# WHY TRAFFIC SIGNAL SYSTEMS NEED RICHER DETECTION INPUTS

Authors: Sabeek Pradhan, Collin Barnwell, Matt Stevens, Eduard Strehlau
March 31st, 2026

## ABSTRACT

Conventional detection systems only indicate presence or absence of vehicles in fixed locations. When many detectors are active, traffic controllers have limited information to resolve competing demands, so engineers depend on historical data to create timing plans that encode assumptions about travel patterns.

In contrast, advanced perception systems, similar to those used in autonomous vehicles, can provide continuous 3D position, velocity, and vehicle classification for all road users, and be used as input for a realistic simulation of future behavior. With this more complete picture of movements through an intersection, control logic can make optimal tradeoffs in real time among phases, rather than relying on rules based on historical patterns.

This approach was deployed at a high-volume intersection in San Anselmo, California, and later was expanded to five additional nearby intersections, with plans to install at eight more nearby intersections. Perception was used to estimate positions, speeds, and trajectories for vehicles, bicyclists, and pedestrians. This data was fed into a planning system that generated optimal phase changes to minimize total delay while respecting the timing plans and avoiding unsafe situations.

Observed single-intersection results included a third-party-verified 25% reduction in total delay and 3% increase in intersection capacity, while also supporting safety features such as dilemma zone protection and dynamic extensions of all-red intervals for cyclist safety.

The approach was extended to use this perception data to optimize nearby intersections in real time as a group using Reinforcement Learning, which yielded further reductions in wait times.

The San Anselmo deployment suggests that richer detection and more flexible, dynamic optimization can improve responsiveness, enhance safety, and reduce delays.

# LIMITATIONS OF EXISTING SYSTEMS

## Detection Limitations and Controller Limitations

Existing detection systems, whether inductive loops, video, lidar, radar, or combinations of the above, share a fundamental limitation: they provide input to traffic controllers as binary detection calls. Even if perfectly accurate and timely, these calls do not encode enough information for any system to make optimal decisions, whether optimizing for wait times, safety, emissions, arrivals on green, or any other metric.

Conventional systems running in free seek to address this by using historical traffic data, such as counts and traffic studies, to set timing parameters: minimum green, maximum green, green extensions, and phase sequence. These parameters encode assumptions about expected traffic patterns that may not always be true.

## Coordinated Systems Limitations

Though not commonly thought of as a workaround for low-information detector inputs, the most common response to detection limitations is coordinated signal control: fixing cycle length and synchronizing offsets across a corridor to optimize delay or arrivals on green along an arterial, which can work well for certain roads where demand is predictable and asymmetric, and does not need detector data to work. Nevertheless, this approach has five major drawbacks.

***Static parameters in a dynamic world.*** Coordination assumes fixed vehicle volumes and speeds within each time-of-day period. Macro changes (shifting commute patterns, new development, road restriping) or micro changes (school hours, special events, weather, construction) can render these parameters obsolete.

***Side streets and pedestrians bear the cost***. The cost of this form of optimization is borne by side street traffic and pedestrians, which is controlled for by setting maximum cycle lengths or maximum greens, often based on subjective criteria and resident complaints. This model generally treats the question of side street and pedestrian wait times as a constraint to satisfy rather than an objective to minimize alongside arterial road users.

***Average demand masks real variation.*** Even with perfectly accurate counts and stable demand, random second-to-second variation means a plan optimized for average demand will consistently over-serve some phases and under-serve others.

***Brittle recovery from interruptions.*** Transit signal priority, emergency vehicle preemption, railroad preemption, and similar events break coordination and force re-synchronization across the corridor. Recovery takes several cycles, creating unnecessary delays.

***Forces simplification.*** The requirement to maintain a fixed cycle length and synchronized offsets forces engineers to make simplifying assumptions that reduce the system's ability to handle real-world complexity. Protected left turn phases, pedestrian crossing intervals, and safety features like dynamic all-red extensions or green extensions for dilemma zone protection and pedestrian safety all add variability to phase durations that cannot always be easily absorbed without losing synchronization. In practice, this means pedestrian intervals are baked into the split whether or not a pedestrian is actually present, safety features are left off or treated as acceptable sources of desynchronization, and complex intersection geometries or overlapping corridor networks are either forced into a single coordination scheme that doesn't quite fit any of them or broken into sub-corridors that are optimized in isolation from each other.

