Safe Temporal Planning for Urban Driving

Bence Cserna, William J. Doyle, Tianyi Gu, Wheeler Ruml

bence, doyle, gu, ruml at cs.unh.edu

University of New Hampshire

Motivation

- The central goal of artificial intelligence is to construct autonomous systems.
- It is becoming more common that these systems interact closely with humans. Autonomous vehicle technology is an auspicious example.
- It is crucial that AI systems humans interact with are designed with safety as a fundamental aspect.

Contributions

- A new technique for urban driving explicitly designed to achieve a safe high-quality plan.
- Empirically demonstrate our safe temporal planner is a significant improvement to conventional methods while still maintaining passenger comfort and safety.
- Theoretical results guaranteeing our approach will find a safe plan and expand only twice as many states as a naïve method.

Problem Setup

- Self-driving vehicles frequently replan to accommodate dynamic obstacles and to ensure safety.
- One way to generate a trajectory is to take a spatial plan, a path, and assign accelerations along the path.
- Previous methods approached this problem by constructing a comfortable trajectory that brings the vehicle to a stop by the end of the spatial plan.

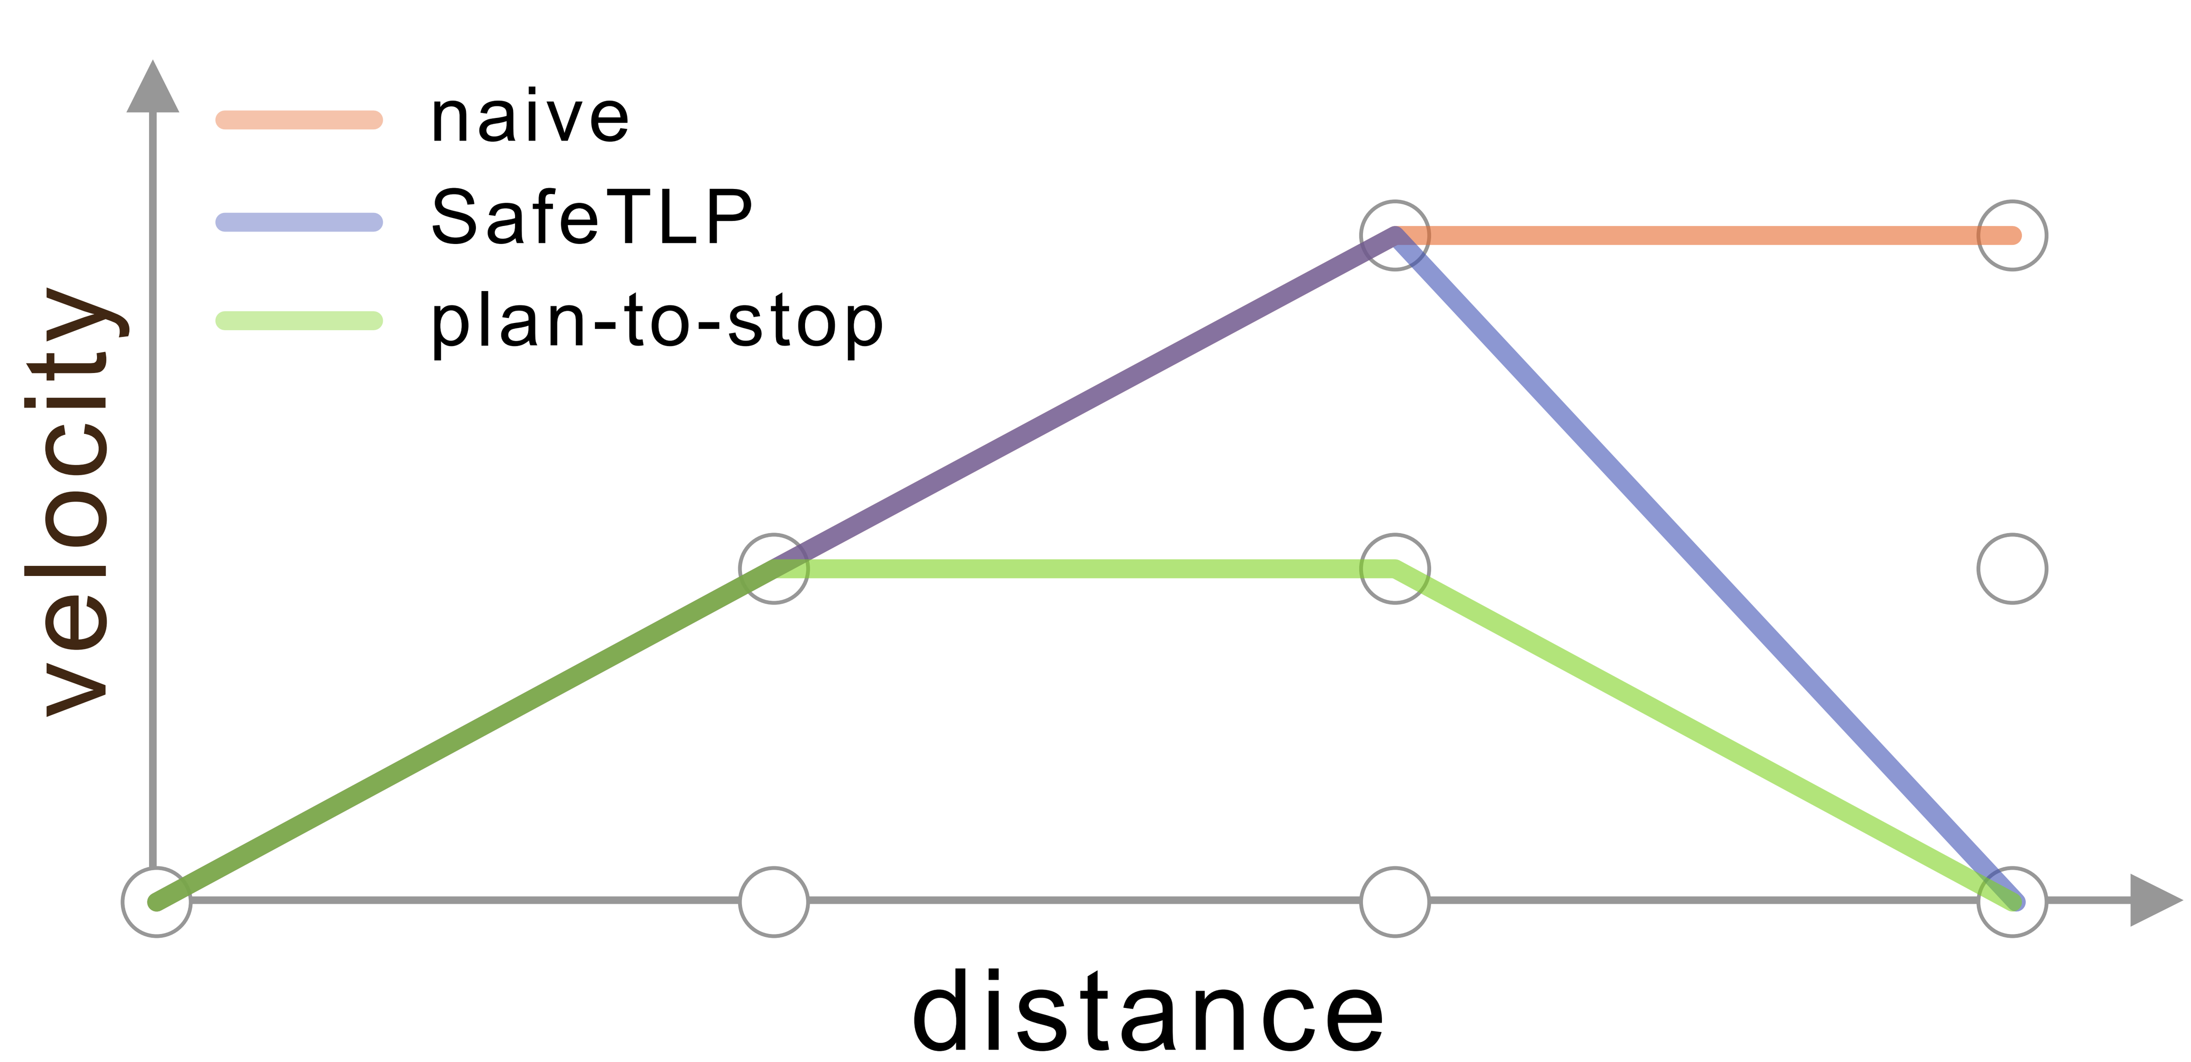

