Behavior of Transient, Exhaustive Polling Systems with n Nodes

Alec Zimmer and Joe Kraisler

August 8, 2012

Abstract

We examine a deterministic exhaustive polling system with n nodes, $n \geq 3$, under a switching rule of generalized greedy type. We consider when the sequence in which the server nodes is periodic. Also examined are the long term behavior of queue lengths relative to each other and when different initial line lengths affect the long-term behavior of the system.

1 Introduction

Consider a deterministic exhaustive polling system with n nodes, $n \geq 3$, with jobs continuously queueing at the nodes. Since the system is exhaustive, the server will process all jobs at one node, and then will switch to another node based on the length of the other queues. The rate at which jobs queue and the server processes jobs will depend on the specific node. We will assume that the rate at which the server processes jobs is large enough such that the server does not get stuck at a single node. We also assume that the switching time is zero; the server will switch between nodes instantaneously.

We follow the approach of [1], where it was shown that for the case $n = 3$ almost all choices of threshold parameters the sequence in which the server visits the nodes is eventually periodic. We will prove many of their results for this deterministic system for $n \geq 3$. However, we had difficulty showing this result for $n > 3$, which is discussed briefly at the end of this paper.

To be more precise about the system, queues will continuously grow at each of the n nodes, at a rate of λ_i at the ith node. One node will be processed at a time by a server until its queue is empy, where the ith queue is processed at a rate μ_i. We will require that $\rho_i := \lambda_i/\mu_i < 1$, as otherwise the system will get stuck processing a single node. We will also require that $\sum_{i=1}^{n} \rho_i > 1$ so that some queue will always have positive length after the process starts.

We can visualize the system as a point travelling in $\mathbb{R}^n_+$, including the boundary, where the ith coordinate is the length of the queue at the ith node. We will consider the system projected onto the n-simplex where the sum of the coordinates is 1, ignoring a start with all the queues having length zero. We can assume the process begins with exactly one queue of length zero, since otherwise we could wait for the server to process the queue it is at and end up with such a system.

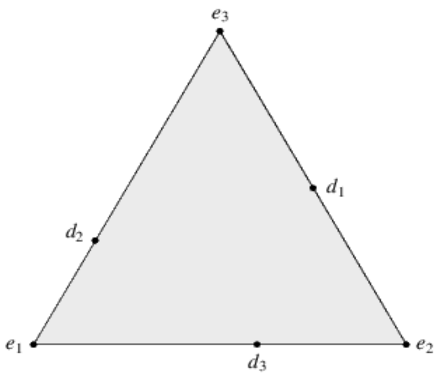

Figure 1: Simplex for n=3

When the ith queue is empty, $x_i = 0$. We will have $b_i > 0$ for $i = 1, 2, ..., n$, where the server will start processing the jth node, where $b_j x_j \geq b_k x_k$ for $j \neq k$. What exactly happens when $b_j x_j = b_k x_k$, a switching boundary, will not matter for our purposes. We could instead use $b_{k,i} > 0$, where the weights to put on the other nodes will depend on which of the k nodes was just processed. We would then lose continuity on the edges of the n-simplex, which would not affect many of our arguments below, but would make further developments more difficult. Let a decision point d_i be the point on A_i such that $b_j x_j = b_k x_k$ for $j, k \neq i$. An example is shown in Figure 1 for $n = 3$.

We will denote the map that sends a point z on the boundary of the n-simplex to the next point on the boundary of the simplex under this process by ϕ. A collection Z of the above parameters on n nodes will be called an n-process. A point z on the n-simplex and its images under ϕ will be called a trajectory. For a trajectory, we will say $z(0) = z$, the initial point, $z(i+1) = \phi(z(i))$ for $i \in \mathbb{N}$.

We say a trajectory is is eventually m-periodic if there is some $T > 0$ such that if $t \geq T$, $z(t+m)$ and $z(t)$ are on the same face of the simplex. A trajectory is a periodic orbit with period m if for any $s, t \in \mathbb{N}$, $z(s + tm) = z(s)$.

2 Basic Properties

When the server is processing the jth node, the queue at the jth node changes at the rate $\lambda_j - \mu_j < 0$, while the ith queue grows at the rate λ_i, $i \neq j$. This corresponds to the particle travelling along a line in $\mathbb{R}^n_+$. If a point starts at $\bar{y}$ and travels toward the jth side, computation gives that the particle reaches $\partial\mathbb{R}^n_+$ at the point

$$y(t_1) = \sum_{i \neq j} \left(\bar{y}_i + \frac{\lambda_i \bar{y}_j}{\mu_j - \lambda_j} \right) e_i$$

where $t_1 = t_0 + \bar{y}_j/(1 - \rho_j)$.