## Adaptive Systems Limitations

Properly tuned adaptive systems can address the first limitation of coordinated systems ("Static parameters in a dynamic world") but not the other four. They preserve the same fundamental control structure (cycle lengths, splits, offsets, and phase sequences) and use detector inputs to tune those parameters. Because the underlying detection data remains binary and zone-based, adaptive systems still lack the information needed to correctly anticipate demand and still optimize primarily for arterial throughput rather than all directions of travel simultaneously.

Adaptive systems also introduce the problem of optimizing on backward-looking demand. Cycle-by-cycle parameter adjustments are derived from counts inferred from detector actuations across the stop line which describes what just happened, not what is about to happen. These adjustments to splits and cycle lengths propagate through the network as updated green waves, but the propagation takes time, and conditions may have already changed before the correction reaches downstream intersections. This lag is at the root of many commonly reported problems with adaptive systems: oscillation between over-serving and under-serving phases, sluggish response to sudden demand changes, and difficulty recovering quickly from disruptions.

## Offline Optimization Limitations

AI systems that take an offline approach use historical detector data or phone GPS data to generate improved timing plans, which are then uploaded to controllers and run as static configurations until the next cycle. The limitations are functionally the same as those of coordinated systems, though offline optimization may reduce the engineering effort required to produce well-tuned plans and could partially address the "forces simplification" problem by handling more parameters than a human could.

## ADDRESSING THESE LIMITATIONS

These limitations share a common root: conventional systems lack sufficiently detailed real-time information to make optimal decisions, and every compensating approach is a workaround for that missing information rather than a solution to it.

For any traffic signal system composed of fixed timing plans and any discrete sequence of approaching road users, there is one sequence of signal changes that will minimize wait times across the network. With advanced 3D perception, we seek to find the exact, discrete demand and with our optimization algorithms, we seek to find that signal change sequence. While work remains in order to more accurately perceive the real world, and to more fully explore the space of all possible signal change strategies, we believe this is the first deployed approach with the theoretical basis to find optimal signal changes across the full range of real-world conditions.

### Advanced 3D Perception: How It Works

*"Give me a lever long enough and a fulcrum on which to place it, and I shall move the world."* *– Archimedes*

Rather than monitoring fixed zones for vehicle presence, the Roundabout system uses cameras and computer vision to continuously track every car, bus, bike, and pedestrian in the intersection, producing a real-time 3D map of individual classifications, positions, velocities, and per-road-user wait times, updated 10 times per second. This draws directly from scene-understanding technology developed for autonomous vehicles.

This richer data stream enables control logic that is impossible with conventional detection. A planning system receiving continuous position and velocity data for every road user can evaluate each phase change permitted by the timing plan, estimate the resulting delay for each road user, and select the sequence that minimizes total delay.

**Figure 1:** *Tesla dashboard visualization. Source: Shutterstock.*

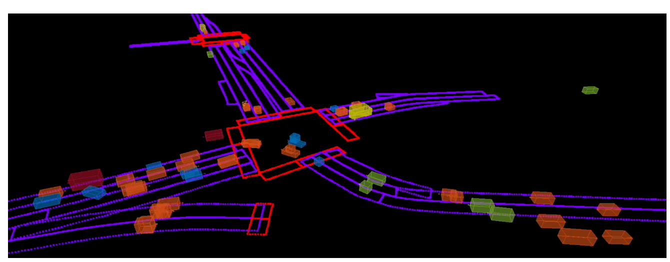

**Figure 2:** *Roundabout visualization*

### Cost-Benefit Analysis

Rich 3D perception allows the system to simulate the per-road-user effects of every available signal change in the timing plan and choose the one with the best outcome.

This can be applied to any observable metric (such as arrivals on green, emissions, and others) but we optimize for wait times because we believe it makes the correct prioritization tradeoffs across all road users and directions of travel.

