A distributed real-time optimisation approach suited to traffic signalling

Ronny Kutadinata, Will Moase, Chris Manzie

Melbourne School of Engineering,
The University of Melbourne, Victoria, Australia

Outline

- Traffic signalling as an optimisation problem
  - Dynamical, time-varying, noisy, large-scale

- Extremum-seeking (ES) for real-time optimisation
  - Single-input single-output (SISO) ES
  - Multiple-input single-output (MISO) ES
  - Distributed ES

- Future work

A case for intelligent traffic signalling

By 2020, traffic congestion is expected to cost Australia $20 billion p.a. cost in wasted productivity and fuel [1].

Significant additional public health costs can also be expected.

Control of signals at freeway on-ramp has achieved a significant increase in both flow rate (~5-9%) and speed (~35-60%) during peak periods on the Monash at very low cost (11 day payback) [2].

[1] Aus. Gov. Dept. of Transport and Regional Economics, 2007
[2] I. PAPAMICHAIL *et al*, *Trans. Research Board Ann. Meeting*, 2010

Optimising urban traffic signalling

Compared to freeways, control of traffic lights for *urban traffic networks* is a more difficult task:

- many more lights to control
- flow not unidirectional
- disturbances due to cars entering and leaving car parks

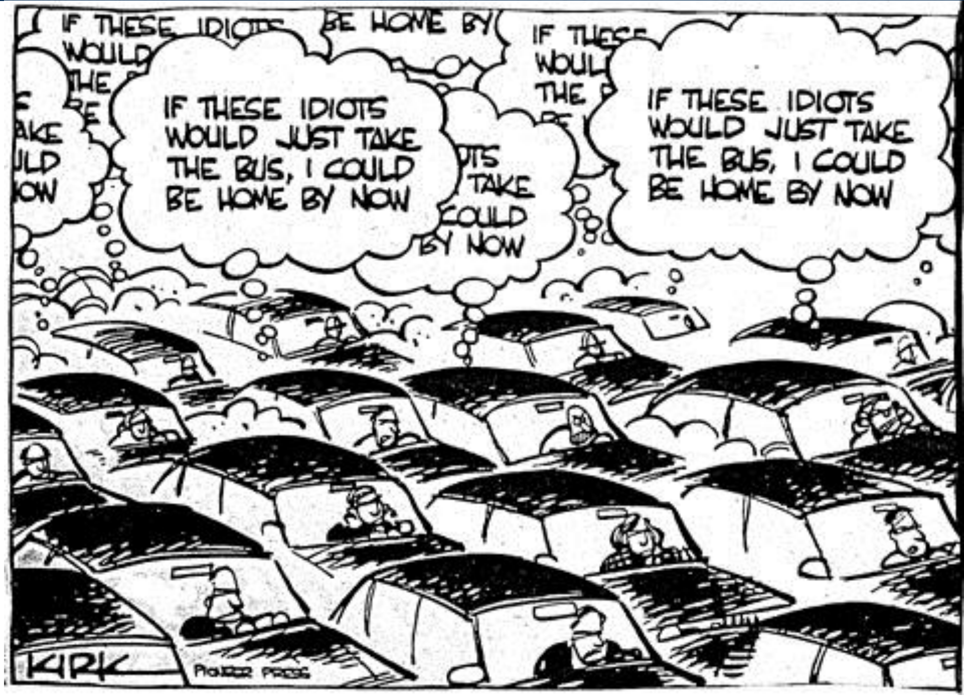

Nonetheless, many strategies exist for urban traffic light control.

These strategies typically have thresholds, parameters or weights that tend to be chosen to achieve “good enough” performance.

Can these quantities be adaptively chosen, using feedback from the traffic network in order to optimise the performance?

A problem with optimisation

A TYPICAL OPTIMISATION

$$y = Q(u)$$

Q is possibly unknown

Goal: Minimise/maximise y

"cost" (*e.g.* estimated total queue length in traffic network) as a function of **"inputs"** (*e.g.* thresholds used in controllers)

Why not use a "standard" optimisation approach?

Optimisation approaches assume a dynamic-less, time-invariant relationship between the input and the cost, but a traffic network is:

- dynamical (*i.e.* when the inputs are changed the cost will pass through some trajectory before settling down to its steady-state)
- time-varying (*i.e.* best inputs may change through the day)
- "noisy"

Outline

- Traffic signalling as an optimisation problem
  - Dynamical, time-varying, noisy, large-scale
- Extremum-seeking (ES) for real-time optimisation
  - Single-input single-output (SISO) ES
  - Multiple-input single-output (MISO) ES
  - Distributed ES
- Future work

What is Extremum-Seeking (ES)?

