



HAL
open science

Experiments in autonomous reversing of a tractor-trailer system

Cédric Pradalier, Kane Usher

► **To cite this version:**

Cédric Pradalier, Kane Usher. Experiments in autonomous reversing of a tractor-trailer system. 6th International Conference on Field and Service Robotics - FSR 2007, Jul 2007, Chamonix, France. inria-00195700

HAL Id: inria-00195700

<https://inria.hal.science/inria-00195700v1>

Submitted on 11 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Experiments in autonomous reversing of a tractor-trailer system

Cédric Pradalier¹ and Kane Usher¹

CSIRO ICT Centre, Autonomous Systems Laboratory, Brisbane, Australia
`firstname.lastname@csiro.au`

Summary. Trailer reversing is a problem frequently considered in the literature, usually with fairly complex non-linear control theory based approaches. In this paper, we overview our simple method for stabilizing a tractor-trailer system to a trajectory based on the notion of controlling the hitch-angle of the trailer rather than the steering angle of the tractor. The performance of this control method, as implemented on the CSIRO Autonomous Tractor, is then experimentally compared against a number of human drivers, showing that this method performs as well as a moderately skilled human driver, even though the system is significantly handicapped in terms of steering actuation speed and by errors in localization.

1 Introduction

Reversing of tractor-trailer systems is a problem frequently addressed in the literature due to the common, real-world occurrences of this problem, along with the interesting mathematical properties of this type of system. Industrial occurrences for reversing tractor-trailer systems include transportation, cargo handling and agriculture. In the main, professional drivers have very high skill levels and experience little problem with this task. However, there are a number of examples of recreational activities which require a driver to position a trailer, boat and caravan owners to name a few. In Australia, the ‘grey-nomad’ phenomenon sees a growing number of retirees, a large portion of whom are inexperienced in towing a caravan, embarking on a journey that takes them around the country. It is hard to imagine that throughout this journey that they would not be required to reverse this tractor-trailer system, a task which can be quite arduous for drivers of limited experience and ability. This is but one example in which a driver-aid or automation system could be of significant use.

Most approaches address the problem of reversing a tractor-trailer system through the development of relatively complex, non-linear controllers. Our approach is based on closing an inner-loop on the angle between the tractor and

trailer using a simple Proportional Integral (PI) controller, and then closing an outer loop on the trajectory using a path tracker developed for articulated vehicles. This paper presents results from experiments which evaluate the performance of this approach in long duration reversing tasks, and compares its performance to that of human drivers with a range of skill levels.

1.1 Background

In the scientific literature, this problem has been approached in three ways: through a *change of variables* which brings the system into a standard form, *optimal control* methods, and *learning-based* approaches.

The *change of variables* approach relies on a transformation which brings the system into a chained (or similar) form, from which standard non-linear control theory can be applied, see for example [2]. These techniques have the advantage of being theoretically generalizable to any number of trailers. Their main disadvantage, with respect to the work presented in this paper, is their complexity and high sensitivity to numerical approximation. This complexity makes it difficult to implement these methods reliably, and also poses difficulties in the tuning of these systems — after the change of variables, many of the control parameters no longer have a physical meaning making the design of a tuning strategy challenging and somewhat non-intuitive. As for numerical stability, the chained-form-based conversion for instance requires several trigonometric operators and inversions, and the resulting control often requires inversion of matrices which can become singular at some point of the navigation.

The *optimal control* based methods use optimization schemes to derive a sequence of demands which will control the tractor-trailer system onto the desired trajectory or path, see for example [1]. In simple terms, these methods use a vehicle model and a simulation process to compute the control commands that will lead to the best tracking of the trajectory. This requires accurate models of the vehicle’s behaviour and a heavy reliance on computational resources. Deviations from the model, which in real-world implementations are inevitable, lead to errors, and these methods are also open-loop meaning that re-planning is necessary to deal with errors in localization.