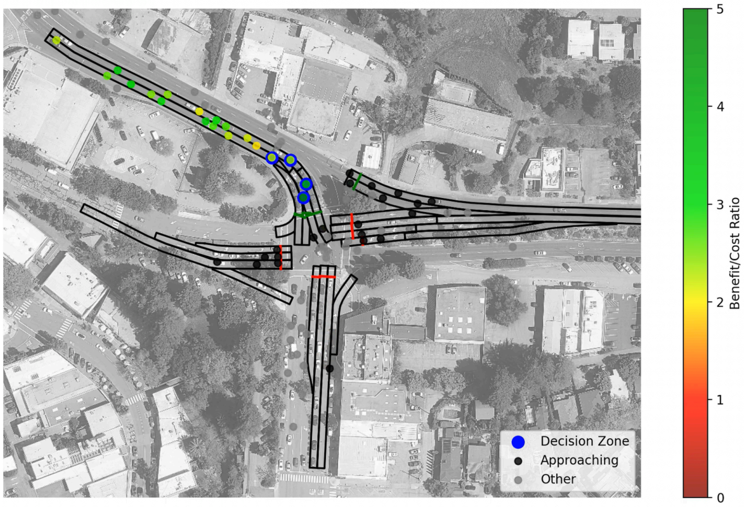


**Figure 3:** *"Cost-benefit analysis" visualization. The benefit realized by each road user is weighed against the cost to all other road users that need to wait in order to realize this benefit. Eventually, the benefit realized by the road user with the green phase is lower than the total cost (cost-benefit ratio < 1) and the system initiates a phase change.*

The most common objection to wait times as a metric is the concern that the system could subject a small minority of road users to extremely long delays in order to reduce total delay. In practice, we have observed the opposite: cycle lengths tend to decrease under this optimization, because the system avoids holding phases longer than the current demand requires. Conventional green extension behaviors driven by gap detectors do not distinguish between a significant volume and a trickle of vehicles, and so they often hold greens longer than optimal.

## Using Perception Across A Corridor With Reinforcement Learning

The benefits of richer detection are amplified across an entire corridor, but so is the complexity. Reinforcement Learning (RL) provides a natural framework for this problem. An RL system operates in a game environment with an observable state, a score, and a set of available actions. Training involves letting an agent play that game millions of times in simulation until it discovers an optimal strategy. This approach has produced superhuman performance in environments as complex as Go,[1] Dota 2,[2] and StarCraft II.[3]

For corridor optimization, Roundabout Technologies built a game environment where the rules are safe but flexible timing plans, the state is the observed positions and velocities of every road user in simulation, and the score is wait times. The agent plays repeatedly and eventually derives strategies that exploit the rich perception data in ways that would be difficult to specify through rule-based logic.

[1] Silver, D., et al. "Mastering the game of Go with deep neural networks and tree search." Nature 529 (2016): 484–489.
[2] Berner, C., et al. "Dota 2 with Large Scale Deep Reinforcement Learning." arXiv:1912.06680 (2019).
[3] Vinyals, O., et al. "Grandmaster level in StarCraft II using multi-agent reinforcement learning." Nature 575 (2019): 350–354.

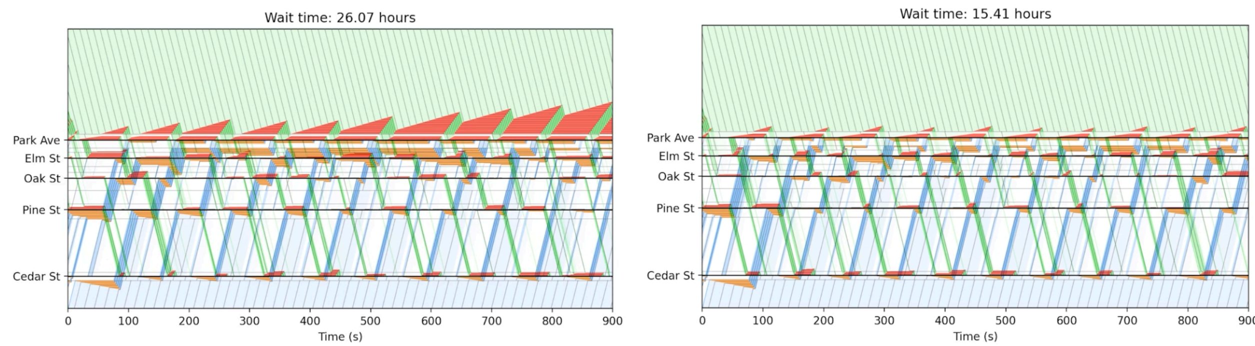


**Figure 4:** *Time-space diagram of the dominant directions of travel for an initial unoptimized run (left) vs fully optimized run (right). Note that these diagrams do not capture wait times experienced by side street traffic or pedestrians which is considered by the optimization algorithm.*

