over the course of one day. The threshold time for teleporting was set to 300 seconds.

The scenario *ACOSTA* is comprised of a section of the Italian city of Bologna and contains 9000 vehicles over the course of 1 hour. It is notable for containing a 2-lane roundabout. The threshold time for teleporting was set to 300 seconds.

The algorithms *old* and *new* correspond to the SUMO vType parameters `laneChangeModel="DK2008"` and `laneChangeModel="JE2013"`. As can be seen in Table 1, the new algorithm brings a significant improvement in all considered scenarios. Additional topics for future improvement are discussed in the next section.

Scenario	Algorithm	Avg. waiting time	Wrong-lane teleports	Jam teleports
Braunschweig	old	89.73	845	464
	new	46.66	7	9
A92	old	17.16	21	1
	new	0.02	0	0
ACOSTA	old	144.59	0	7
	new	76.69	0	0

Table 1 Performance metrics for old and new lane-changing model and different scenarios.

Compared to the old lane-changing model described in [4], the new model shows improvements in the following areas:

- Fine grained control over speed adjustments to ego vehicle and blockers lead to higher fulfillment rate of change request. In the old model, vehicles always reacted to blocking leaders by slowing down and they always slowed down when being a blocking follower (improved all metrics and all scenarios).
- Extrapolation of dynamics over multiple steps allows better choices between overtaking blockers and allowing to be overtaken (also improves the success rate and thus improves all metrics).
- Improved checking for deadlock-prone situations avoids deadlocks in more cases. In the old model, some cases of deadlock were avoided by allowing neighboring vehicles with opposite change requests to swap their positions instantly. This oversimplification is no longer necessary (primary impact on wrongLaneTeleports but secondary effect for the other metrics).
- Asymmetrical behavior when helping other vehicles with lane-changing (depending on the direction of change) prevents the main flow from breaking down at busy highway on-ramps (improved avgWaitingTime especially in *Braunschweig*).
- Special behavior within multi-lane roundabouts ensures that all lanes are used whereas in the earlier model only the outer lane was ever used (improved avgWaitingTime and jamTeleports, only *ACOSTA*).

- The explicit discrimination between the 4 different motivations for lane changing allows fine grained control for integrating model dynamics with external change requests (TraCI). This was necessary to successfully complete a project which simulated automated platooning (not discussed here).

Note, that a large number of model changes were tested in isolation using the above metrics. To simplify the presentation of our results we only show the effect of all combined model changes. It can be seen that all metrics improved for all scenarios except for the few cases where they had the best possible value to begin with, thus validating the usefulness of the new lane changing model.

The *A92* scenario is based on fine grained detector measurements which showed no jamming in the real world data (in contrast to simulation with the old model). Also, the *Braunschweig* scenario exhibited deadlocks and jamming with a frequency that was utterly implausible for the demand model of a normal working day when simulation with the old model. Although improved traffic flow is not generally a sign of a more realistic model (after all, jams are a fact of life), for the above scenarios an increase in realism can be posited.

10.2 Lane usage

In the old model, regulatory changing was conflated with tactical lane changing by giving an additional decrement to `speedGainProbability` whenever the right lane allowed a high enough speed (for the next simulation step). This way, the decision for performing a regulatory lane change was strongly based on accumulated knowledge about the past rather than anticipated behavior in the future. The noticeable disadvantage of this design was the high delay between passing a slower vehicle and pulling back into the rightmost lane which had a detrimental effect on average road speed. The new model is designed to change to the right shortly after overtaking whenever the desired speed can expectedly be maintained on the right lane for a sufficient time.

To investigate lane usage properties a simulation scenario based on measurements from the German motorway A3 was set up. The input data consists of single vehicle detections from a detector cross-section with 3 lanes measured over a whole day². The motorway experiences an average flow of 1800 vehicles per hour with a peak flow of 3500 vehicles per hour. The data is described in detail in [10]. Each data point contains a high-resolution time stamp, as well as the speed and the length of the vehicle. This data was fed into a simulated 3 lane motorway with a length of 30km. The simulated vehicles were inserted with normally distributed values for *tau* and *sigma* by generating an individual *vType* element for each vehicle. The maximum speed, departure time and vehicle length were taken from the real-world measurements.

