RC24980 (W1004-069) April 19, 2010
Electrical Engineering

IBM Research Report

Controlling Autonomous Data Ferries: Proof of Selected Theorems

Ting He, Kang-Won Lee
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Ananthram Swami
U.S. Army Research Laboratory
Adelphi, MD 20783

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Controlling Autonomous Data Ferries: Proof of Selected Theorems

Ting He, Kang-Won Lee, and Ananthram Swami

I. Introduction

This report contains detailed proofs for the theorems and corollaries in [1]. We first list the theorems in Section II and then give the proofs in Section III. See the original paper for terms and definitions.

II. Theorems

Proposition 2.1: The optimal value function $V_T(\mathbf{b})$ can be written as

$$V_T(\mathbf{b}) = \max\left(\max_{\boldsymbol{\alpha}\in\Gamma_T^f} \mathbf{b}\cdot\boldsymbol{\alpha},\ \max_{\boldsymbol{\alpha}\in\Gamma_T^s} \mathbf{b}_0\cdot\boldsymbol{\alpha}\right). \quad (1)$$

Moreover, Γ_T^f and Γ_T^s ($T \geq 1$) satisfy the recursion[1]:

$$\Gamma_T^f = \mathbf{P}\tilde{\Gamma}_T, \qquad \Gamma_T^s = \arg\max_{\boldsymbol{\alpha}\in\tilde{\Gamma}_T} \mathbf{b}_0\cdot\boldsymbol{\alpha}, \quad (2)$$

where $\mathbf{P}\tilde{\Gamma}_T \triangleq \{\mathbf{P}\boldsymbol{\alpha} : \forall\boldsymbol{\alpha}\in\tilde{\Gamma}_T\}$, and $\tilde{\Gamma}_T \triangleq \{\tilde{\boldsymbol{\alpha}}^{u,\boldsymbol{\alpha}'},\ \tilde{\boldsymbol{\alpha}}^u : \forall u\in\mathcal{S},\ \boldsymbol{\alpha}'\in\Gamma_{T-1}^f\}$ for

$$\tilde{\alpha}^{u,\boldsymbol{\alpha}'}(s) = \begin{cases} \gamma(1+c_1) & \text{if } s = u, \\ \gamma\alpha'(s) & \text{o.w.,} \end{cases} \quad (3)$$

$$\tilde{\alpha}^{u}(s) = \begin{cases} \gamma(1+c_1) & \text{if } s = u, \\ \gamma c_2 & \text{o.w.,} \end{cases} \quad (4)$$

$c_1 \triangleq \max_{\boldsymbol{\alpha}'\in\Gamma_{T-1}^f\cup\Gamma_{T-1}^s} \mathbf{b}_0\cdot\boldsymbol{\alpha}'$, and $c_2 \triangleq \max_{\boldsymbol{\alpha}'\in\Gamma_{T-1}^s} \mathbf{b}_0\cdot\boldsymbol{\alpha}'$.

Theorem 2.2: At horizon T and belief state $\mathbf{b}$, it is optimal to follow the myopic policy if

$$1-\gamma \geq \frac{1-b^{(1)}}{(1-b^{(2)})\left[1+b_0^*\gamma(1-\gamma^{T-1})/(1-\gamma)\right]}, \quad (5)$$

where $b^{(1)} = \max_{u\in U} b'(u)$ (achieved at u_1) is the largest immediate reward, and $b^{(2)} = \max_{u\in U\setminus u_1} b'(u)$ the second largest immediate reward.

T. He and K-W Lee are with IBM T.J. Watson Research Center, Hawthorne, NY 10532, USA. E-mail: {the,kangwon}@us.ibm.com.

A. Swami is with the Army Research Laboratory, Adelphi, MD 20783, USA. E-mail: aswami@arl.army.mil.

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

[1]For $T = 0$, $\Gamma_0^f = \Gamma_0^s = \{(0,\ldots,0)\}$.

Proposition 2.3: Each belief-based dynamic policy π corresponds to a unique trajectory $\mathbf{u}^\pi \triangleq (u^\pi(t))_{t=1}^d$ (d can be ∞) within one domain such that π is equivalent to a policy of following $\mathbf{u}^\pi$ and repeating it in different domains until contact, after which this process is repeated in a new domain.

