Provably safe composition of two autonomous driving functions via set-based contracts

Stanley Smith, Petter Nilsson, Necmiye Ozay
EECS, University of Michigan, Ann Arbor

This poster highlights our recent work [6], in which we present a method of designing multiple controllers for a set of systems that interact with one another. We design the controllers such that the composition of the systems has guaranteed safe behavior. Rather than considering the behavior of the systems as a group during our analysis, which can be cumbersome when the number of systems becomes very large, we opt for a modular design technique in which we consider each system individually. For every system, we select a specification (referred to as a contract) for that system to satisfy, and then design a controller accordingly, under the assumption that all the other systems also fulfill their own contracts. With this assumption, the effects that the other systems have on the system in consideration are restricted to an allowable range. We then quantify this bounded system interdependence and use the result during the controller design process. The resulting controllers are able to ensure that all system contracts are met, as well as other objectives.

The advantage in such an approach is that it can solve issues that arise when systems are integrated into a final product. Often times, products which require massive system integration suffer recalls due to inherent difficulties in their design process [3, 5], which is part of the motivation for this research. In order for engineers to trust that these systems will work properly when integrated, our techniques can be applied. We can think of the *system contracts* as a form of *mutual trust* between the separate systems. In the case where these systems are designed by different companies, the contracts are comparable to how companies place trust in each other in the form of coordinated design specifications. One company promises that it will fulfill certain criteria, while the other company trusts that said criteria will be met, and vice versa. In our paper, we apply this technique to two autonomous driving functions, and these system contracts come in the form of controlled invariant sets. Ultimately, these sets allow us to provide safety guarantees regarding the behavior of the vehicle, which is crucial in order to show the public that these vehicles are reliable and can pass federal regulations.

Proving to the public that autonomous systems are reliable is a vital first step towards a society which trusts autonomous systems. We believe that in order for society to place its trust in autonomy, it must be fully assured that any automated systems in question will operate as expected. Provable safety guarantees are a convenient and accurate method of providing these assurances. In products where safety is critical, such as autonomous vehicles, these guarantees may help provide a measure of whether or not a vehicle passes certain vehicle regulations. This is becoming increasingly important as more states are starting to pass legislation regarding the safety standards of autonomous vehicles. Furthermore, in light of rapid increases in the complexity of such automated systems, we suggest the contract-based design approach so that this technique will also scale well.

We apply our methods to two autonomous driving functions: adaptive cruise control (ACC) and lane keeping (LK). The ACC system ensures that the vehicle either remains a safe distance from cars ahead or maintains a desired velocity when it is safe to do so. The LK system keeps the vehicle centered in its lane. These systems exhibit interdependence on each other in that the steering dynamics of the vehicle are affected by its forward speed, and vice versa [7, 4, 1]. Therefore, we model the systems as being linear parameter varying, where the parameter of each system is the state of the other system. This implies that our system contracts are in the form of state constraints for individual systems. Thus, we seek controlled invariant sets for the systems, described as polyhedra [2], such that each system's state can remain inside its controlled invariant set as long as the other system does the same. For example, we assume that the ACC controller keeps the vehicle's velocity within a certain range and then design a LK controller while adhering to this assumption.

To approach this problem, we first choose safe sets corresponding to some initial state constraints for each system and then tighten these constraints as necessary in order to obtain controlled invariant sets. This allows us to use the safe sets as system contracts (since the state of each system will remain in its controlled invariant set) which restrict the interdependence between the two systems. That is, the safe sets are state constraints, and therefore impose constraints on the systems' linear parameter varying models, since each uses the state of the other system as its parameter. We seek to quantify this interdependence in a way that is useful during the process of finding controlled invariant sets. This quantification is as follows: for each system, we seek a family of systems such that its convex hull covers the range of the linear parameter varying model, given that the domain of the model is the safe set of the other system. We reduce this problem to that of finding a covering for the range of a properly defined function. Given this new formulation, we suggest a few methods for its solution based on whether or not the resulting function has certain properties - specifically, convexity and monotonicity. This quantification is then used during the process of shrinking