Given the richness of the perception data and a wait-time minimization objective, the agent readily discovers green wave strategies with well-calibrated offsets in any direction, similar to what a traditional coordinated or adaptive timing plan produces. But the system is also capable of finding novel strategies that selectively create and break green waves to handle complex scenarios that coordination-based strategies cannot.

## RESULTS

The results of any comparative optimization study are highly dependent on the baseline being compared against. A well-tuned system will show smaller improvements than a poorly-tuned one, and some geometries and traffic patterns are particularly difficult for conventional systems to handle. With that caveat, two things are testable with the data we have:

1. That a system with real-time 3D perception of all road users can reduce delay compared to conventional detection at the same intersection under the same timing plan.
2. That cross-intersection communication using richer data can produce additional delay reductions beyond what each intersection achieves independently, even for signals that would not be coordinated using conventional methods.

### Single Intersection Optimization

The system was deployed at a high-volume regional bottleneck in San Anselmo, California, which accommodates approximately 50,000 vehicles per day. This intersection, nicknamed “The Hub”, is the intersection of the Sir Francis Drake Boulevard and Red Hill Avenue corridors.

A third-party analysis by the Town's contracted engineering firm, Parametrix, used Synchro and SimTraffic software against two pre-deployment baseline counts (spring 2021 and fall 2025) and found a 3% increase in intersection capacity, a 5–8 second-per-cycle reduction in average pedestrian delay, and an 11%-43% delay reduction for vehicles and cyclists, depending on the methodology and time period, averaging 25%[4]. Roundabout's internal metrics had suggested a 30% reduction at afternoon peak and 10% across most of the rest of the day. Actual improvements were likely larger than either figure captures because, as noted in the Parametrix study, during the baseline period, queues frequently extended far beyond the intersection's field of view but neither the simulation nor the Roundabout system before neighboring intersections came online were able to take this into account.

Table 5 Synchro Model Results, Fall 2025 Counts

| | HCM 2000 (Synchro) | | | | SimTraffic | | |
|---|---|---|---|---|---|---|---|
| Peak Hour and Timing Plan | Actuated Cycle (s) | Level of Service | Delay(s) | Percentage change | Level of Service | Delay(s) | Percentage change |
| AM Peak (8 AM - 9 AM) | | | | | | | |
| Pre-project | 153.5 | D | 46.1 | -12% | C | 34.5 | -32% |
| Post- project (Roundabout) | 136.6 | D | 40.5 | | C | 23.3 | |
| PM Peak (4 PM - 5 PM) | | | | | | | |
| Pre-project | 161.3 | D | 49.8 | -16% | D | 37.6 | -35% |
| Post-project (Roundabout) | 131.9 | D | 41.9 | | C | 24.3 | |

Parametrix, 2026.

Table 6 Synchro Model Results, April 2021 Escalated Estimate

| | HCM 2000 (Synchro) | | | | SimTraffic | | |
|---|---|---|---|---|---|---|---|
| Peak Hour and Timing Plan | Actuated Cycle (s) | Level of Service | Delay(s) | Percentage change | Level of Service | Delay(s) | Percentage change |
| AM Peak (8 AM - 9 AM) | | | | | | | |
| Pre-project | 170.3 | F | 99.1 | -14% | E | 79.7 | -43% |
| Post- project (Roundabout) | 160.8 | F | 85.5 | | D | 45.2 | |
| PM Peak (4 PM - 5 PM) | | | | | | | |
| Pre-project | 179.1 | F | 78.2 | -11% | F | 99 | -27% |
| Post-project (Roundabout) | 168.7 | F | 69.5 | | E | 72 | |

Parametrix, 2026.