Safe Temporal Lattice Planner

- Our work addresses urban driving at the level of motion planning: we send trajectories to the vehicle such that the motions are safe and comfortable for passengers.
- Instead of doing an exhaustive search we explicitly optimize for the safe solution.

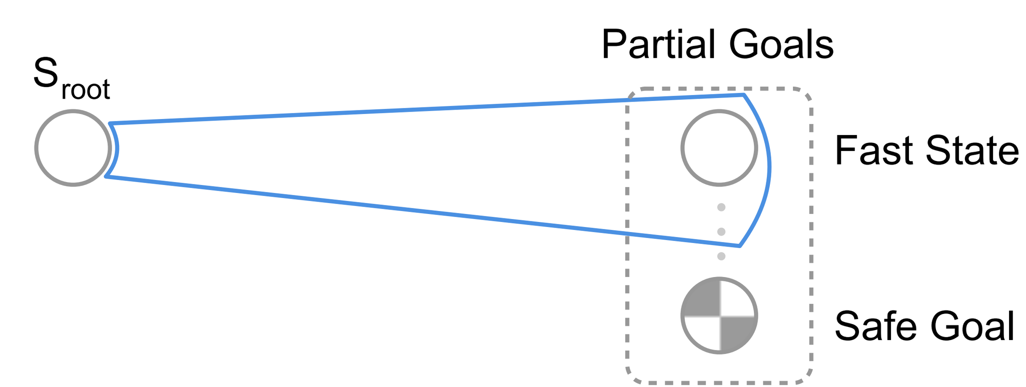

SafeTLP phase 1: search optimized for performance (velocity, comfort, etc.)

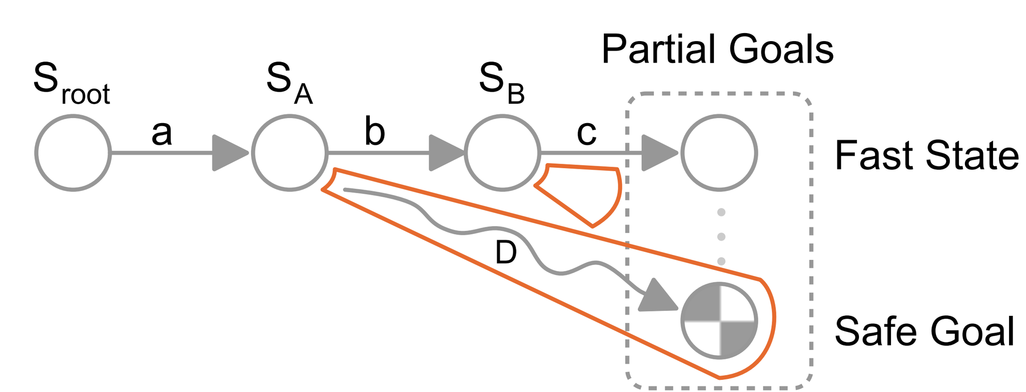

SafeTLP phase 3: search optimized for safety.

Algorithm 1: SafeTLP

Input: s_{root}

perform Best-First Search on lattice to find a potentially unsafe trajectory T from s_{root} to a partial goal