Claim 2.4: The optimal predetermined policy is to keep switching among the most likely cells of different domains (called the *switching policy*).

Theorem 2.5: The average inter-contact time $\mathbb{E}[K^{\text{DY}}]$ of the optimal dynamic policy satisfies

$$K_0' + \frac{d'(1-p_0')}{p_0'} \leq \mathbb{E}[K^{\text{DY}}] \leq K_0 + \frac{d(1-p_0)}{p_0}, \quad (6)$$

where $p_0 \triangleq 1 - \prod_{t=1}^{d}(1-p_t)$, $K_0 \triangleq \frac{1}{p_0}\sum_{t=1}^{d} t p_t \prod_{i=1}^{t-1}(1-p_i)$, and p_0', K_0' are similarly defined with p_t, d replaced by p_t', d'.

Corollary 2.6: If $\mathbf{b}_0$ is uniform, then $\mathbb{E}[K^{\text{SW}}] = n$, and

$$\frac{n}{2} \leq \mathbb{E}[K^{\text{DY}}] \leq n - \frac{d(n)}{2}, \quad (7)$$

where $d(n)$ is the length of the longest dominating trajectory.

Corollary 2.7: If the node mobility is sufficiently biased such that $b_0^{(1)} > b_0^{(2)}/(1-b_0^{(1)})$, then $\mathbb{E}[K^{\text{DY}}] = \mathbb{E}[K^{\text{SW}}]$, *i.e.,* the switching policy is optimal.

Lemma 2.8: Let $\mathbf{P}$ be a base transition matrix and $\mathbf{P}_\beta$ its scaled version according to an activeness parameter β. Then their associated steady-state distribution $\mathbf{b}_0$ and $\mathbf{b}_0'$ are equal for all $\beta \in (0,\ 1/(1-\min_i P_{ii}))$.

Theorem 2.9: If we scale the transition matrix in Theorem 2.5 by an activeness parameter β, then besides the bounds in (6), $\mathbb{E}[K^{\text{DY}}]$ also satisfies $\mathbb{E}[K^{\text{DY}}] \leq \tilde{K}_0 + d'(1-\tilde{p}_0)/\tilde{p}_0$, where

$$\tilde{p}_0 \triangleq \sum_{t=1}^{d'} b_0^{(t)} - \frac{\beta\delta}{2}\sum_{t=1}^{d'-1} b_0^{(t)}(d'-t)(d'-t+1), \quad (8)$$

$$\tilde{K}_0 \triangleq \frac{1}{\tilde{p}_0}\sum_{t=1}^{d'} t\left(b_0^{(t)} - \beta\delta\sum_{i=1}^{t-1} b_0^{(i)}(t-i)\right), \quad (9)$$

for $\delta \triangleq \max_{i\neq j} P_{ij}$ and d' defined as in Theorem 2.5. In particular, as $\beta \to 0$, $\mathbb{E}[K^{\text{DY}}]$ converges to the lower bound at $O(\beta)$.

Corollary 2.10: Under $L \geq 1$ nodes per domain, Corollary 2.7 and 2.6 can be extended as follows: for biased mobility, the switching policy is still optimal if $b_0^{(1)} > b_0^{(2)}/(1-b_0^{(1)})$, and its inter-contact time will converge to 1 as L increases at $O((1-b_0^{(1)})^L)$; for symmetric mobility and large domains (*i.e.,* $n \gg 1$), $\mathbb{E}[K^{\text{SW}}] \approx n/L$, and

$$\frac{n}{L+1} \le \mathbb{E}[K^{\text{DY}}] \le \frac{n[1-(1-d(n)/n)^{L+1}]}{(L+1)[1-(1-d(n)/n)^L]}, \quad (10)$$

where $d(n)$ is defined as in Corollary 2.6.