Simulated detectors were placed with a spacing of 1 km to accumulate the absolute number of vehicles passing on each lane as well as the average speed. The investigation assumes that vehicles at the measurement location are in an equilibrium state in regard to their speeds and lane usage. Under that assumption, the absolute count of vehicles on each lane should remain nearly constant over the length of the motorway. Likewise, the average speed on each lane should remain constant over the length of the motorway.

² 1995, 11th of March, a Saturday

Figure 6 shows the evolution of lane usage for different versions of the lane-changing model. It can be seen that older versions of the lane changing model deviate much stronger from the equilibrium assumption whereas newer versions deviate less. Specifically, the old model DK2008 placed excessive weight on regulatory lane changing whereas the intermediate version of the JE2013 model in revision 0.19.0 had no regulatory changing at all due to an implementation bug.




Figure 6 Lane usage measurements for different versions of the lane changing model. Top left: DK2008 sumo 0.18.0, top right: JE2013 sumo 0.19.0, bottom left: JE2013 sumo 0.20.0, bottom right: JE2013 sumo 0.22.0³. Version 0.21.0 is not shown due to being the same as 0.20.0.

In addition to the lane usage and speeds, the average number of lane changes per vehicle and driven kilometer was measured for different lane change models. A survey by Lee et al. measured an average number of **0.26** lane changes per kilometer and vehicle while driving on American highways and interstates [7]⁴. Table 2 gives the corresponding value for different model versions. A large jump can be seen from sumo version 0.19.0 to 0.20.0 when the new formula for effecting regulatory lane

³ The lane changing model slated for the upcoming release version 0.22.0 is already available as source code revision 17102.

⁴ Lee et al. report 0.36 lane change maneuvers per mile or 0.22 per km with 16% of the maneuvers consisting of more than one lane change.

changes was introduced. However, due to a bug that was causing oscillations between tactical and regulatory lane changing⁵ the number of lane changes was higher than intended. This bug is fixed in the upcoming version 0.22.0 resulting in a reduced number of lane changes. Overall an improvement in realism can be posited based on the measured values.

Table 2 Lane changing statistics for different model versions

Model	SUMO version	Lane changes per vehicle and km.
DK2008	0.18.0	0.03
JE2013	0.19.0	0.04
JE2013	0.20.0	0.45
JE2013	0.21.0	0.45
JE2013	0.22.0 ⁶	0.26

11 Outlook

The focus of the recent improvements of the lane-changing model was on deadlock-prevention and success rate of lane-changing as well as lane usage. The next goal is to perform calibration in order to reproduce lane-changing frequency.

To perform model calibration, several hard-coded model parameters shall be exposed to the end user. There should at the very least be one parameter for each of the four motivations:

- Urgency of strategic changes
- Tradeoff between altruistic and egoistical behavior
- Eagerness to realize speed-gains
- Eagerness to clear the overtaking lane

Using these parameters the model should be calibrated and validated using real world measurements.

The type of deadlock discussed in section 4.4 could be considered as a network modelling error. In reality, vehicles have the option of moving “diagonally” from their current lane to the target lane on a subsequent edge without requiring free space for their full length on the target lane. Typical traffic regulations would prohibit this kind of maneuver except for preventing deadlock. These connections are currently not modelled because SUMO currently cannot handle multiple connections from the same edge to the same target lane. It may be necessary to extend the network model to achieve more realistic traffic flow in deadlock-prone scenarios.

Cooperative lane-changing has not been extensively looked at and is a probable candidate for model improvements. Currently, only blocking followers change cooperatively whereas real-world situations are conceivable in which blocking leaders change as well. A typical situation which is not yet considered by the lane-changing model is the coercion to change to the right because a faster vehicle is approaching from the rear on the same lane. Another point is the usage of multi-lane roundabouts where currently, some vehicles become stuck on a dead-end inside lane. Additional checks should be done to prevent these situations from

⁵ Vehicles were changing to the left in anticipation of overtaking a slower vehicle which was still very far away and shortly afterwards changing back to the right by obligation.