OPTIMISATION

$$y = Q(u)$$

Q is possibly unknown

Goal: Minimise/maximise y

ES

$$\left.\begin{aligned} \dot{x} &= f(x,u) \\ y &= g(x,u) \end{aligned}\right\} \text{ but } y \to Q(u) \text{ if } u \text{ held constant}$$

f, g and Q are typically unknown

Regulate y to steady-state min/max

i.e. **ES is for steady-state optimisation of dynamical systems (such as traffic networks)**

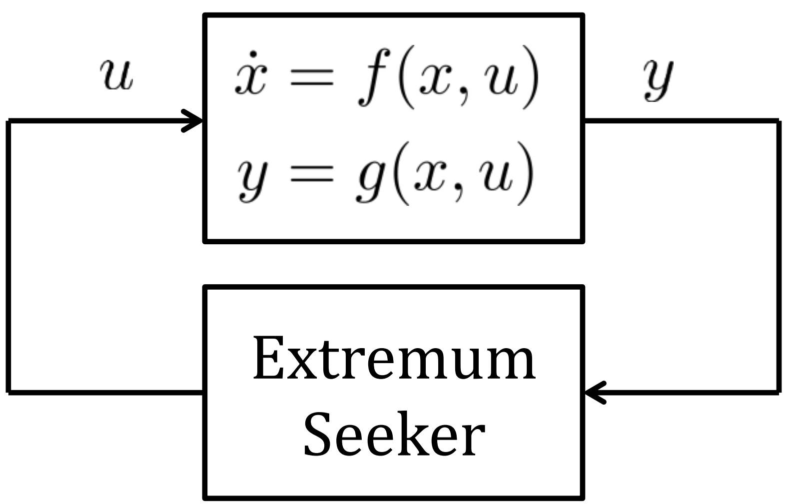

Examples of ES

CONTACTLESS POWER (1922)

