When Optimal Sensor Interrogation is Rhythmic

Alan C. O'Connor

Harvard School of Engineering and Applied Sciences

December 15, 2010

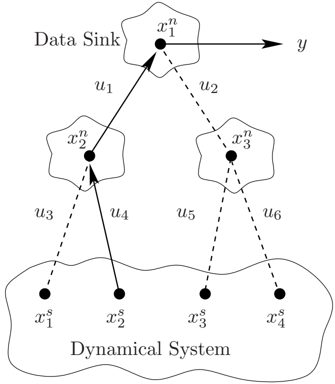

Central question of this presentation:

When can a time-varying policy for controlling the network get us a gain in useful information supplied by the output y ?

Two different approaches to control of a sensor network:

- (**Direct Querying**) The network is transparent from the point-of-view of the central observer. By selecting which network nodes are active, the central observer **directly** obtains observations of different states of the system.

- (**Data Aggregating**) Links in the network are controlled so that intrmediate nodes **aggregate and process data** from lower levels before forwarding summaries to the central observer.

Direct Querying (DQ)

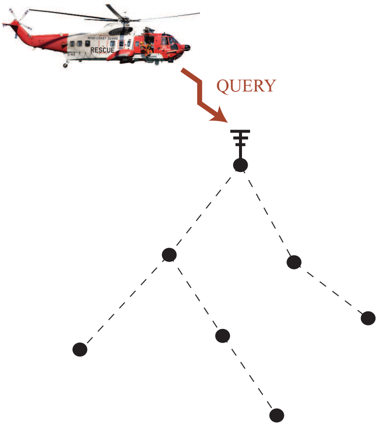

Control Task: select the sensor with which to interrogate the state of the dynamical system

Network: intermediate nodes forward queries and responses between the data-sink and the active sensor

Direct Querying (DQ)

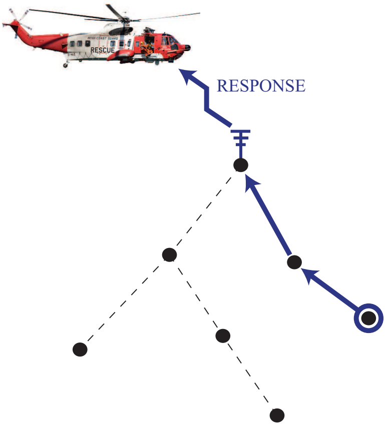

Control Task: select the sensor with which to interrogate the state of the dynamical system

Network: intermediate nodes forward queries and responses between the data-sink and the active sensor

Equations Describing DQ-Type Problem

- $x^s(t)$, system state vector
- A^s, time-invariant linear dynamics
- $D^s\,dw$, standard Wiener noise process
- y observations corrupted by zero-mean Gaussian noise with covariance $R_{ij}\,dt$.

$$dx^s = A^s x^s\,dt + D^s\,dw$$

$$dy = \sum_i u_i(t)\big(C_i\,x^s\,dt + d\nu_i\big), \quad u(t) \in \Omega.$$

Data Aggregating (DA)

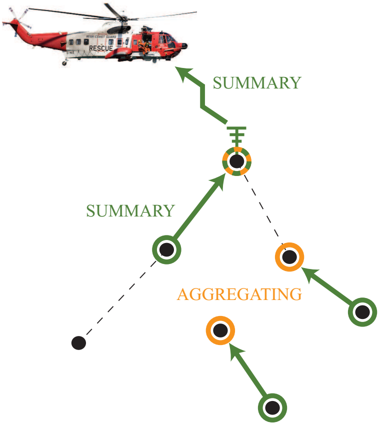

Control Task: schedule the interconnections between the network nodes

Network: nodes have internal states (memory) and may communicate and gather data even when they are not connected to the data sink

Equations Describing DA-Type Problem

- $x^s(t)$, state of the underlying system
- $x^n(t)$, internal state of the network nodes
- B_i, controllable interconnections
- $D^s\,dw$, standard Wiener noise process
- y, observations corrupted by zero-mean Gaussian noise with covariance $R_{ij}\,dt$.

$$
\begin{aligned}
dx^s &= A^s x^s\,dt + D^s\,dw \\
dx^n &= \left(A^n x^n + \sum u_i B_i\,x^s\right)dt, \quad u(t) \in \Omega \\
dy &= C x^n\,dt + d\nu,
\end{aligned}
$$

Time-invariant versus time-varying

We do not consider the problem of fine timing control of the observations, rather the fraction of time that each link is active.

