A Knowledge-Based Analysis of Intersection Protocols

DISC 2024

Kaya Alpturer[1] Joseph Y. Halpern[2] Ron van der Meyden[3]

[1]Princeton University

[2]Cornell University

[3]UNSW Sydney

Intersections

- Driverless cars will be able to communicate with each other.
    - Going through intersections should become much more efficient, without the need for traffic lights.

Intersections

- Driverless cars will be able to communicate with each other.
    - Going through intersections should become much more efficient, without the need for traffic lights.
- Fault-tolerance is desirable. Solutions need to handle:
    - communication failures of driverless cars
    - human drivers without wireless communication

Intersections

- Driverless cars will be able to communicate with each other.
  - Going through intersections should become much more efficient, without the need for traffic lights.
- Fault-tolerance is desirable. Solutions need to handle:
  - communication failures of driverless cars
  - human drivers without wireless communication
- There are many types of intersections:
  - The model should capture all of them, and solutions need to apply broadly.

Related work

Inefficiency of traffic-light-based intersection management has long been recognized.

Related work

Inefficiency of traffic-light-based intersection management has long been recognized.

To the best of our knowledge, prior approaches mainly focused on:

- Specific intersection scenarios [RBS21, SSP17].
- Leader-election protocols without communication failures [FVP$^+$13, FFCa$^+$10].
- Simulation of scenarios [FFCa$^+$10, RBS21].

Intersection model

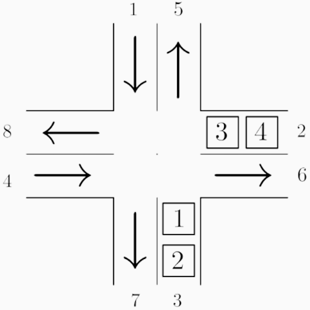

Intersection model

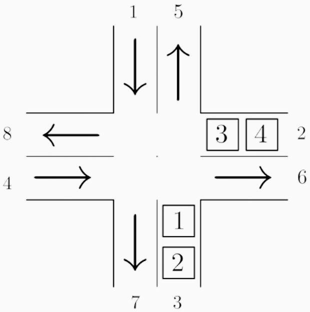

- Vehicles are numbered $\mathsf{Ag} = \{1, 2, \ldots, \ldots\}$ (potentially ∞)

Intersection model

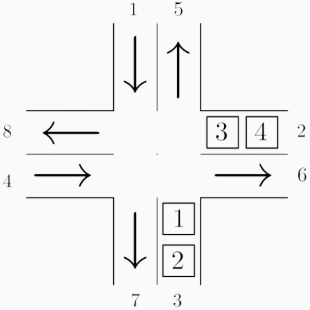

- Vehicles are numbered $\text{Ag} = \{1, 2, \ldots, \ldots\}$ (potentially ∞)
- Incoming lanes are $\mathcal{L}_{in}$, outgoing lanes are $\mathcal{L}_{out}$.

Intersection model

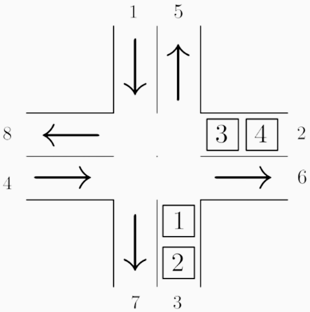

- Vehicles are numbered $\mathrm{Ag} = \{1, 2, \ldots, \ldots\}$ (potentially ∞)
- Incoming lanes are $\mathcal{L}_{in}$, outgoing lanes are $\mathcal{L}_{out}$.
- Each vehicle when in front makes a move (ℓ_{in}, ℓ_{out}).

Intersection model

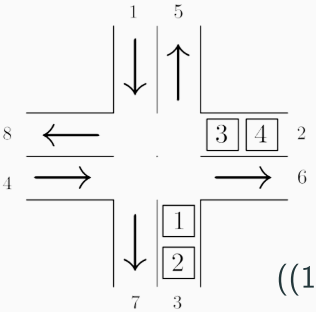

$((1,7),(3,5)) \in \mathcal{O}$

- Vehicles are numbered $\mathrm{Ag} = \{1, 2, \ldots, \ldots\}$ (potentially ∞)
- Incoming lanes are $\mathcal{L}_{in}$, outgoing lanes are $\mathcal{L}_{out}$.
- Each vehicle when in front makes a move (ℓ_{in}, ℓ_{out}).
- A relation $\mathcal{O}$ determines which moves are *compatible*.