Maximise: power transmission

u: inductance

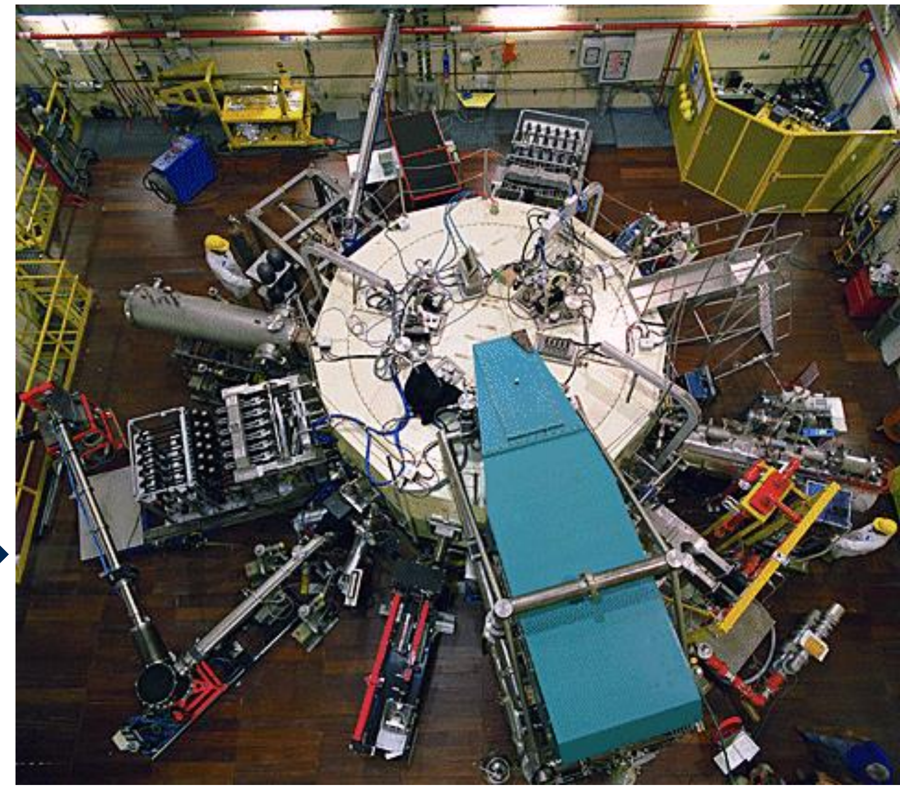

TOKAMAK

Maximise: plasma temperature

u: plasma location

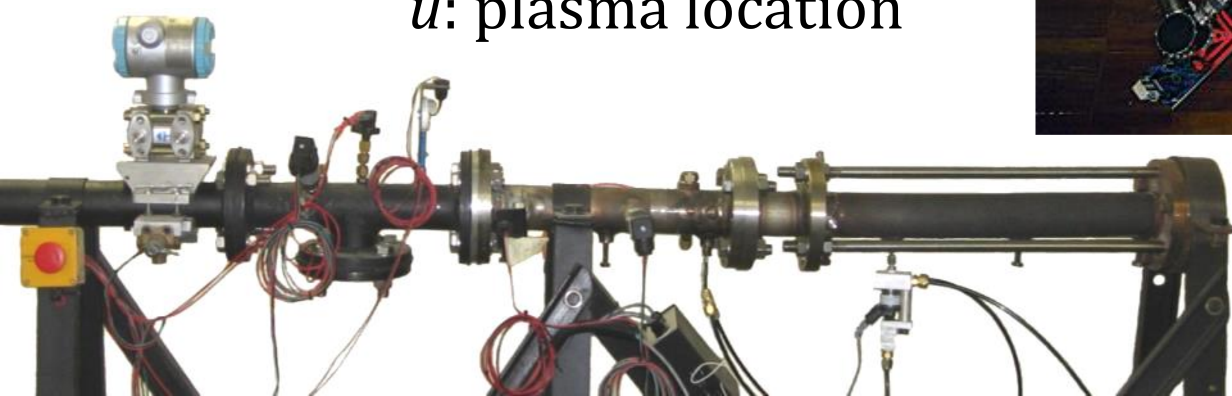

COMBUSTOR

Minimise: “noise”

u: fuel variation

Gradient system

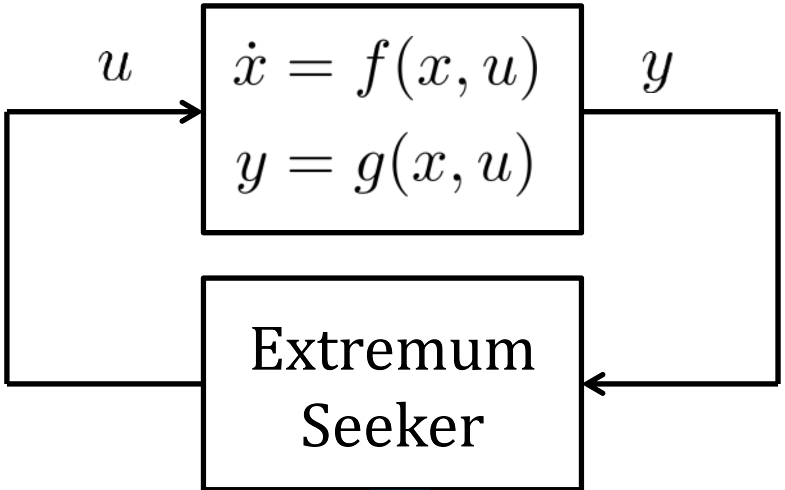

$y \to Q(u)$ if u held constant

How can we make the above closed-loop system approximate this?