(as in, *e.g.*, Mourikis *et al.*, 2006)

A **time-invariant policy** is a constant average allocation of sensing resources, (even if this is implemented by dithering).

Time-varying policies may include relatively slow variations in the average amount of attention paid to the different sensors.

Error covariance of the Kalman estimator of the system state:

$$\Sigma = \mathcal{E}(x - \hat{x})(x - \hat{x})'$$

Σ obeys a controlled Riccati equation:

DQ–Type $\quad \dot{\Sigma} = A\Sigma + \Sigma A' + DD' - \Sigma\, U(t)\, \Sigma$

DA–Type $\quad \dot{\Sigma} = \big(A + UB\big)\Sigma + \Sigma\big(A' + B'U'\big) + DD' - \Sigma\, C'C\, \Sigma$

Control objective: Pick $U(t)$ to minimize

$$\eta = \frac{1}{T}\int_0^T \mathrm{tr}\,\big(Q\Sigma(t)\big) + \mathrm{tr}\,\big(KU(t)\big)\,dt$$

For the finite horizon problem, the sensor or link allocation that minimizes the cost:

$$\eta = \frac{1}{T}\int_0^T \mathrm{tr}\left(Q\Sigma(t)\right) + \mathrm{tr}\left(KU(t)\right)dt$$

can be found by applying the maximum principle to the control Hamiltonian:

$$\mathcal{H}(\Sigma, \Lambda, U) = \mathrm{tr}\left(\Lambda'\dot{\Sigma}\right) + \frac{1}{T}\mathrm{tr}\left(Q\Sigma + KU(t)\right)$$

Because $U(t)$ appears linearly in $\mathcal{H}$ and takes values in a closed set, maximum principle leads generically to bang-bang controls.

Optimal *constant* policies with $\bar{U}$ in the interior of the control set can occur only if the problem is singular.

(*c.f.* Brockett, *CIS*, 2010)

No Optimal Rhythmic Policies for DQ-Type Problem

Theorem 1. *For DQ-type problems, the average cost*

$$\eta = \lim_{T\to\infty} \frac{1}{T}\int_0^T \operatorname{tr}\big(Q\Sigma(t)\big) + \operatorname{tr}\big(KU(t)\big)\,dt$$

for a time-varying policy $U(t)$ is bounded below by the cost of the steady-state policy

$$\bar{U} = \lim_{T\to\infty} \frac{1}{T}\int_0^T U(t)\,dt.$$

Assumptions: $(A, \bar{U})$ *is observable, the process noise satisfies* $(-A, D)$ *is stabilizable, and* $\Sigma(0)$ *is bounded.*

DQ-Type Impossibility Result: Sketch of the proof

For the DQ-type system, the variance of the optimal estimator of the system state obeys the controlled Riccati equation:

$$\dot{\Sigma} = A\Sigma + \Sigma A' + DD' - \Sigma\, U(t)\, \Sigma$$

The proof also makes use of the differential equation for the inverse of the covariance $X = \Sigma^{-1}$

$$\dot{X} = -A'X - XA - XDD'X + U(t)$$

DQ-Type Impossibility Result: Sketch of the proof

The proof establishes an inequality (in the sense of the partial-ordering of positive-definite matrices) between the time-average of the estimator variance, $\bar{\Sigma}$, and the steady-state variance achieved by the constant policy $\bar{U}$.

The assumptions

- $(A, \bar{U})$ is observable
- $(-A, D)$ is stabilizable
- $\Sigma(0)$ is bounded

are sufficient to ensure that $\bar{X}$ exists and is invertible for all T, and that $X(t)$ converges to a positive-definite steady-state.