III. Proofs

A. Proof of Proposition 2.1

We prove by induction. The result clearly holds for $T = 0$ and $\Gamma_0^f = \Gamma_0^s = \{(0,\ldots,0)\}$. At horizon T, we plug the induction result for horizon $T-1$ into the value iteration:

$$V_T(\mathbf{b}) = \gamma \max_u \Big[b'(u) + b'(u) \max_{\boldsymbol{\alpha}' \in \Gamma^f_{T-1} \cup \Gamma^s_{T-1}} \mathbf{b}_0 \cdot \boldsymbol{\alpha}' + (1-b'(u)) \max\big(\max_{\boldsymbol{\alpha}' \in \Gamma^f_{T-1}} \mathbf{b}'_{\setminus u} \cdot \boldsymbol{\alpha}', \max_{\boldsymbol{\alpha}' \in \Gamma^s_{T-1}} \mathbf{b}_0 \cdot \boldsymbol{\alpha}' \big) \Big] \quad (11)$$

Given Γ^f_{T-1} and Γ^s_{T-1}, $\max_{\boldsymbol{\alpha}' \in \Gamma^f_{T-1} \cup \Gamma^s_{T-1}} \mathbf{b}_0 \cdot \boldsymbol{\alpha}'$ and $\max_{\boldsymbol{\alpha}' \in \Gamma^s_{T-1}} \mathbf{b}_0 \cdot \boldsymbol{\alpha}'$ are constants, denoting them by c_1, c_2.

For $u \in \mathcal{U}_f$, if $\max_{\boldsymbol{\alpha}' \in \Gamma^f_{T-1}} \mathbf{b}'_{\setminus u} \cdot \boldsymbol{\alpha}' \ge c_2$, then

$$\begin{aligned} V_T(\mathbf{b}) &= \gamma \left[b'(u) + b'(u)c_1 + (1-b'(u))\mathbf{b}'_{\setminus u} \cdot \tilde{\boldsymbol{\alpha}}' \right] \\ &= \mathbf{b}' \cdot \tilde{\boldsymbol{\alpha}}^{u,\tilde{\boldsymbol{\alpha}}'} = \mathbf{b} \cdot (\mathbf{P}\tilde{\boldsymbol{\alpha}}^{u,\tilde{\boldsymbol{\alpha}}'}), \end{aligned}$$

where $\tilde{\boldsymbol{\alpha}}^{u,\tilde{\boldsymbol{\alpha}}'}$ is defined as in (3) for $\tilde{\boldsymbol{\alpha}}' = \arg\max_{\boldsymbol{\alpha}' \in \Gamma^f_{T-1}} \mathbf{b}'_{\setminus u} \cdot \boldsymbol{\alpha}'$; otherwise,

$$\begin{aligned} V_T(\mathbf{b}) &= \gamma \left[b'(u) + b'(u)c_1 + (1-b'(u))c_2 \right] \\ &= \mathbf{b}' \cdot \tilde{\boldsymbol{\alpha}}^u = \mathbf{b} \cdot (\mathbf{P}\tilde{\boldsymbol{\alpha}}^u), \end{aligned}$$

where $\tilde{\boldsymbol{\alpha}}^u$ is defined as in (4).

For $u \in \mathcal{U}_s$, since $\mathbf{b}'$ is always reset to $\mathbf{b}_0$, $V_T(\mathbf{b})$ no longer depends on $\mathbf{b}$, and similar arguments as above will show that $V_T(\mathbf{b}) \equiv \mathbf{b}_0 \cdot \boldsymbol{\alpha}^*_T$, where $\boldsymbol{\alpha}^*_T = \arg\max \mathbf{b}_0 \cdot \boldsymbol{\alpha}$ over all $\tilde{\boldsymbol{\alpha}}^{u,\boldsymbol{\alpha}'}$, $\tilde{\boldsymbol{\alpha}}^u$. Combining the above cases proves the induction, where Γ^f_T contains all $\mathbf{P}\tilde{\boldsymbol{\alpha}}^{u,\boldsymbol{\alpha}'}$ and $\mathbf{P}\tilde{\boldsymbol{\alpha}}^u$, and Γ^s_T only contains $\boldsymbol{\alpha}^*_T$. ■

B. Proof of Theorem 2.2

The myopic action $\pi_1(\mathbf{b}) = u_1$ is optimal if and only if its reward is greater than all alternative actions, *i.e.,*

$$(b'(u_1) - b'(u))(1 + V_{T-1}(\mathbf{b}_0)) \ge (1-b'(u))V_{T-1}(\mathbf{b}'_{\setminus u}) - (1-b'(u_1))V_{T-1}(\mathbf{b}'_{\setminus u_1}). \quad (12)$$