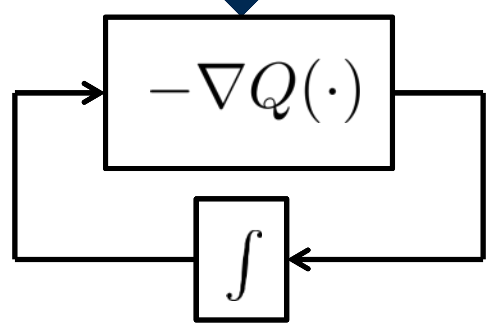

This simple "gradient descent" approach is well-known to minimise $Q(\cdot)$

Single-input single-output (SISO) ES

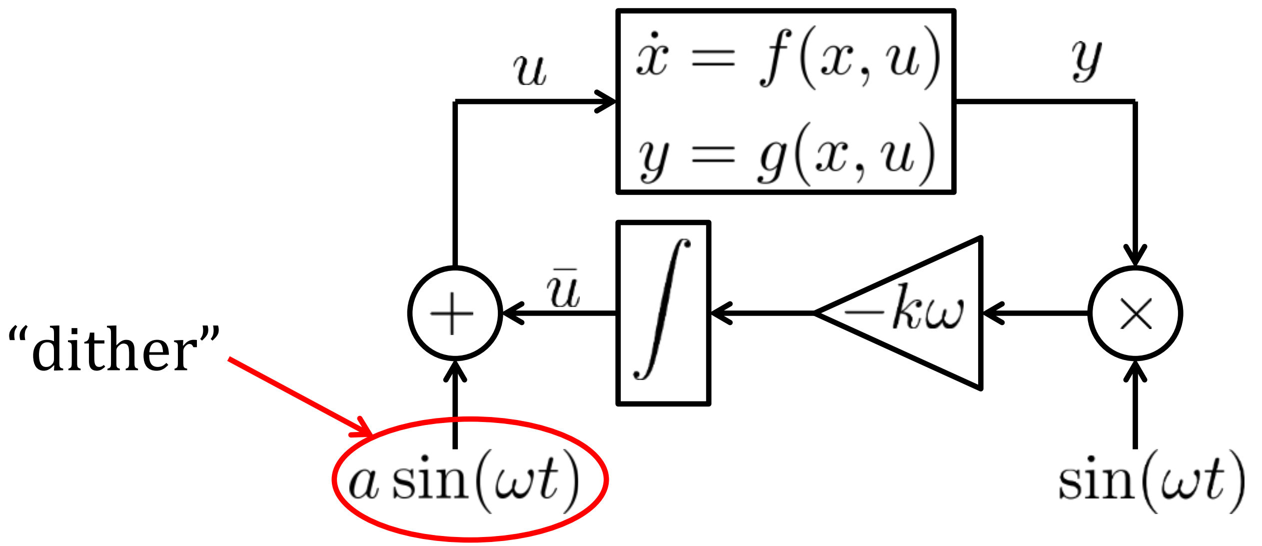

(this is one of many possible ES schemes)

k, a and ω are all small positive reals

$$\frac{\mathrm{d}x}{\mathrm{d}t} = f(x, \bar{u} + a\sin(\omega t))$$
$$\frac{\mathrm{d}\bar{u}}{\mathrm{d}t} = -k\omega g(x, \bar{u} + a\sin(\omega t))\sin(\omega t)$$

SISO ES: sketch of analysis

$\tau := \omega t$

$$\omega \frac{\mathrm{d}x}{\mathrm{d}\tau} = f(x, \bar{u} + a \sin \tau)$$
$$\frac{\mathrm{d}\bar{u}}{\mathrm{d}\tau} = -kg(x, \bar{u} + a \sin \tau) \sin \tau$$

Singular perturbation theory (small *ω*)

$$\frac{\mathrm{d}\bar{u}_{\mathrm{r}}}{\mathrm{d}\tau} = -kQ(\bar{u}_{\mathrm{r}} + a \sin \tau) \sin \tau$$

Periodic averaging theory (small *k*)

$$\frac{\mathrm{d}\bar{u}_{\mathrm{av}}}{\mathrm{d}\tau} = \frac{-k}{2\pi} \int_0^{2\pi} Q(\bar{u}_{\mathrm{av}} + a \sin \tau) \sin \tau \mathrm{d}\tau$$

Taylor series expansion (small *a*)

$$\frac{\mathrm{d}\bar{u}_{\mathrm{av}}}{\mathrm{d}\tau} \approx \frac{-k}{2\pi} \int_0^{2\pi} \left(Q(\bar{u}_{\mathrm{av}}) + a\nabla Q(\bar{u}_{\mathrm{av}}) \sin \tau\right) \sin \tau \mathrm{d}\tau$$
$$\frac{\mathrm{d}\bar{u}_{\mathrm{av}}}{\mathrm{d}t} \approx -\tfrac{1}{2} ka\omega \nabla Q(\bar{u}_{\mathrm{av}})$$

So closed-loop behaves like a gradient-descent optimiser

Multi-input single-output (MISO) ES

$u = [u_1\ u_2\ \ldots\ u_n]^\mathsf{T}$

For each input u_i:

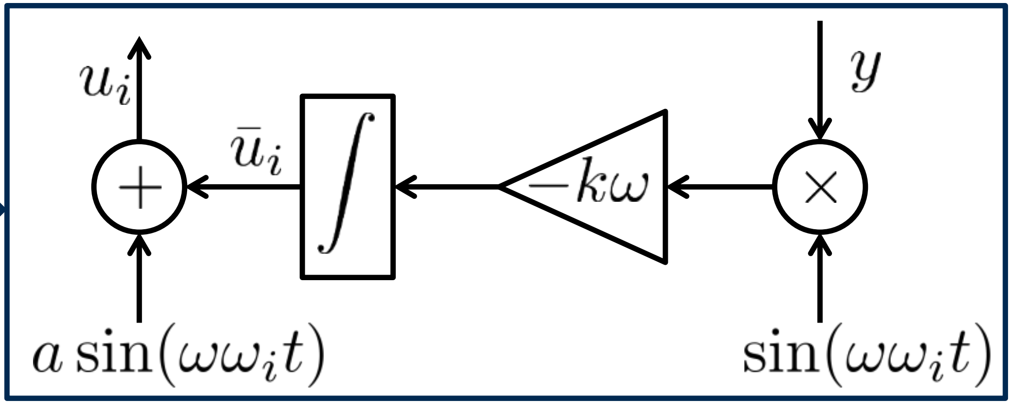

1) ω_i are all positive rationals; 2) $\omega_i \neq \omega_j,\ \forall\, i \neq j$

Analysis follows similar steps to SISO ES. We eventually find:

$$\frac{\mathrm{d}\bar{u}_{\mathrm{av},i}}{\mathrm{d}\tau} = \frac{-k}{T}\int_0^T \left(Q(\bar{u}_{\mathrm{av}}) + a[\sin\omega_1\tau\ \ldots\ \sin\omega_n\tau]\nabla Q(\bar{u}_{\mathrm{av}})\right)\sin\omega_i\tau\mathrm{d}\tau$$

Since $\quad \omega_i \neq \omega_j \implies \int_0^T (\sin\omega_i\tau)(\sin\omega_j\tau)\mathrm{d}\tau = 0,$

Then $\quad \dfrac{\mathrm{d}\bar{u}_{\mathrm{av},i}}{\mathrm{d}\tau} = -\tfrac{1}{2}ka\nabla_i Q(\bar{u}_{\mathrm{av}})$ and $\dfrac{\mathrm{d}\bar{u}_{\mathrm{av}}}{\mathrm{d}t} = -\tfrac{1}{2}ka\omega\nabla Q(\bar{u}_{\mathrm{av}})$

Again, the closed-loop behaves like a gradient-descent scheme

A problem with MISO ES

$u = [u_1 \; u_2 \; \ldots \; u_n]^\mathsf{T}$

For each input u_i:

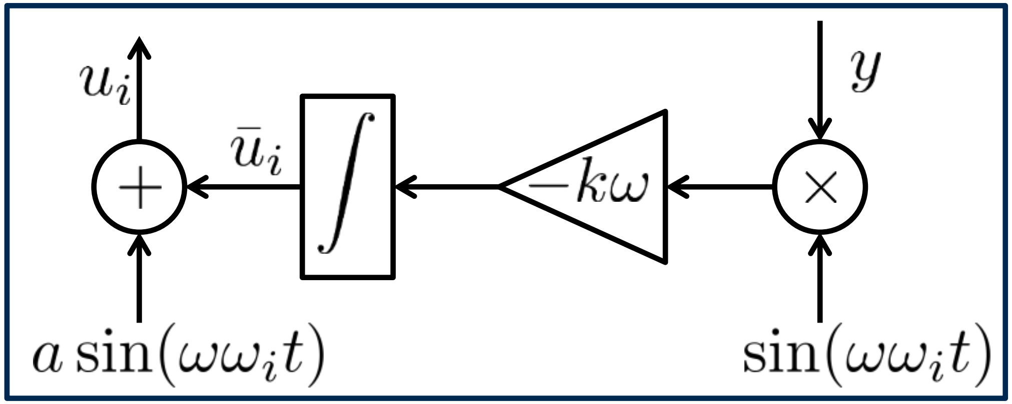

1) ω_i are all positive rationals; 2) $\omega_i \neq \omega_j, \; \forall\, i \neq j$

intuition: unique ω allow the effect of each input on the cost to be distinguished

- In a traffic network, there are many inputs to control.
- For each input, the designer must select a unique dither frequency, each affecting the convergence rate of the optimisation.
- Selecting "good" dither frequencies would be an overwhelmingly laborious task. Can this process be simplified?