Intersection model

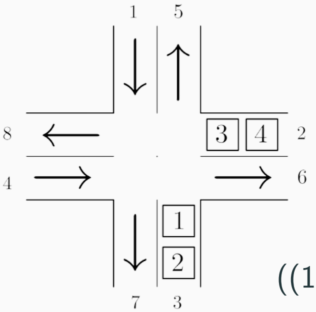

$((1,7),(3,5)) \in \mathcal{O}$

- Vehicles are numbered $\mathrm{Ag} = \{1, 2, \ldots, \ldots\}$ (potentially ∞)
- Incoming lanes are $\mathcal{L}_{in}$, outgoing lanes are $\mathcal{L}_{out}$.
- Each vehicle when in front makes a move (ℓ_{in}, ℓ_{out}).
- A relation $\mathcal{O}$ determines which moves are *compatible*.
- The arrival schedule of vehicles and their planned moves are adversarially chosen.

Intersection model

Two actions available to vehicles are **go** and **noop**.

Intersection model

Two actions available to vehicles are **go** and **noop**.

Communication is broadcast based and potentially faulty:

- Each car can broadcast a message within some radius.
- For our protocols, we only assume that this range is large enough for agents in front to communicate with each other.
- We consider benign communication failures.

Intersection model

Two actions available to vehicles are **go** and **noop**.

Communication is broadcast based and potentially faulty:

- Each car can broadcast a message within some radius.
- For our protocols, we only assume that this range is large enough for agents in front to communicate with each other.
- We consider benign communication failures.

Our goal is to find protocols that satisfy these properties:

- **Validity:** a vehicle executes **go** only when it is in front
- **Safety:** if two vehicles execute **go**, their moves are compatible
- **Liveness:** if a vehicle is in front, then at some point in the future it executes **go**

*Another way to think about this problem is as a generalization of distributed mutual exclusion

Overview

- we use the knowledge-based framework of [FHMV95]

Overview

- we use the knowledge-based framework of [FHMV95]
- we define optimality criteria for intersection protocols:
  → intuitively, vehicles should never wait *unnecessarily*

Overview

- we use the knowledge-based framework of [FHMV95]
- we define optimality criteria for intersection protocols:
  → intuitively, vehicles should never wait *unnecessarily*
- we give constructions that result in optimal protocols:
  → via a "global" to "local" reduction

Reasoning about knowledge

To model knowledge we use the standard runs-and-systems model [FHMV95].

Reasoning about knowledge

To model knowledge we use the standard runs-and-systems model [FHMV95].

- A run r in our system $\mathcal{I}$ is a function mapping a time m to a *global state* $r(m)$ which describes the state of the environment and each vehicle

Reasoning about knowledge

To model knowledge we use the standard runs-and-systems model [FHMV95].

- A run r in our system $\mathcal{I}$ is a function mapping a time m to a *global state* $r(m)$ which describes the state of the environment and each vehicle
- At a point (r, m), we can interpret formulas about the state of the system.

Reasoning about knowledge

To model knowledge we use the standard runs-and-systems model [FHMV95].

- A run r in our system $\mathcal{I}$ is a function mapping a time m to a *global state* $r(m)$ which describes the state of the environment and each vehicle
- At a point (r, m), we can interpret formulas about the state of the system.
- Key definition: $\mathcal{I}, (r, m) \models K_i\phi$ if $\mathcal{I}, (r', m') \models \phi$ at all points (r', m') such that i has the same local state in (r, m) and (r', m')

Reasoning about knowledge

We will design knowledge-based programs where the tests involve the knowledge of the agents:

```
if $K_i$($i$ is the only agent in front of a lane) then go
else follow traffic lights
```

Reasoning about knowledge

We will design knowledge-based programs where the tests involve the knowledge of the agents:

```
if $K_i$($i$ is the only agent in front of a lane) then go
else follow traffic lights
```

We can then find implementations of these programs that only use standard tests on the local states of each agent

```
if $i$ did not receive a reply then go
else follow traffic lights
```

Reasoning about knowledge

We will design knowledge-based programs where the tests involve the knowledge of the agents:

```
if $K_i$($i$ is the only agent in front of a lane) then go
else follow traffic lights
```

We can then find implementations of these programs that only use standard tests on the local states of each agent

```
if $i$ did not receive a reply then go
else follow traffic lights
```

Some useful formulas:

- $\mathcal{I}, (r, m) \models \mathit{going}_i$ if i executes go at (r, m)
- $\mathcal{I}, (r, m) \models \mathit{front}_i$ if i is in front its lane at (r, m)

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Domination-based optimality criteria (from knowledge-based BA literature):

- A protocol is optimal if there is no other protocol that strictly dominates it:
  - P dominates P' if agents in P under the same adversary (i.e. under the same message failures and vehicle arrival schedule) always go through the intersection earlier or at the same time as in P'.