Finally, the *learning-based* controllers seek to ease the computational burdens of the previous methods by providing a mapping between the current vehicle state, the desired state and the required inputs to reach the desired state, see for example [7]. Essentially, a simulated model of the tractor-trailer system is used to try many different possible methods and parameters. ‘Learning’ occurs by searching the parameter space for the best set of methods/parameters which are then encoded into, for example, a neural network or look-up table, which maps the current to desired configuration. The main drawback of these techniques is the learning itself: if learning occurs from a model of the system, then errors in the model are clearly problematic; if learning occurs on the real vehicle, then there are clearly safety issues since it is not possible to predict

the behaviour of the vehicle in the learning phase. Also, even if the best set of parameters performs well in practice, it is hard to guarantee its performance, which can be problematic in an application where reliability is an issue.

Several patents have been proposed for automatically reversing a tractor-trailer system [3, 5, 6, 10]. However, these patents are focused on reversing on straight trajectories.

1.2 Originality

The control scheme we propose and evaluate in this paper has several advantages over the schemes proposed in the scientific and patent literature. It is able to control the tractor-trailer along straight and curved trajectories and it does not require any mechanical modification of the tractor-trailer system other than the ability to sense the hitch-angle. Our control scheme is simple to implement and tune, numerically stable, and easy to predict and understand. Its main drawback is that, in its current state, it is not generalizable to any number of trailers.

2 Experimental platform

The platform used in these experiments is the CSIRO Autonomous Tractor (AT), as shown in Fig. 1(a). It is an Ackerman steered, ride-on mower which has been retro-fitted with an array of actuators, sensors, and a computer system enabling the implementation and testing of control and navigation algorithms. For full details of the vehicle's design, refer to [11].

The trailer hitch-angle is sensed using a set of string-pot encoders. Much of the existing research presumes that the trailer hitch-point lies over the centre of the rear-axle, as seen on a prime-mover and trailer combination. Here we consider the case where the hitch-point is offset from the rear-axle, an example being a standard automobile towing a trailer, which represents a further distinction of this work from the existing literature.

3 Hitch-angle stabilization


Briefly, the method used to control the tractor-trailer hitch-angle when reversing is simply a PI controller:

$$\phi = K_p(\psi^* - \psi) + K_I \int_0^t (\psi^* - \psi) du \quad (1)$$

where ψ^* is the demanded hitch-angle, ϕ is the vehicle steering angle and K_p and K_I are the proportional and integral gains. Due to the non-linear nature of the system, the demand to the system has to be modified in order to reduce



(a) The Autonomous Tractor with its trailer.



(b) Kinematic model of the tractor-trailer system.

Fig. 1. The Autonomous Tractor-Trailer system used in these experiments.

the steady state error when stabilizing to non-zero set-points. The demanded hitch-angle ψ^d that will make the system converge to a desired angle ψ^* can be approximated by:

$$\psi^d(\psi^*) = \frac{K_P L_1 - L + K_P L_2}{K_P (L_1 + L_2)} \psi^* \quad (2)$$

and the controller becomes:

$$\phi = K_P(\psi^d(\psi^*) - \psi) + K_I \int_0^t (\psi^* - \psi) du \quad (3)$$

In this system, the integral term is not essential for the stability of the control. It is added to account for minor inaccuracies remaining after the proportional control. Such inaccuracies can result from the linearization leading to $\psi^d(\psi^*)$, or from errors in the vehicle model (e.g L , L_1 , and L_2 in fig. 1(b)). The hitch-angle controller is described more completely in [8].

4 Stabilization of an articulated vehicle to a trajectory

To stabilize the vehicle to a trajectory – *i.e.* a sequence of vehicle states indexed by time – we use a path tracking controller derived for an articulated vehicle.