DQ-Type Impossibility Result: Sketch of the proof

We also use the inequality

$$\bar{\Sigma} = \frac{1}{T}\int_0^T \Sigma(t)\,dt = \frac{1}{T}\int_0^T X^{-1}(t)\,dt \succeq \Big(\frac{1}{T}\int_0^T X(t)\,dt\Big)^{-1} = \bar{X}^{-1}$$

This follows from Jensen's inequality and the convexity of the inverse when applied to positive-definite matrices.

DQ-Type Impossibility Result: Sketch of the proof

Let Z be the steady-state estimation error covariance if the constant observation policy $\bar{U}$ is used.

$$0 = AZ + ZA' + DD' - Z\bar{U}Z$$

We may show, by some manipulations of the Riccati equation, that

$$\bar{X}^{-1} \succeq Z$$

DQ-Type Impossibility Result: Sketch of the proof

$$
\begin{array}{ccccc}
\bar{\Sigma} & \succeq & \bar{X}^{-1} & \succeq & Z \\
\operatorname{tr}\left(Q\,\bar{\Sigma}\right) & \geq & \operatorname{tr}\left(Q\,\bar{X}^{-1}\right) & \geq & \operatorname{tr}\left(Q\,Z\right) \\
+\operatorname{tr}\left(K\,\bar{U}\right) & & & & +\operatorname{tr}\left(K\,\bar{U}\right) \\
\hline
\eta\big(U(t)\big) & & & & \eta(\bar{U}) \\
\uparrow & & & & \uparrow \\
\text{cost of the time-varying policy} & & \geq & & \text{cost of the constant policy}
\end{array}
$$

where

$$
\bar{U} = \lim_{T \to \infty} \; \frac{1}{T} \int_0^T U(t)\, dt
$$

So optimal, infinite-horizon sensing policies for DQ-type sensor networks **never involve rhythmic variations** in the amount of attention paid to the different sensors.

What about data-aggregating (DA-type) sensor networks?

Optimal Rhythmic Policies for DA-Type Systems?

In the paper, we examine the second variation of the cost with respect to time-varying changes in the sensor network link allocation to determine optimality for DA-type systems.

A second-order necessary condition for optimality of a constant policy is given.

We now consider an example for which this condition is **not satisfied** and for which a periodic variation in link allocation improves performance beyond the best constant allocation.

Example: DQ-Type System

The control selects the fraction of the time spent monitoring each of the two states

$$d\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} dt + \begin{pmatrix} dw_1 \\ dw_2 \end{pmatrix}, \qquad 0 > \alpha_1 > \alpha_2$$

$$dy = \big(u(t)x_1 + (1-u(t))x_2\big)\, dt + d\nu$$

$\boxed{u=0}$

$$dy = x_2\, dt + d\nu$$

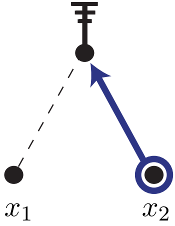

$\boxed{u=1}$

$$dy = x_1\, dt + d\nu$$

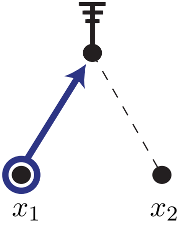

Example: DA-Type System

The data-sink node acts as an integrator, and can be connected to either system state. The output is the integrator state plus noise

$$d\begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 & u(t) & 1-u(t) \\ 0 & \alpha_1 & 0 \\ 0 & 0 & \alpha_2 \end{pmatrix}\begin{pmatrix} x_0 \\ x_1 \\ x_2 \end{pmatrix} dt + \begin{pmatrix} 0 \\ dw_1 \\ dw_2 \end{pmatrix}, \quad 0 > \alpha_1 > \alpha_2$$