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Domination-based optimality criteria (from knowledge-based BA literature):

- A protocol is optimal if there is no other protocol that strictly dominates it:
  - P dominates P' if agents in P under the same adversary (i.e. under the same message failures and vehicle arrival schedule) always go through the intersection earlier or at the same time as in P'.

Are they equivalent?

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Domination-based optimality criteria (from knowledge-based BA literature):

- ~~A protocol is optimal if there is no other protocol that strictly dominates it:~~
  - ~~P dominates P' if agents in P under the same adversary (i.e. under the same message failures and vehicle arrival schedule) always go through the intersection earlier or at the same time as in P'.~~

Are they equivalent?

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
    - Also used in the distributed mutual exclusion literature

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Lexicographical optimality:

- A protocol is **lexicographically** optimal if there is no other protocol that strictly **lexicographically** dominates it:
  - P **lexicographically** dominates P' if *at the first time the set of agents leaving the intersection differ*, agents in P go through the intersection earlier compared to those in P'.

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Lexicographical optimality:

- A protocol is **lexicographically** optimal if there is no other protocol that strictly **lexicographically** dominates it:
  - P **lexicographically** dominates P' if *at the first time the set of agents leaving the intersection differ*, agents in P go through the intersection earlier compared to those in P'.

Are they equivalent?

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
    - Also used in the distributed mutual exclusion literature

Lexicographical optimality:

- A protocol is **lexicographically** optimal if there is no other protocol that strictly **lexicographically** dominates it:
    - P **lexicographically** dominates P' if *at the first time the set of agents leaving the intersection differ*, agents in P go through the intersection earlier compared to those in P'.

Are they equivalent? Yes, under some conditions.

Optimality criteria

Intuitive criteria:

- *Eliminate unnecessary waiting:* there is no point (r, m) where some vehicle i waits but its move is compatible with all other vehicles that went
  - Also used in the distributed mutual exclusion literature

Lexicographical optimality:

- A protocol is **lexicographically** optimal if there is no other protocol that strictly **lexicographically** dominates it:
  - P **lexicographically** dominates P' if *at the first time the set of agents leaving the intersection differ*, agents in P go through the intersection earlier compared to those in P'.

Are they equivalent? Yes, under some conditions.
Eliminating unnecessary waiting is not always possible but lexicographical optimality can be achieved under weaker conditions.

Optimality vs. Lexicographical optimality

Proposition 1: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ then P is optimal with respect to γ.

Optimality vs. Lexicographical optimality

Proposition 1: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ then P is optimal with respect to γ.

Proposition 2: There exists an intersection context γ with full information exchange and no failures and an intersection protocol P such that P has unnecessary waiting and is optimal with respect to γ.

Optimality vs. Lexicographical optimality

Proposition 1: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ then P is optimal with respect to γ.

Proposition 2: There exists an intersection context γ with full information exchange and no failures and an intersection protocol P such that P has unnecessary waiting and is optimal with respect to γ.

Proposition 3: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ, then P is lexicographically optimal with respect to γ.

Optimality vs. Lexicographical optimality

Proposition 1: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ then P is optimal with respect to γ.

Proposition 2: There exists an intersection context γ with full information exchange and no failures and an intersection protocol P such that P has unnecessary waiting and is optimal with respect to γ.

Proposition 3: If an intersection protocol P has no unnecessary waiting with respect to an intersection context γ, then P is lexicographically optimal with respect to γ.

- The converse is also true (Proposition 4 and 5) if information exchange is “sufficiently rich” and the decision rule only depends on agents in front of each queue.

Intersection policies

We define a “god’s eye” view of the intersection:

- Let an **intersection policy** σ be a mapping from histories to sets of “allowed” moves.

Intersection policies

We define a “god’s eye” view of the intersection:

- Let an **intersection policy** σ be a mapping from histories to sets of “allowed” moves.
- conflict-free: σ always gives compatible moves
- fairness: all possible moves are allowed infinitely often

Intersection policies

We define a “god’s eye” view of the intersection:

- Let an **intersection policy** σ be a mapping from histories to sets of “allowed” moves.
- conflict-free: σ always gives compatible moves
- fairness: all possible moves are allowed infinitely often

Knowledge-based program $\mathbf{P}^{\sigma}$ of an intersection policy:

```
if $K_i(\mathit{front}_i \wedge i$'s move is in $\sigma)$ then go
else noop
```

Intersection policies

We define a "god's eye" view of the intersection:

- Let an **intersection policy** σ be a mapping from histories to sets of "allowed" moves.
- conflict-free: σ always gives compatible moves
- fairness: all possible moves are allowed infinitely often

Knowledge-based program $\mathbf{P}^{\sigma}$ of an intersection policy:

```
if $K_i(\mathit{front}_i \wedge i\text{'s move is in } \sigma)$ then go
else noop
```

Proposition 6: Under some assumptions on the context, if σ is correct, then the knowledge-based program has implementations that satisfy Validity and Safety.