Our control law is inspired from [9], which aims at stabilizing the vehicle on a path, *i.e.* a 2D curve in the plane. It relies on three error measurements, as depicted in Fig. 2(a): ξ_y the lateral error, ξ_θ the heading error, and ξ_κ the curvature error. In the original work, the authors controlled the hitch-angle rate of the articulated vehicle using these measurements. Here we control the hitch-angle and the control law is defined as:

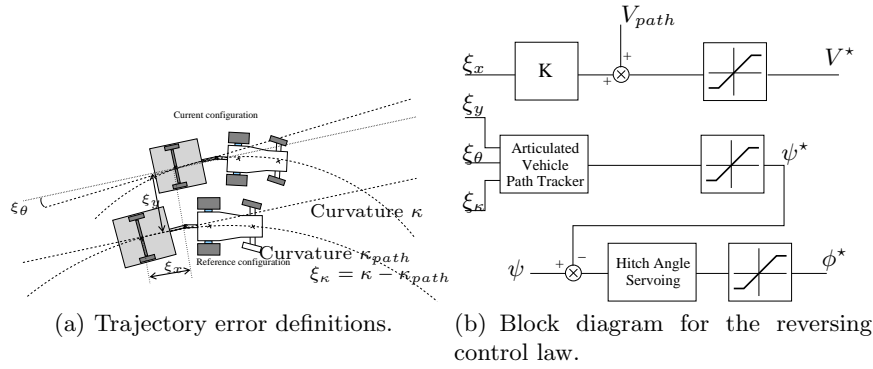


Fig. 2. The trajectory controller and the definition of path errors.

$$\psi = K_y \xi_y + K_\theta \xi_\theta + K_\kappa \xi_\kappa \tag{4}$$

where K_y , K_θ and K_κ are tuning parameters. The resulting hitch-angle is then fed to the hitch-angle stabilization law described in Section 3. Speed is controlled using a standard pursuit approach to maintain the vehicle’s progression along the trajectory and switching logic is used to deal with any changes of direction, as occurs when the onset of jack-knifing is detected or there is a cusp in the trajectory (see [8] for details). The trajectory-tracking system is described in block diagram form in Fig. 2(b).

5 Experimental results

This section presents results obtained when tracking various reversing trajectories, first using odometry based localization, then several tests using an external localization estimate. The external localization estimates are provided by a particle-filter based method using the vehicle odometry and sparse reflective beacons sensed with a front-mounted laser range-finder (details on the localization system can be found in [4]). This localization estimate is ‘drift-free’ but comes at the cost of localization discontinuities when corrections are applied on spotting a beacon. Such discontinuities are especially challenging for the trajectory tracking system.

5.1 Reversing on a circle with fixed radius

In this first set of experiments, we define reference trajectories as arcs of circles of various radii. The pose of the vehicle is computed using odometry information. The advantage of this setting is that the vehicle pose is a very clean and smooth signal. The disadvantage is that odometry localization is known to drift over time, and consequently, is not suitable for real-life, long range, robotic operations.




Fig. 3. Trajectory tracking with odometry based localization

Fig. 3 shows tracking results for a circular path, with radii of 7.5 m and 15 m. As can be observed, the tracking is very accurate for radii of 15 m. For a radius of 7.5 m, the maximum achievable curvature is reached and the system cannot converge to the required trajectory. However, it should be noted that the system converges to a stable orbit, which is the best it can do to follow the required curvature.

5.2 Reversing on a badly planned trajectory

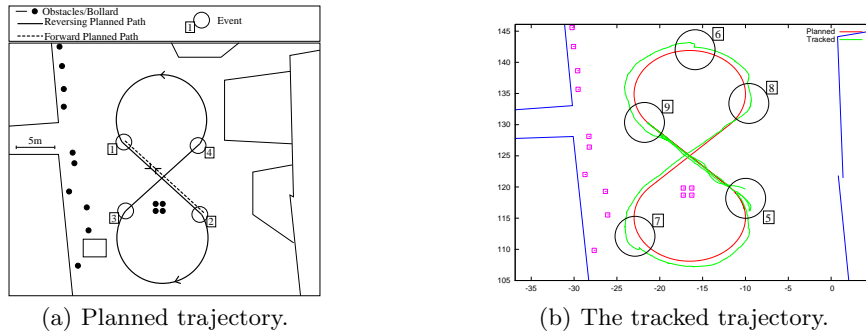


Fig. 4. Performance for a poorly planned trajectory which does not account for curvature limitations or path discontinuities. *events* are denoted by the boxed numbers.