$s_{current} \leftarrow T.last$

while $s_{current}$ *exists* **do**

  perform Best-First Search on d_{safe} from the node $s_{current}$ to s_{goal}

  if s_{goal} *is found* **then**

    cache the safe partial path from $s_{current}$ to s_{goal}

    return $\langle s_{root} \ldots s_{current} \rangle \langle s_{current} \ldots s_{goal} \rangle$

  else

    $s_{current} \leftarrow s_{current}.predecessor$ in P

Best-First Search to find a safe trajectory T from s_{root} to s_{goal}

return T

Empirical Validation

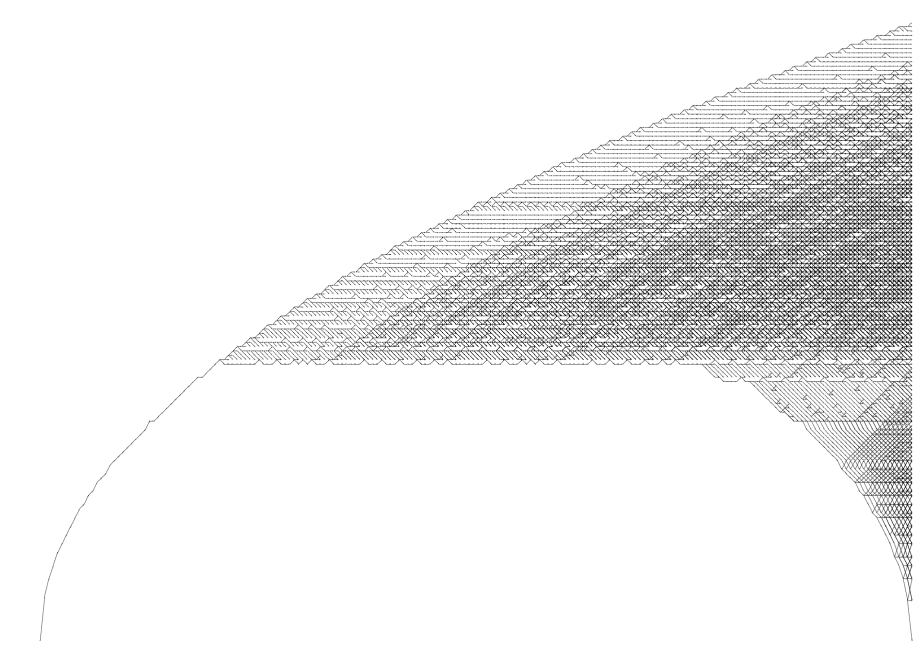

Distance-velocity projection of the plan-to-stop search tree. The horizontal axis represents space and the vertical axis velocity. Note that the temporal aspect of the states is not captured by the projection.