Since it is preferable to earn rewards earlier due to discount, the value function is upper bounded by the case when there is an immediate contact: $V_{T-1}(\mathbf{b}'_{\setminus u}) \le \gamma(1 + V_{T-1}(\mathbf{b}_0))$. Moreover, the value function is lower bounded by its value at $\mathbf{b}_0$ since we have the option of switching domain: $V_{T-1}(\mathbf{b}'_{\setminus u_1}) \ge V_{T-1}(\mathbf{b}_0)$. Applying these bounds to the right-hand side of (12) gives a sufficient condition

$$1-\gamma \ge \frac{1-b'(u_1)}{(1-b'(u))(1+V_{T-1}(\mathbf{b}_0))} \quad (13)$$

for all $u \neq u_1$, which can be reduced to a single inequality with $b'(u) = b^{(2)}$ (note $b'(u_1) = b^{(1)}$). Finally, the value function is lower bounded by that of deterministically switching among the most likely cells of different domains, yielding $V_T(\mathbf{b}) \ge b_0^* \gamma(1-\gamma^T)/(1-\gamma)$. Replacing $V_{T-1}(\mathbf{b}_0)$ in (13) by its lower bound proves the theorem. ■

C. Proof of Proposition 2.3

It suffices to show that the original policy π and the policy of following $\mathbf{u}^\pi$ give the same action in every step. We construct $\mathbf{u}^\pi$ as follows. Define $\mathbf{b}_1 = \mathbf{b}_0$ and $\mathbf{b}_{t+1} = [\mathbf{P}^T \mathbf{b}_t]_{\setminus u^\pi(t)}$ for $t \ge 1$; define $u^\pi(t) = \pi(\mathbf{b}_t)$ and d to be the last step before switching, *i.e.,* $\pi(\mathbf{b}_{d+1}) \in \mathcal{U}_s$. It can be verified that each $\mathbf{b}_t$ is indeed the belief after following π and having $t-1$ consecutive misses, and by the above construction the two policies will give the same action. ■

D. Proof of Claim 2.4

Suppose that the optimal predetermined policy is to let the ferry stay with one node for τ slots, at states $u_1,\ldots,u_\tau$ (possibly with repetitions), for some $\tau \ge 1$. We claim that the expected reward r_t of u_t satisfies $r_t \le b_0(u_t) \le \max_u b_0(u)$, and thus the switching policy is optimal.

Since $r_t = \Pr\{\text{first contact occurs in step } t\}$ which can be written as $(1-\sum_{j=1}^{t-1} r_j)\Pr\{s_t = u_t | (s_j \neq u_j)_{j=1}^{t-1}\}$, it suffices to show

$$\Pr\{s_t = u | (s_j \neq u_j)_{j=1}^{t-1}\} \le \frac{b_0(u)}{1-\sum_{j=1}^{t-1} r_j}$$

for any state u. We prove the claim by induction. For $t = 1$, the result holds trivially. For $t > 1$,

$$\Pr\{s_t = u | (s_j \neq u_j)_{j=1}^{t-1}\} \le \frac{\Pr\{s_t = u | (s_j \neq u_j)_{j=1}^{t-2}\}}{1 - \Pr\{s_{t-1} = u_{t-1} | (s_j \neq u_j)_{j=1}^{t-2}\}} \quad (14)$$

because $P(A|B) \le P(A)/P(B)$. By induction and the invariance of $\mathbf{b}_0$ under state transition, the numerator is bounded by $b_0(u)/(1-\sum_{j=1}^{t-2} r_j)$, and the denominator is equal to $1 - r_{t-1}/(1-\sum_{j=1}^{t-2} r_j)$ by definition. Plugging these into (14) yields the result. ■

E. Proof of Theorem 2.5

For the upper bound, consider the inter-contact time K^{MY} of the myopic policy. If we call the process of completing the policy trajectory once a *round*, then the mean inter-contact time can be decomposed into

$$\mathbb{E}[K^{\text{MY}}] = \sum_{k=1}^{\infty} \Pr\{k \text{ rounds}\} \cdot \mathbb{E}[K^{\text{MY}} | k \text{ rounds}]. \quad (15)$$