Intersection policies

We define a “god’s eye” view of the intersection:

- Let an **intersection policy** σ be a mapping from histories to sets of “allowed” moves.
- conflict-free: σ always gives compatible moves
- fairness: all possible moves are allowed infinitely often

Knowledge-based program $\mathbf{P}^{\sigma}$ of an intersection policy:

```
if $K_i(\mathit{front}_i \land i\text{'s move is in } \sigma)$ then go
else noop
```

Proposition 6: Under some assumptions on the context, if σ is correct, then the knowledge-based program has implementations that satisfy Validity and Safety.

Proposition 7: For every protocol P satisfying Validity and Safety there exists an intersection policy σ such that P implements $\mathbf{P}^{\sigma}$.

Lexicographically optimal protocols

A lexicographically optimal protocol **P**:

```
if $K_i(\mathit{front}_i \wedge (i\text{'s move is in } \sigma \vee V_i))$ then go
else noop
```

where V_i holds, roughly speaking, if i is in a lane where its move is compatible with all other vehicles that are going (breaking ties in cyclic order).

Lexicographically optimal protocols

A lexicographically optimal protocol **P**:

```
if $K_i(\mathit{front}_i \wedge (i\text{'s move is in } \sigma \vee V_i))$ then go
else noop
```

where V_i holds, roughly speaking, if i is in a lane where its move is compatible with all other vehicles that are going (breaking ties in cyclic order).

Intuitively, **P** allows the following moves:

- all moves permitted by σ
- other moves not in σ in cyclic priority order if each agent knows that its move is compatible with the moves of all agents of higher priority (including agents permitted to go by σ).

Lexicographically optimal protocols

Suppose the following hold in the context

Lexicographically optimal protocols

Suppose the following hold in the context

- synchrony,

Lexicographically optimal protocols

Suppose the following hold in the context

- synchrony,
- at all times m, a cyclic priority order is known by all agents and top priority is assigned to all lanes infinitely often,

Lexicographically optimal protocols

Suppose the following hold in the context

- synchrony,
- at all times m, a cyclic priority order is known by all agents and top priority is assigned to all lanes infinitely often,
- σ is known,

Lexicographically optimal protocols

Suppose the following hold in the context

- synchrony,
- at all times m, a cyclic priority order is known by all agents and top priority is assigned to all lanes infinitely often,
- σ is known,

we get the following result:

Proposition 8, 9: There exists a unique implementation of **P** that satisfies lexicographical optimality, validity, safety, and liveness.

Lexicographically optimal protocols

Suppose the following hold in the context

- synchrony,
- at all times m, a cyclic priority order is known by all agents and top priority is assigned to all lanes infinitely often,
- σ is known,

we get the following result:

Proposition 8, 9: There exists a unique implementation of **P** that satisfies lexicographical optimality, validity, safety, and liveness.

*For example, if the precedence cycle in a given point (r, m) is from $k = m \mod |\mathcal{L}_{in}|$ to $k+1 \mod |\mathcal{L}_{in}|$, $k+2 \mod |\mathcal{L}_{in}|, \dots,$ and $\sigma = \emptyset$ we satisfy these assumptions.

Future directions

- Tolerating stronger adversaries
- Evaluating implementations in other contexts
- Considering strategic behavior

Thank you!

M. Ferreira, R. Fernandes, H. Conceição, W. Viriyasitavat, and O. K. Tonguz.
Self-organized traffic control.
In *Proceedings of the Seventh ACM International Workshop on VehiculAr InterNETworking*, page 85–90, 2010.

R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi.
Reasoning About Knowledge.
MIT Press, Cambridge, MA, 1995.
A slightly revised paperback version was published in 2003.

N. Fathollahnejad, E. Villani, R. Pathan, R. Barbosa, and J. Karlsson.
On reliability analysis of leader election protocols for virtual traffic lights.
In *2013 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop (DSN-W)*, pages 1–12, 2013.

E. Regnath, M. Birkner, and S. Steinhorst.
CISCAV: consensus-based intersection scheduling for connected autonomous vehicles.
In *2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, pages 1–7, 2021.

V. Savic, E. M. Schiller, and M. Papatriantafilou.
Distributed algorithm for collision avoidance at road intersections in the presence of communication failures.
In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1005–1012, 2017.