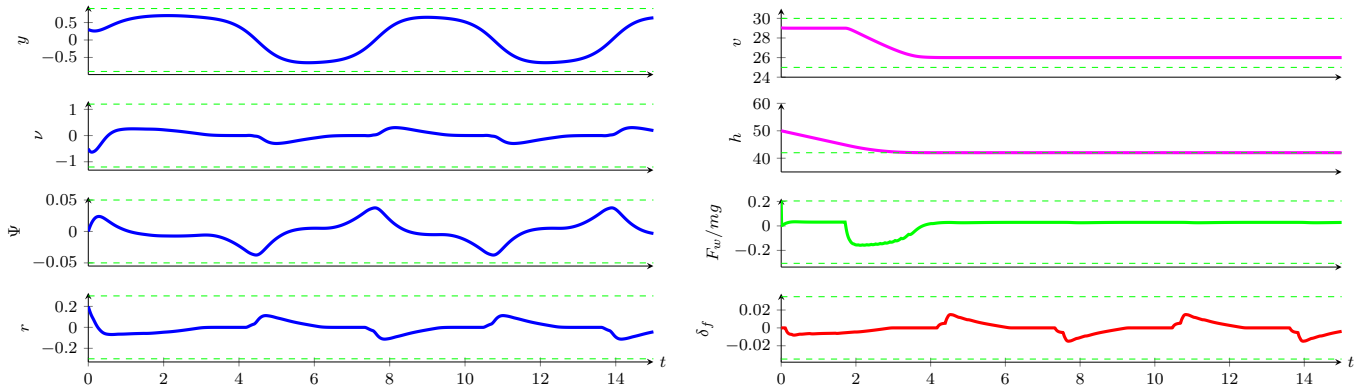

Fig. 1. Simultaneous simulation of the ACC and LK systems on a curvy road. The ACC system has states h and u, which are headway and velocity, respectively. The LK system has states y, ν, $\triangle\Psi$, and r, which are lateral displacement from the center of the lane, lateral velocity, yaw angle, and yaw rate, respectively. The input to the ACC system is the applied longitudinal force F_w and the input to the LK system is the steering angle δ_f. The dashed green lines indicate the initial safe sets which are guaranteed to bound the states at all times.

the initial safe sets down to controlled invariant sets. We also prove that this approach does not introduce any conservatism into our computation of controlled invariant sets.

To keep the ACC and LK states in their respective controlled invariant sets, we enforce the corresponding state constraints with two model-predictive controllers with quadratic costs. At every time step, the controllers find an input to each subsystem that minimizes the cost of its predicted next state, while also keeping it within its controlled invariant set. By keeping the predicted state of both systems in their controlled invariant sets, the composition is guaranteed to be safe. Furthermore, the ACC and LK controllers take into account the state of both systems so that their interdependencies can be accounted for (see Figure 2). In short, finding control inputs amounts to enforcing certain linear constraints which by construction are feasible.

Our synthesized controllers are tested in Simulink (also in Carsim, not shown), where we add disturbances to the systems corresponding to a curvy road. We also add a slower lead car, forcing the ACC controller to decrease the vehicle's speed halfway through the simulation. Our results show that the states of both systems remain within their initial safe sets as expected. The state constraints corresponding to these initial safe sets are shown in Figure 1 alongside the state of each subsystem during the simulation (note that the signals are color coded the same in each figure for clarity). These results show that our method allows us to make safety guarantees for complex integrated systems as the theory suggests. Providing these safety guarantees is a necessary step towards a society that fully trusts autonomous systems.

Acknowledgments: This work was supported in part by NSF grants CNS-1239037, CNS-1446298 and ECCS-1553873.

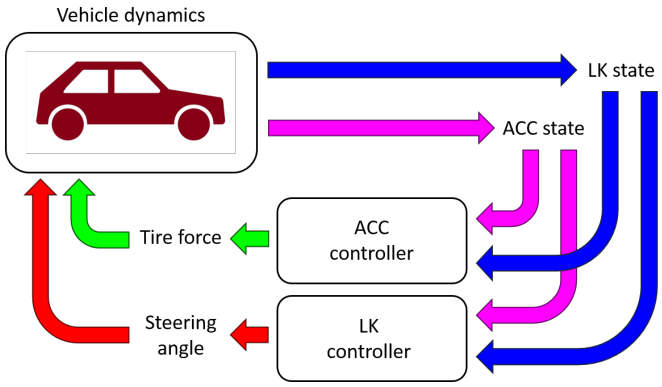

Fig. 2. Block diagram demonstrating simultaneous implementation of the two controllers. Color coding of the signals follows those in Figure 1.

References

[1] A. D. Ames, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *Proc. of the IEEE CDC*, pages 6271–6278, 2014.

[2] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[3] D. Greising and J. Johnsson. Behind Boeing's 787 delays. *Chicago Tribune*, 10:2007, 2007.

[4] P. Nilsson, O. Hussien, Y. Chen, A. Balkan, M. Rungger, A. Ames, J. Grizzle, N. Ozay, H. Peng, and P. Tabuada. Preliminary results on correct-by-construction control software synthesis for adaptive cruise control. In *Proc. of the IEEE CDC*, pages 816–823, 2014.

[5] R. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. Cyber-physical systems: the next computing revolution. In *Proc. of the ACM DAC*, pages 731–736, 2010.

[6] S. Smith, P. Nilsson, and N. Ozay. Interdependence quantification for compositional control synthesis: An application in vehicle safety systems. 2016. preprint.

[7] K. L. Talvala, K. Kritayakirana, and J. C. Gerdes. Pushing the limits: From lane keeping to autonomous racing. *Annual Reviews in Control*, 35(1):137–148, 2011.