Within one round, its distribution satisfies

$$\Pr\{K^{\text{MY}} = t\} = p_t \prod_{j=1}^{t-1}(1-p_j), \quad 1 \le t \le d,$$

where $p_t \triangleq \Pr\{K^{\text{MY}} = t | K^{\text{MY}} > t-1\}$. Let p_0 denote the probability of having a contact within the first round, given by

$$p_0 \triangleq \Pr\{K^{\text{MY}} \le d\} = 1 - \prod_{t=1}^{d}(1-p_t),$$

and K_0 be the expected inter-contact time if having a contact within the first round, given by

$$K_0 \triangleq \mathbb{E}[K^{\text{MY}} | K^{\text{MY}} \le d] = \frac{1}{p_0}\sum_{t=1}^{d} t \Pr\{K^{\text{MY}} = t\}.$$

Then we have $\Pr\{k \text{ rounds}\} = (1-p_0)^{k-1}p_0$, and $\mathbb{E}[K^{\text{MY}}|k \text{ rounds}] = (k-1)d + K_0$ since different rounds are *i.i.d.* Plugging the results into (15) yields the upper bound $\mathbb{E}[K^{\text{MY}}] = K_0 + d(1-p_0)/p_0$.

For the lower bound, we first show that the inter-contact time is lowered bounded when nodes are static. For any given policy with trajectory $(u^*(t))_{t=1}^{d^*}$, the mean inter-contact time $\mathbb{E}[K]$ can be written as

$$\mathbb{E}[K] = \sum_{t=1}^{\infty} t p_t^* \prod_{j=1}^{t-1}(1-p_j^*), \tag{16}$$

where $p_t^* \triangleq \Pr\{K = t | K > t-1\}$ as before, with the property that $p_t^* = p_{t-d^*}^*$ for $t > d^*$. We observe that $\mathbb{E}[K]$ is a decreasing function with each p_t^* by examining the partial derivative

$$\begin{aligned}\frac{\partial \mathbb{E}[K]}{\partial p_t^*} &= t\prod_{j=1}^{t-1}(1-p_j^*) - \sum_{s=t+1}^{\infty} s p_s^* \left[\prod_{j=1:s-1,\, j\ne t}(1-p_j^*)\right] \\ &= \left[\prod_{j=1}^{t-1}(1-p_j^*)\right](t - \mathbb{E}[K|K>t]) < 0.\end{aligned}$$

Moreover, within each round (*i.e.,* $1 \le t \le d^*$), $p_t^* \le b_0(u^*(t))/(1-\sum_{j=1}^{t-1} b_0(u^*(j)))$, achievable only if the node is static and the trajectory does not have loops. This is from the property that $b_t(i) \le b_0(i)/(1-\sum_{j=1}^{t-1} b_0(u^*(j)))$, where $\mathbf{b}_t = \mathbf{P}^T(\mathbf{b}_{t-1})_{\setminus u^*(t-1)}$ ($\mathbf{b}_1 \triangleq \mathbf{b}_0$). For $t=1$, it holds trivially. For $t > 1$,

$$\begin{aligned} b_t(i) &= \sum_{j\ne u^*(t-1)} \frac{b_{t-1}(j)}{1-b_{t-1}(u^*(t-1))}\cdot P_{ji} \\ &\le \frac{1}{1-\frac{b_0(u^*(t-1))}{1-\sum_{l=1}^{t-2} b_0(u^*(l))}} \cdot \frac{\sum_{j\ne u^*(t-1)} b_0(j)P_{ji}}{1-\sum_{l=1}^{t-2} b_0(u^*(l))} \\ &\le \frac{b_0(i)}{1-\sum_{j=1}^{t-1} b_0(u^*(j))}.\end{aligned}$$

We finish by noting that $p_t^* = b_t(u^*(t))$.

Next, we find the minimum inter-contact time for static nodes. Since for given values $x_1, x_2, \ldots,$ the permutation ϕ that minimizes $\sum_t t x_{\phi(t)} \prod_{j=1}^{t-1}(1-x_{\phi(j)})$ is the descending order $x_{\phi(1)} \ge x_{\phi(2)} \ge \ldots$, the best policy for static nodes is to tour each domain in descending order of $b_0(u)$ and switch domain if the conditional contact probability $< b_0^{(1)}$. The corresponding mean inter-contact time can be computed as in the upper bound. ∎