⁶ Upcoming release. Source code revision 17102

arising. It would also be helpful to know the degree in which inner lanes are used in reality. The fact that cooperative lane changes towards the blocked vehicle benefit simulation performance should also be investigated.

Some of the above issues might be resolved by extending the concept of *urgency* to all four motivations. A cooperative lane change is more urgent if the supported vehicle is about to suffer a bigger speed loss unless it receives help. Likewise a tactical lane change is more urgent if the ego vehicle is about to suffer a bigger speed loss (due to a slow leader on its lane). Changes with the intent of clearing the overtaking lane are more urgent if the follower on this lane is about to suffer a bigger speed loss as well.

Another point that should be addressed in the future is the interaction between lane-changing and car-following. In SUMO vehicles always maintain sufficient gaps to allow safe stopping if their leader vehicle were to brake with maximum deceleration until stopped. Likewise, the ego vehicle only changes to a new lane when the follower vehicle has enough space to stop safely if the ego vehicle were to brake with maximum deceleration until stopped.

In reality, drivers may accept much lower front-headways during lane-changing which may be justified by any of the following arguments:

- Their leader vehicle will usually not start to brake hard (especially when there is no apparent blockade)
- They may change the lane again to avoid collision if necessary
- They have a lower reaction time by concentrating on a critical maneuver and thus are able to use smaller gaps safely

Likewise, drivers in reality may accept much lower rear-headways during lane-changing by using the following justifications:

- They will not suddenly start braking hard when there is no obstacle (which they could see in advance if it was there)
- They may plan to continue their lane change maneuver one lane further which only requires them to remain safely clear of the follower vehicle for a brief time window

Some of these justifications may be applicable to SUMO vehicles and could be used for altering the parameters of the car-following model when evaluating the safety of a lane change maneuver. This would go a long way towards increasing the realism in scenarios such as highway on- and off-ramps where urgent strategic changes are needed.

12 References

- [1] Behrisch, M., Bieker, L., Erdmann, J., Krajzewicz, D.: SUMO – Simulation of Urban MObility: An Overview. In: SIMUL 2011, The Third International Conference on Advances in System Simulation (2011)
- [2] DLR and contributors: SUMO homepage. <http://dlr.de/ts/sumo/> (2014)
- [3] SUMO source code corresponding to this document. <http://sumo-sim.org/trac.wsgi/browser/trunk/sumo/src/?rev=17102>

- [4] Krajzewicz, D.: Traffic Simulation with SUMO – Simulation of Urban Mobility. In: Barceló, Jaume (Ed.): Fundamentals of Traffic Simulation, Series: International Series in Operations Research & Management Science, Vol. 145, Springer, ISBN 978-1-4419-6141-9 (2010)
- [5] Krauß, S., Wagner, P., Gawron, C.: Metastable states in a microscopic model of traffic flow. Phys. Rev. E, American Physical Society, 1997, 55, pp. 5597-5602, (1997)
- [6] Erdmann, Jakob (2014) Lane-Changing Model in SUMO. In: Proceedings of the SUMO2014 Modeling Mobility with Open Data, 24, Seiten 77-88. Deutsches Zentrum für Luft- und Raumfahrt e.V.. SUMO2014, 16.-16. Mai 2014, Berlin, Deutschland. ISSN 1866-721X
- [7] Lee S. E., Olsen E. C. B., Wierwille W. W. (2004). A Comprehensive Examination of Naturalistic Lane-Changes. Report No. DOT HS 809 702. Washington, D.C.: National Highway Traffic Safety Administration.
- [8] § 7 StVO (German Straßenverkehrsordnung)
- [9] § 2 StVO (German Straßenverkehrsordnung)
- [10] Knospe, W.; Santen, L.; Schadschneider, A. & Schreckenberg, M. Single-vehicle data of highway traffic: Microscopic description of traffic phases Physical Review E, American Physical Society, 2002, 65, 056133