In this second set of experiments, we stretch the system by introducing a more complex path as depicted in Fig. 4(a). We consider this trajectory “badly planned” since it does not take into account the limitations of the vehicle, especially the limited turning radius and the very strong limitation on the hitch-angle rate. In addition, the trajectory contains path discontinuities

in which straight lines join circular segments — such discontinuities are challenging for the control law since they require a very fast change of hitch-angle. Our objective here is to show the performance of our system at the limits of its nominal specification. To further challenge the controller, we also use the external localization system rather than the vehicle odometry.

Fig. 4(b) illustrates how the planned path is tracked by our tractor-trailer systems. From this experiment, we note that on straight line segments the system converges reasonably fast. On the curved segments, the curvature is beyond the vehicle capabilities and the tracked trajectory is offset from the planned one. Transitions between the curved and straight segments of the path (events 5, 8 and 9) introduce a discontinuity which the system struggles to deal with but nevertheless recovers from. Localization discontinuities (events 6 and 7) are handled much more gracefully than the path discontinuities. Finally the effects of the switching logic are illustrated at event 5.

From these experiments, it is clear that in order to obtain very accurate tracking, the path planner needs to take into account the strong constraints of this system, in particular, the small maximum trackable curvature and the need for smooth curvature profiles. However, the trajectory control performed extremely well within these constraints.

5.3 Reversing on a pre-planned trajectory

In this set of experiments, our objective was to demonstrate the performance of our system in a real-world situation. To this end, we designed a path across our research centre. This path is approximately 170 m long and was planned manually using hand-selected way points and interpolating splines. This approach guarantees a smooth and continuous trajectory. For these experiments, the system relied on the external localization estimates.

The resulting path is depicted in Fig. 5(a). It is composed of two segments: a short forward motion to help align the system with the trajectory, and a long reversing segment of approximately 160 m. The main constraints when designing this trajectory were to keep the vehicle close to the middle of the road, and to minimize the path curvature.

Fig. 5(b) and 5(c) show the trajectories realized by our controller on two separate runs. More accurately, the trajectories represent the path of the centre of the AT's rear axle. Examples of localization jumps (corrections) can be seen at the top end of all trajectories, and around (55, 0) in Run *B* and *C*. It is important to note that our system stays stable around the reference trajectory, even when the localization estimation is very unstable. In some situations, e.g. close to $(-10, 10)$ in run *A* and *C*, or close to $(75, 0)$ in run *B*, the localization estimate oscillates between both sides of the reference trajectory. If this occurs at the wrong frequency, it can bring the system to a jack-knife situation. In all cases, the preliminary signs of this event were detected by the system, and a short forward motion was initiated to realign the vehicle with the reference trajectory.




Fig. 5. Reversing trajectories of the autonomous system on two independent runs, compared with a human driver on the same trajectory (Subject 4).

5.4 Evaluation: comparison with human drivers

As a final experimental evaluation of our reversing system, we tested a number of human drivers in performing a trajectory similar to that performed by the autonomous system outlined in Section 5.3. Skill levels of these drivers ranged from novice to professional. For practical reasons, the drivers were required to keep the vehicle in the middle of the road, rather than following the same trajectory as the automation system.

We evaluated their ability to stabilize the trailer, to avoid jack-knife situations, and the statistical properties of their steering input. Figure 5.d shows an example of human controlled trajectory, in comparison with autonomous system ones.

Hitch-angle stabilization

Fig. 6(a) depicts the statistical properties of the measured hitch-angle, for each run, as box diagrams. Each line represents one run, with numbers referring to




Fig. 6. Comparisons of the capabilities of human drivers and the autonomous system.

the different subjects. Lines 10 to 12 refer to the autonomous runs. On these box diagrams, the red circles are points considered outliers, in the context of a normal distribution. They are particularly important since they represent measurements resulting from unusual events, such as the correction of a jack-knife situation.

