Passivity based Tools for Intelligent Transportation

Panos J. Antsaklis and Vijay Gupta, Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556
Shige Wang, General Motors Research and Development Center, Warren, MI 48090
Emails: antsaklis.1@nd.ed, vgupta2@nd.edu, shige.wang@gm.com

While transportation networks and large scale systems studied traditionally in distributed control share many similarities – in the sense that both consist of multiple sub-systems that need to be controlled locally to obtain a global objective – there are significant differences between the two. The level of autonomy enjoyed by individual vehicles in a transportation network (whether human driven or completely autonomous) is much more and the global performance arises from a competition among selfish agents. Nevertheless, we contend that several tools from distributed control will be very useful in designing transportation networks that perform well.

In particular, the classical tools of dissipativity and passivity have shown great promise in the design of transportation networks and need to be developed and understood more fully in this context. Passivity (and more generally dissipativity) is an energy based notion that has its roots in circuit theory. Roughly speaking, a system is passive if the energy stored internally in the system is less than the energy supplied to it by any external input; thus, the system does not generate energy. A passive system is stable, which is the first requirement in any designed system. Moreover, passivity is compositional in the sense that two passive systems interconnected in feedback or parallel structure remain passive – notice that stability may not be preserved under such an interconnection. There is a large literature on designing control laws that ensure passivity of the closed loop system.

In traffic networks, passivity is likely to be useful in at least two aspects. The first is intervehicular – how to guarantee safety and comfort as multiple vehicles interact on the road while being controlled locally? As an example, consider cars moving in a single lane with each car trying to maintain a constant distance from its predecessor. Many control laws can be designed to achieve this goal in the presence of disturbances acting on every vehicle. However, as shown in Fig 1, if the controller is additionally designed to be passive, then the response of every vehicle is very smooth. This feature is useful both for safety (an aggressive controller can lead to inter-vehicular distances being too small at least transiently) and human comfort (the vehicle does not accelerate and decelerate rapidly). Such features of the response are intricately related to the structure of the passive plant; for instance, to the pole-zero difference of the plant. However, much work needs to be done to outline this relation precisely:

1. What is the tradeoff between safety / comfort on one hand and transient performance as measured by the inter-vehicular distance on the other?
2. How do we design passivating controllers that allow us to span this trade-off?

3. What are the fundamental limits on the performance achievable using a passivating controller for a given automobile model?

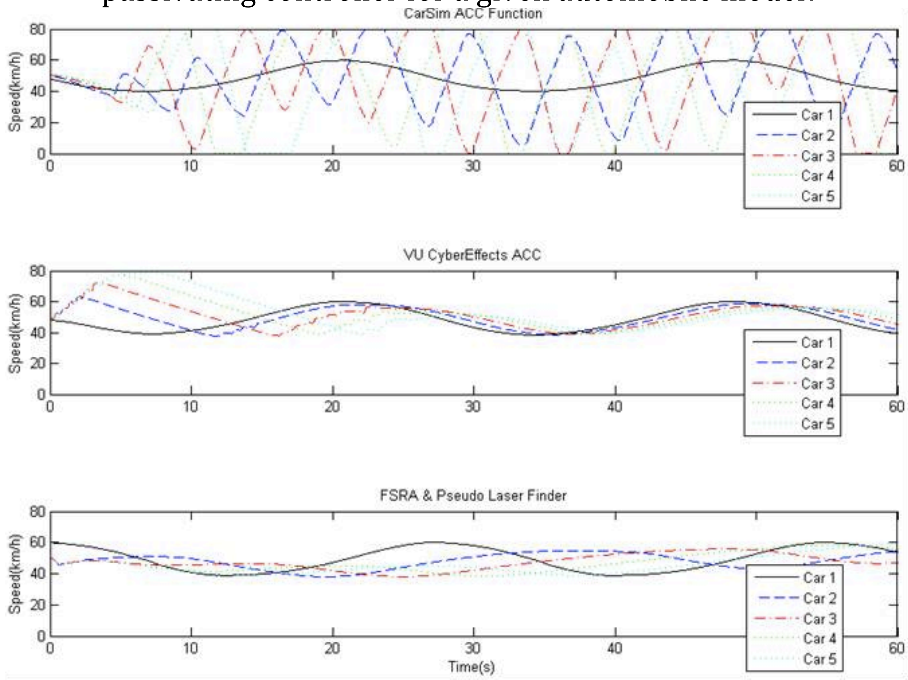

Fig 1: Velocity in lane following maneuver with a passive controller employed by the vehicles. The first subplot is a non-passive controller and the bottom plot is the passive controller. The passive controller leads to much smoother changes in velocity.

The second aspect in which passivity provides great benefits is ensuring compositionality within the design of an individual vehicle. Even as more and more functions in a modern automobile are being automated (even to the extent that a car is now a 'network of computers'), it has been hard to prove that no unintended consequences will arise because of interactions between the various control loops. Previously such interactions might have arisen due to shared microprocessor or communication resources. As more autonomy is imparted to the automobile, and as many such automobiles interact, these interactions will proliferate. Even the decision making of various controllers may now be in direct conflict. As an instance, the automatic cruise controller may conflict with the lane changing controller if the only way to change a lane while avoiding a collision is to reduce speed. If such interactions can be reasoned out, a logical supervisory layer can be developed; however, clearly that is not a scalable solution. Once again, the notion of passivity can be expected to be useful. What is really required is a compositional property, so that each individual controller only takes actions that keep the automobile in a safe zone. Then, no matter how two controllers interact, the car can be expected to be safe. Many research questions need to be addressed in this direction:

1. How do we guarantee safe operation (perhaps using barrier certificates from Lyapunov function theory) when the controller is restricted to be passivating?

2. What is the relation between passivity and traditional concerns in software design, such as verification?
3. How can passivity be guaranteed in the presence of communication buses and software modules that are described using a different mathematics?

Another challenge in intelligent automobiles is that the human driver interacts with the actions taken by other controllers. Many important actions can only be taken at a high enough level by the human driver and the effect of these actions needs to be modeled and studied. Modeling of human as a driver or a pilot has a rich history. One of the most basic effects introduced by a human is that of delay (primarily reaction delay). It is well known that delay can hurt passivity of the closed loop system; thus, the presence of a human can render all previous considerations futile. There may be many mechanisms to deal with this situation, and research needs to be done on many fronts. For instance, there is much research available on ensuring passivity in the presence of communication network induced delay. Roughly speaking, the solution is to transform the variables that are transmitted and received over the network through a wave transformation matrix. In the case of a human driver, these transformations can take the shape of an interface across which the driver controls the automobile. In a mechanical setting of an automobile, how to implement such transformations is an important open question. A different method could be to restrict the human to take actions that do not move the car out of the safe passive zone of operation. This may be achieved through some reward functions that are ingrained in the driver during a training phase. There has been some research in sequential decision making by humans for given reward structures, both in psychology and in control theory. This provides a lead into this exciting possibility. More generally, passivity with a hierarchical controller is an open question that needs to be developed further.

In summary, intelligent automobiles and transportation networks provide many challenges and opportunities. Many tools from diverse areas will be needed to fully design such systems to operate safely and reliably. We believe that traditional tools of passivity and dissipativity, as developed to meet new challenges, are promising. However, much research is needed to solve fully all the questions that are raised.