$$dy = x_0\,dt + d\nu$$

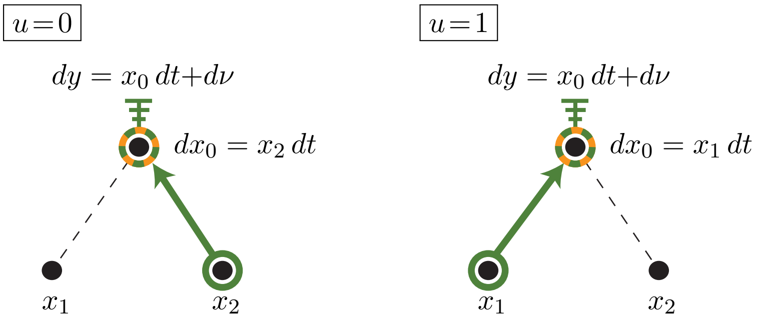

Example: Sensor control objective

Error covariance of the best linear estimator of the system state:

$$\Sigma = \mathcal{E}(x - \hat{x})(x - \hat{x})'$$

The objective is to minimize the average uncertainty of the two system states:

$$\eta = \frac{1}{T} \int_0^T \left(\Sigma_{11} + \Sigma_{22}\right) dt$$

Example: Parameter values

Eigenvalues of the dynamics:

$$\alpha_1 = -1/5, \quad \alpha_2 = -1/4$$

Noise covariance:

$$\mathcal{E}(dw\,dw') = \begin{pmatrix} 10 & 5 \\ 5 & 10 \end{pmatrix}$$

The best constant controls are:

DQ: $\bar{u} = 0.51$ **DA:** $\bar{u} = 0.57.$

As expected, the less stable state x_1 is given more attention than x_2.

Example: But are the constant policies optimal?

Periodic variations about the best constant policies:

$$u(t) = \bar{u} + 0.1\cos(\omega t)$$

Direct Query-type: Variations about the best constant policy give **worse** performance for all frequencies.

Data Aggregating-type: For a range of frequencies, periodic variations increase data throughput and reduce the average cost.

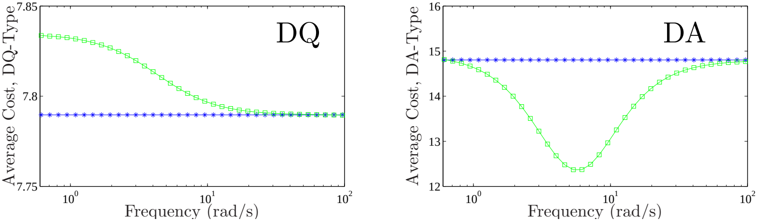

Conclusions

- For the DQ-type system, the optimizing constant sensing policy has the best time-average performance for the infinite-horizon problem.
- Introduced the data-aggregating (DA) type sensor network control problem, where there is in-network processing capability and reconfigurable connections.
- Gave an example DA-type system for which rhythmic oscillations improve performance beyond the best constant policy.
- In the paper, an additional condition for a constant policy to be optimal for a DA-type system is given.

Thanks to...

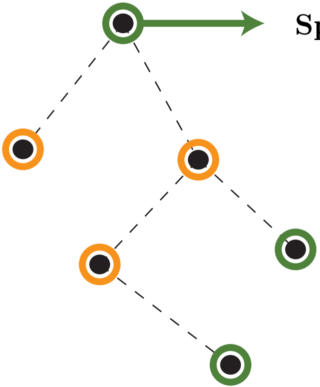

Sponsor: Army Research Office

Roger W. Brockett,
who suggested this line of investigation

Mohammed Ali Belabbas,
for many useful discussions

My contact information:

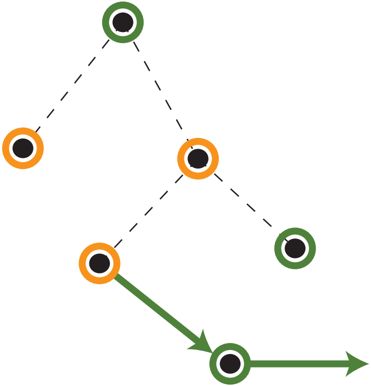

Alan C. O'Connor
aoconnor@eecs.harvard.edu