Distributed ES

Suppose: instead of there being one globally defined cost function, each input has a corresponding "cost" to optimise.

$$y = g(x,u) \qquad\qquad y_i = g_i(x,u)$$
$$y \to Q(u) \quad \Rightarrow \quad y_i \to Q_i(u)$$

In way, this new problem is no longer centralised. Perhaps it can be solved in a distributed fashion…

The price of de-centralising the problem

- Now each input is "competing" against other the inputs in an attempt to minimise its associated cost.
- We seek a *"Nash equilibrium"* where the change of any given input would result in an increase in its associated cost.
- The Nash equilibrium isn't necessarily the same as the original "global" minimum. However, careful design of g_i can ensure proximity of the Nash and original minimum.

Nash Equilibrium Seeking (NES)

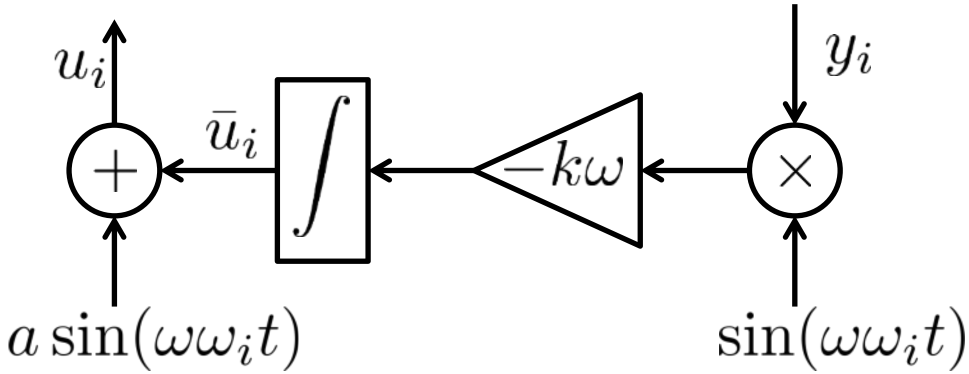

STANKOVIC *et al, IEEE Trans. Automatic Control* (2012)

- Let $\mathcal{N}_i$ be called the "neighbourhood" for i
- $j \in \mathcal{N}_i$ if j affects Q_i
- Must ensure $\omega_i \neq \omega_j$ whenever $i \in \mathcal{N}_j$ or $j \in \mathcal{N}_i$
- But may allow $\omega_i = \omega_j$ otherwise
- Only a few distinct dither frequencies required if $\mathcal{N}_i$ are small

Not the case for traffic networks, where a change in behaviour of a given set of traffic lights can have a far-reaching effect!

Our contribution

KUTADINATA, MOASE and MANZIE, *IEEE CDC*, 2012

Consider systems where an input may affect all measured costs, but its effect dissipates as the "distance" from the input grows.

Define $\mathcal{N}_i(R)$ such that:

- For small R, $\mathcal{N}_i(R)$ contains the "closest" neighbours of i
- As we increase R, $\mathcal{N}_i(R)$ grows to include more neighbours
- For large enough R, $\mathcal{N}_i(R) = \{1, 2, \ldots, n\}$

Result: For large enough R, we may allow $\omega_i = \omega_j$ if $j \notin \mathcal{N}_i(R)$ and $i \notin \mathcal{N}_j(R)$

Early testing on toy (non-traffic) systems indicates in suitably dissipative networks, a small number of dither frequencies may be used without a significant reduction in the fidelity of the optimisation.

Outline

- Traffic signalling as an optimisation problem
  - Dynamical, time-varying, noisy, large-scale

- Extremum-seeking (ES) for real-time optimisation
  - Single-input single-output (SISO) ES
  - Multiple-input single-output (MISO) ES
  - Distributed ES

- Future work

Future work

- Demonstration of distributed ES on a simple traffic grid by adaptively tuning thresholds in "Self-Organising Traffic Lights" (SOTL) control algorithm.
- Increasing convergence speed using "fast extremum-seeking" (for centralised fast ES see MOASE & MANZIE, *Automatica*, 2012)
- Higher-fidelity simulations on more realistic traffic networks (multi-modal traffic).
- Other applications (irrigation and power networks)


THANK YOU
FOR YOUR
PATIENCE
KELLY BROS. SOLAR SIGNS 01454 312675
LR2102
Not in Service
P674 XGJ
LF60 ZSN
LB09 MRY