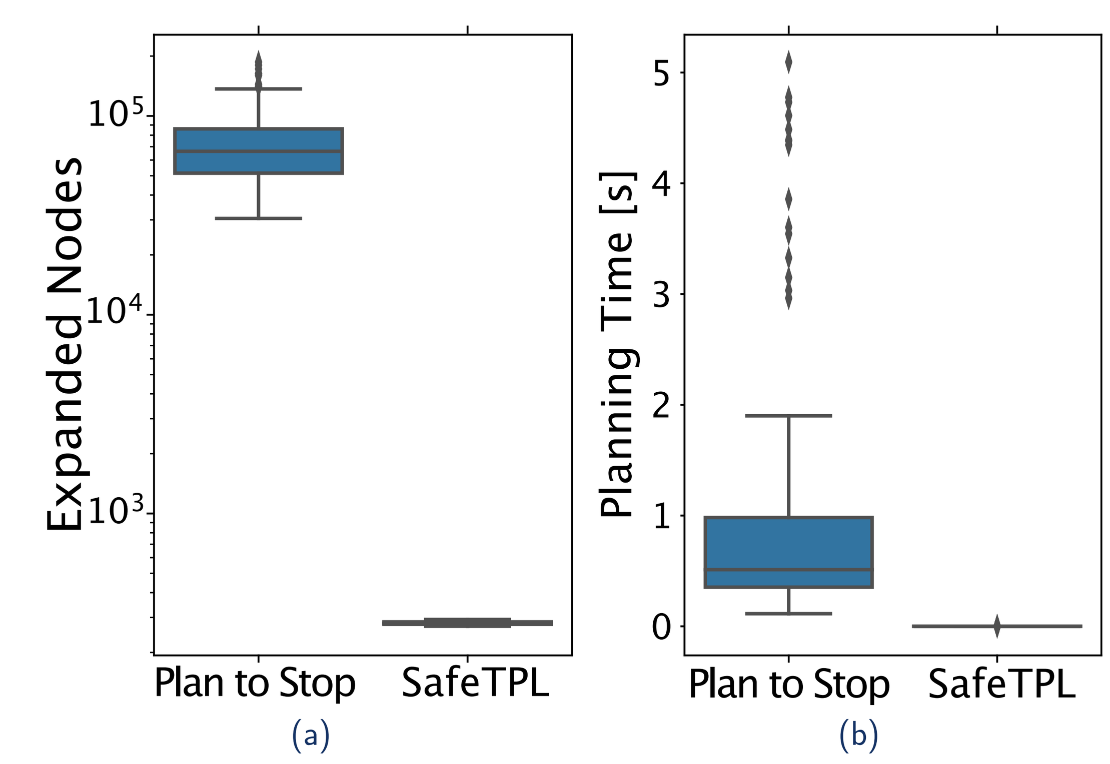

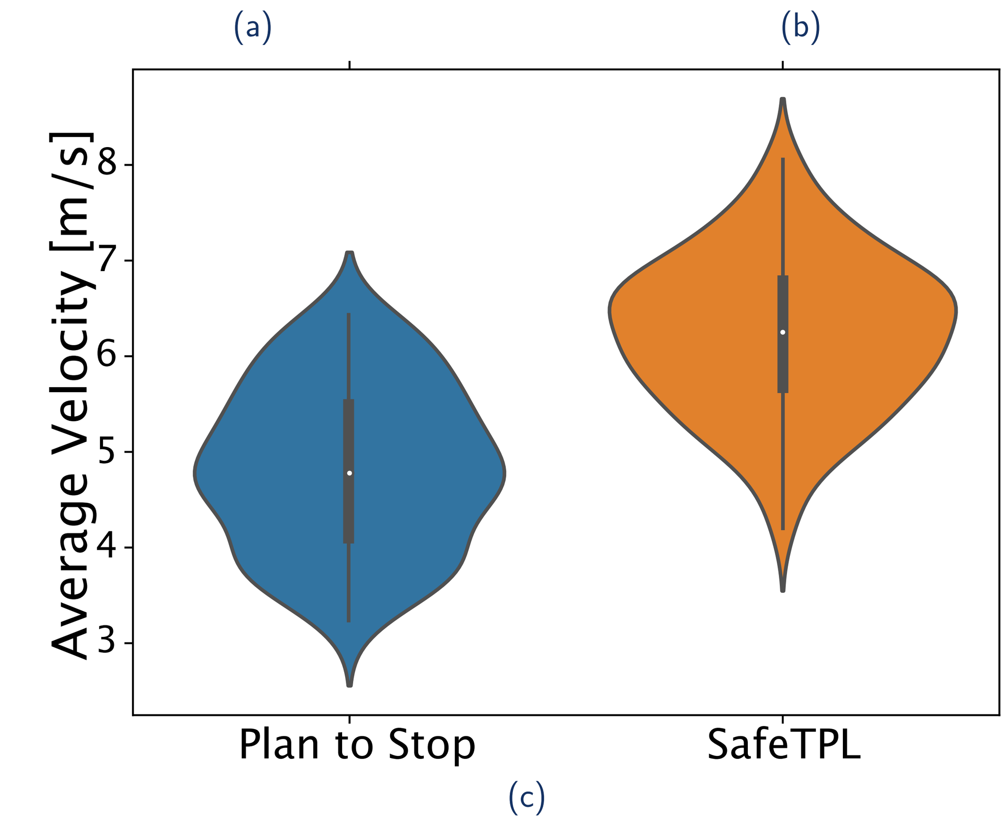

(a): Distribution of the number of expanded nodes during search. (b): Planning time distribution. (c): Average vehicle velocity distribution.

Conclusion

- Introduced a new and more effective method for safe action selection in spatiotemporal planning.
- Other planners expand more nodes than required to guarantee safety; our approach quickly generates safe and comfortable plans online.
- Drastically reduced planning time while maintaining a higher average velocity of the vehicle.
- We've demonstrated that safe real-time heuristic search (Cserna et al., 2018) has important benefits in autonomous vehicles.