F. Proof of Corollary 2.6

The lower bound follows directly by plugging $b_0(u) \equiv 1/n$ into the lower bound in (6). For the upper bound, consider a policy with the longest dominating trajectory as its policy trajectory. By definition of the dominating trajectory, the conditional contact probability is $p_t = 1/(n-t+1)$ for $t = 1, \ldots, d(n)$. Following similar arguments as in the proof of Theorem 2.5, we can compute the inter-contact time of this policy by $K_0 + d(n)(1-p_0)/p_0$, where

$$\begin{aligned} p_0 &= 1 - \prod_{t=1}^{d(n)}\left(1-\frac{1}{n-t+1}\right) = \frac{d(n)}{n}(1+o(1)), \\ K_0 &= \frac{1}{p_0}\sum_{t=1}^{d(n)} \frac{t}{n-t+1}\left[\prod_{j=1}^{t-1}(1-\frac{1}{n-j+1})\right] \\ &= \frac{d(n)}{2}(1+o(1)).\end{aligned}$$

Plugging in the approximations yields the bound. ∎

G. Proof of Corollary 2.7

Since $\mathbb{E}[K^{\text{DY}}] \le \mathbb{E}[K^{\text{SW}}]$ by definition, it suffices to show that the lower bound in Theorem 2.5 is equal to $\mathbb{E}[K^{\text{SW}}] = 1/b_0^{(1)}$. Specifically, $b_0^{(1)} > b_0^{(2)}/(1-b_0^{(1)})$ implies that $d' = 1$, which leads to a lower bound of $1/b_0^{(1)}$. ∎

H. Proof of Lemma 2.8

By definition of steady-state distribution, we have that for any $j \in \{1, \ldots, n\}$,

$$b_0'(j) = \sum_{i=1}^{n} b_0'(i)(P_\beta)_{ij} = b_0'(j)(1-\beta\sum_{i\ne j}P_{ji}) + \beta\sum_{i\ne j} b_0'(i)P_{ij},$$

which gives $b_0'(j)\sum_{i\ne j}P_{ji} = \sum_{i\ne j} b_0'(i)P_{ij}$. On the other hand, $\mathbf{b}_0$ also satisfies

$$b_0(j) = \sum_{i=1}^{n} b_0(i)P_{ij} = b_0(j)(1-\sum_{i\ne j}P_{ji}) + \sum_{i\ne j} b_0(i)P_{ij}.$$

Combining the above (and the uniqueness of steady-state distribution) yields $\mathbf{b}_0' = \mathbf{b}_0$. ∎

I. Proof of Theorem 2.9

Let $(v(t))_{t=1}^{d'}$ be the optimal policy trajectory for static nodes, *i.e.,* $v(t)$ is the cell with the tth largest steady-state probability $b_0^{(t)}$ and d' is defined as in (6). The upper bound is derived from the performance of this policy under the scaled transition matrix $\mathbf{P}_\beta$. We claim that the conditional contact probability satisfies

$$p_t \geq \frac{b_0^{(t)} - \beta\delta \sum_{i=1}^{t-1} b_0^{(i)}(t-i)}{\prod_{i=1}^{t-1}(1-p_i)}. \quad (17)$$

It holds trivially for $t=1$. For $t>1$, we can write p_t as

$$\begin{aligned} p_t &= \Pr\{\text{node in } v(t) \text{ in slot } t | \text{not in } v(i) \text{ in slot } i<t\} \\ &= \frac{1}{\prod_{i=1}^{t-1}(1-p_i)} \Bigg(\Pr\{v(t) \text{ in slot } t\} \\ &\quad - \Pr\Big\{ \bigcup_{i=1}^{t-1} \{v(i) \text{ in slot } i,\ v(t) \text{ in slot } t\}\Big\}\Bigg). \quad (18) \end{aligned}$$

Now note that $\Pr\{v(t) \text{ in slot } t\} = b_0^{(t)}$, and $\Pr\{v(i)$ in slot i, $v(t)$ in slot $t\} = b_0^{(i)}(P_\beta^{t-i})_{v(i),\, v(t)} \leq b_0^{(i)}(t-i)\beta\delta$. The last inequality is from $(P_\beta^k)_{ij} \leq k\beta\delta$ for $i \neq j$, where $\mathbf{P}_\beta^k$ is the k-step transition matrix, shown by induction:

$$\begin{aligned} (P_\beta^k)_{ij} &= \sum_{l\neq j} (P_\beta)_{il} \cdot (P_\beta^{k-1})_{lj} + (P_\beta)_{ij} \cdot (P_\beta^{k-1})_{jj} \\ &\leq (k-1)\beta\delta + \beta\delta. \end{aligned}$$

Applying these results and union bound into (18) yields (17).