**Figure 5:** *Dual methodology (Synchro and SimTraffic) calculations by Parametrix.*

In March 2026, the system was extended to five additional nearby intersections, with an ongoing rollout planned for 14 intersections across Sir Francis Drake Blvd and Red Hill Ave.[5]

## Multi-Intersection Optimization

When the single-intersection system was deployed to 5 additional intersections, wait times at each were reduced by 10–20% depending on the intersection compared to the single-intersection pre-Roundabout baseline. Preliminary data suggests enabling cross-intersection communication adds a further 10% reduction in delay across the network on top of those gains. Further corridor optimization improvements are expected prior to the June presentation.

The system also supported dilemma zone protection, dynamic all-red extensions, and dynamic pedestrian phase extensions. These features likely carry some wait-time cost.

[4] Andrew Lee, PE, TE, PTOE and Sharanya Varma. "The Hub Signal Timing Post-Project Evaluation." Technical Memorandum, Parametrix, March 16, 2026. Prepared for Scott Schneider, PE, Town of San Anselmo. Project No. 474-8968-001.
[5] This project was funded by a Transportation Authority of Marin Innovation Grant using Measure AA tax funds.

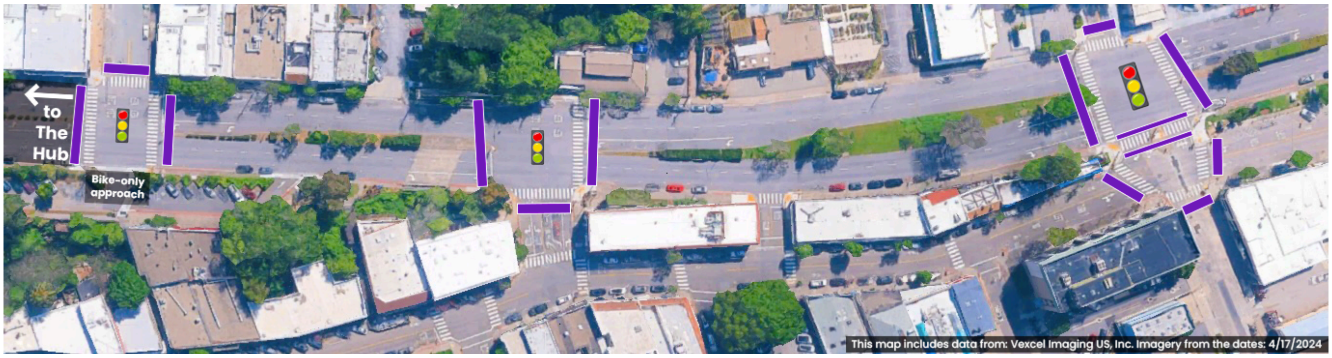


**Figure 6:** *The area's unusual geometry, lane configurations, myriad overlaps, and highly variable pedestrian crossing intervals do not lend themselves to conventional coordination. Prior to deployment, traffic was primarily bottlenecked at the Sir Francis Drake / Red Hill intersection; surrounding intersections ran time of day based coordination.*

## Conclusion

The limitations of conventional signal control share a common cause: binary detection inputs do not provide enough information to make optimal decisions. This is compensated for by using historical data to make assumptions about demand. Though the assumptions differ, this is true for signals running in free, coordinated systems, and adaptive systems.

We seek to reformulate the problem as such: for any traffic signal system composed of fixed timing plans and any discrete sequence of approaching road users, there is one sequence of signal changes that will minimize wait times across the network.

With advanced 3D perception, we seek to find the exact, discrete demand and with our optimization algorithms, both for single intersections and across complex networks, we seek to find that signal change sequence.

The San Anselmo deployment provides evidence that this approach works in practice, with meaningful reductions in wait times for all intersections that were deployed, and further meaningful reductions when cross-intersection communication was enabled, all while implementing safety features at an assumed cost to wait times.

The broader lesson is that signal optimization has historically been constrained not by optimization approaches, which may be optimal according to available data, but by the lack of detailed, real-time information that all of these optimization methodologies assume. With the right sensing layer in place, the problem of optimization takes on a new form, and better performance can be achieved.