For the autonomous runs we note that the mean hitch angle is slightly shifted toward positive values. This is expected, since the trajectory is close to a long right-hand turn. When comparing the hitch-angle range of the autonomous system with the human drivers, it is clear that the autonomous system keeps much tighter control of the angle than most of the human drivers and is comparable with the professional driver (driver 1). Nevertheless, it must be noted that this performance is achieved at the cost of a much lower ground speed ($0.3m/s$ instead of $1.0m/s$).

Fig. 6(b) and 6(c) illustrate the properties of the steering angle and its first derivative, recorded for the human drivers and the autonomous system. These graphs demonstrate the reduced control range and capabilities available to the autonomous system in comparison with the capabilities of human drivers — the accessible steering range available to the automation system is reduced by 40% in comparison to the human drivers and the accessible steering speed and steering acceleration are reduced by about 70%. From these observations, it is clear that human drivers use their additional range of control to drive faster while demonstrating similar hitch-angle stabilization performance.

Finally, it is interesting to note that with driving capabilities reduced by 70%, the performance of our system is similar to a professional driver, albeit with a velocity reduced by about 70%. We are yet to determine if this is a coincidence.

6 Conclusion

This article evaluated a new control scheme for a tractor-trailer system. This scheme is based on a two layer control loop: first a hitch-angle stabilization loop controls the angle between tractor and trailer, then a path tracking con-

trol loop, initially designed for an articulated mining vehicle, is adapted to our tractor-trailer system.

The main advantage of this approach over traditional methods is its simplicity of implementation. Only a few parameters need to be tuned and they all have a clear physical meaning. Although simple, this control scheme relies on a mathematically sound background.

Finally, this control law has been implemented on a real vehicle and experiments were conducted on challenging trajectories. Given the limited dynamic performance of our platform (slow response time, loose components, low speed actuation), the control law exhibited excellent convergence and stability properties. Furthermore, our controller also compared well with a range of human drivers on a similar trajectory, even though the control system is significantly handicapped in terms of its actuation capabilities.

Acknowledgments: This work was funded by the CSIRO ICT Centre under the ROVER and Dependable Field Robotics projects. The authors would like to thank the ASL team for their support. Special thanks go to Polly Alexander, Stephen Brosnan, Peter Corke, Elliot Duff, Paul Flick, Leslie Overs, Jonathan Roberts, Ashley Tews, John Whitham and Graeme Winstanley who all contributed to the development of the AT.

References

1. C. Altafani, A. Speranzon, and B. Wahlberg. A feedback control scheme for reversing a truck and trailer vehicle. *IEEE Transactions on Robotics and Automation*, 2001.
2. P. Bolzern, R.M. DeSantis, and A. Locatelli. An input-output linearization approach to the control of an n-body articulated vehicle. *Journal of Dynamic Systems, Measurement, and Control*, 2001.
3. D.N. Carter and J.R. Lormor. Vehicle steering aid system, 2004. GB patent GB2398050.
4. Elliot Duff, Kane Usher, and Peter Ridley. Swing loader traffic control. Technical Report ICT 06/142, CSIRO ICT Centre, Brisbane, Australia, January 2006.
5. R.G. Kandall. Vehicular steering system for reverse paths, 1993. US patent US5247442.
6. S. Karl. Automatic reverse steering system for passenger vehicles with trailer, 1987. DE patent DE3538338.
7. J.R. Koza. A genetic approach to finding a controller to back up a tractor-trailer truck. In *Proceedings of the 1992 American Control Conference*, 1992.
8. C. Pradalier and K. Usher. A simple and efficient control scheme to reverse a tractor-trailer system on a trajectory. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, Roma, Italy, 2007.
9. P. Ridley and P. Corke. Load haul dump vehicle kinematics and control. *Journal of Dyn. Sys, Measurement and Control*, 2003.
10. S.D. Robert. Trailer backing up device and method, 2004. US patent US2004215374.
11. K. Usher. *Visual homing for a car-like vehicle*. PhD thesis, Queensland University of Technology, Brisbane, Australia, June 2005.