Now that $\mathbb{E}[K]$ is a decreasing function of p_t (see the proof of Theorem 2.5), plugging the lower bound (17) into the formula in Theorem 2.5 gives an upper bound on $\mathbb{E}[K]$ of policy $(v(t))_{t=1}^{d'}$, which in turn upper bounds $\mathbb{E}[K^{\text{DY}}]$.

Moreover, the gap between the new upper bound and the lower bound is

$$\frac{\beta\delta}{2\tilde{p}_0 p_0'} \Bigg[\Bigg(\sum_{t=1}^{d'-1} b_0^{(t)}(d'-t)(d'-t+1) \Bigg) \Big(d' + \sum_{t=1}^{d'} t b_0^{(t)}\Big) - 2p_0' \sum_{t=1}^{d'} t \Bigg(\sum_{i=1}^{t-1} b_0^{(i)}(t-i) \Bigg) \Bigg].$$

Since $\tilde{p}_0 \to p_0'$ as $\beta \to 0$ and the other quantities are independent of β, we conclude that the gap decays at $O(\beta)$ as $\beta \to 0$. ■

J. Proof of Corollary 2.10

For biased mobility, the optimality of the switching policy follows from the same argument as in Corollary 2.7. For $L \gg 1$, its mean inter-contact time satisfies

$$\mathbb{E}[K^{\text{SW}}] = \frac{1}{1-(1-b_0^{(1)})^L} = 1 + (1-b_0^{(1)})^L + o\left((1-b_0^{(1)})^L\right).$$

For symmetric mobility, we extend Corollary 2.6 as follows. For the switching policy,

$$\mathbb{E}[K^{\text{SW}}] = \frac{1}{1-(1-\frac{1}{n})^L} = \frac{1}{1-e^{-L/n+o(1/n)}} = \frac{1}{\frac{L}{n}+o(\frac{1}{n})}.$$

For dynamic policy, the proof is similar to that of Corollary 2.6.

Specifically, for the lower bound, assume the policy is to sweep the field in the order of cells $1, \ldots, n$ (all orders are equivalent). The probability of having the first contact in slot t can be written as $\Pr\{\exists$ node in cell t, $\nexists$ node in cells $<t\}$, which is equal to $\left[1-(1-1/(n-t+1))^L\right]((n-t+1)/n)^L$. Therefore, the mean inter-contact time for static nodes is

$$\sum_{t=1}^{n} t \left[1-(1-\frac{1}{n-t+1})^L\right] \left(\frac{n-t+1}{n}\right)^L = n^{-L}\sum_{t=1}^{n} t^L. \quad (19)$$

By the inequality $\int_0^n x^L dx \leq \sum_{t=1}^n t^L \leq \int_0^n (x+1)^L dx$, we have that (19) is bounded between $n/(L+1)$ and $[(n+1)(1+1/n)^L - n^{-L}]/(L+1)$, where the upper bound approximates $n/(L+1)$ for large n.

For the upper bound, still consider the policy of following the longest dominating trajectory. Its conditional contact probability is $p_t = 1-[(n-t)/(n-t+1)]^L$. Using the formula in (6), we have

$$\begin{aligned} p_0 &= 1 - \left(1-\frac{d(n)}{n}\right)^L, \\ K_0 &= \frac{1}{p_0}\left[n^{-L} \sum_{t=n-d(n)+1}^{n} t^L - d(n)\left(\frac{n-d(n)}{n}\right)^L\right]. \end{aligned}$$

Applying the previous bounds on $\sum_{t=1}^n t^L$, we can show that for large n, $K_0 p_0$ can be approximated by $n/(L+1) - ((n-d(n))/n)^L (n+Ld(n))/(L+1)$. Plugging these into the formula gives the upper bound. ■

References

[1] T. He, K.-W. Lee, and A. Swami, "Flying in the Dark: Controlling Autonomous Data Ferries with Partial Observations," 2010. draft.