We'd like to analyze the process py projecting down to the simplex where $\sum_i y_i = 1$. When we consider this projected system, section 2 of [1] derives that when the server is processing the jth node, the particle goes in a straight line

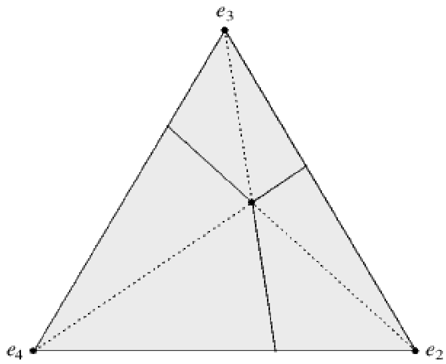

Figure 2: n=4

in the direction of the point

$$v_j = \frac{1}{\mu_j \theta_j} \left(\sum_{i \neq j} \lambda_i e_i + (\lambda_j - \mu_j) e_j \right) \tag{1}$$

if $\theta_j := \mu_j^{-1} \sum_i \lambda_i - 1 \neq 0$. It then stops once $x_i = 0$, at which point the server switches to a new node. If $\theta = 0$, the different possible particles going towards the jth side go parallel to each other.

Denote the n-simplex the system gets projected to by A, and let A_j be the face of A where the jth coordinate is zero. From now on we will only be considering the version of the n-process on A for our analysis, though it is isomorphic to the original system. Let B_j be the subset of A consisting of points that get sent to A_j under the forward map ϕ. Let f_j be the restriction of ϕ to B_j. While we won't make use of the explicit formula, it is given by [1] for f_j is:

$$f_j(z) = \sum_{i \neq j} \frac{(\mu_j - \lambda_j) z_i + \lambda_i z_j}{(\mu_j - \lambda_j) + \mu_j \theta_j z_j} e_i.$$

It is helpful for many of the following arguments to think of the $n = 3$ and $n = 4$ as special cases. For $n = 4$, on the ith side there will be a point where $b_j x_j = b_k x_k = b_l x_l$, and from this point we will have three line segments going out to an edge corresponding to $b_j x_j = b_k x_k$. These line segments, if continued into lines, will intersect one of the e_j. See Figure 2 for one face of the tetrahedron. They also move continuously into the corresponding lines on adjacent faces (as long as b_i does not depend on which face we are on). Similar statements could be made about $n > 4$.

3 The n-Process

We first simplify the choice of parameters.

Lemma 3.1. *Consider the n-process Z with parameters λ_i and μ_i for $i, j = 1, ..., n$, $i \neq j$. There is some n-process Z' isomorphic to Z with parameters $\mu'_i = 1$ and $\lambda'_i = \lambda_i / \mu_i = \rho_i$.*

Proof. Multiply the coordinates of each point in Z by

$$\begin{pmatrix} 1/\mu_1 & 0 & \dots & 0 \\ 0 & 1/\mu_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1/\mu_n \end{pmatrix}$$

and then project back to the unit simplex A. Scaling the ith coordinate by $1/\mu_i$ preserves lines in the n-process, meaning it preserves trajectories and creates an isomorphic n-process. Projecting to the unit simplex then gives us an isomorphic process Z'.

We now see that Z' has the desired parameters. For any i, a queue of length 1, which takes time μ_i to be processed in Z, will be sent to a queue of length $1/\mu_i$, and thus take time 1 to be processed. This means the new n-process has $\mu'_i = 1$. Similarly, λ_i gets sent to $\lambda'_i = \lambda_i/\mu_i$. □

This lemma means that from now on we will assume that $\mu_i = 1$ for all i. We then get $\theta := \theta_j = \sum_i \rho_i - 1 > 0$ does not depend on the specific j.

We get from Equation 1

$$v_k = \frac{1}{\theta}\left(\sum_{i \neq k} \rho_i e_i + (\rho_k - 1)e_k\right) \tag{2}$$

giving

$$v_i - v_j = \frac{1}{\theta}(e_j - e_i).$$

The v_i then form an n-simplex with edges of length $1/\theta$ parallel to those of A, with each v_i outside A. We will now use A^V, the intersection of the v_i-simplex with the boundary of A. This set will be central to our analysis. Let A^V_j be the jth side of A^V (where the jth coordinate is zero). We immediately have that A^V is closed under the map ϕ.

Lemma 3.2. *For any i, there is some $\gamma \in (0,1)$ such that for $z, z' \in A^V_j \cap B_i$, $j \neq i$, $|f_i(z) - f_i(z')| \leq \gamma |z - z'|$.*

Proof. Let x be the maximum distance from v_i to $f_i(B_i)$. Let d be the minimum distance from v_i to B_i. Then $x/d < 1$, and we'll show that $\gamma := x/d$ suffices.

For $i \neq j$, consider the line L connecting v_i and v_j. The angle, which we will call α, of the unit normal at $L \cap A^V_j$ to L is the same as the angle of the unit normal at $L \cap A^V_i$ to L since the edges of the n-simplex formed by the nodes are parallel to the edges of A. For any point $z \in A^V_j$, consider the unit normal u_z on the face into the interior of our simplex. The angle u_z makes with the line from z to v_i will be greater than α by the geometry of the simplex, and the angle the unit normal of $f_i(z)$ makes with the line from z to v_i will be less than α.

Thus, on A^V_j, the transformation a region undergoes by being projected at an angle will never expand the region. Thus f_i contracts by the factor γ. □

Figure 3 gives an example for $n = 3$.

Note that in the above we can find a slightly larger set on which there is a $\gamma < 1$ such that ϕ contracts by at least γ, a property we will use later. Let $C(\gamma)$ be some such set, containing points within some ϵ distance from A^V. Also, for

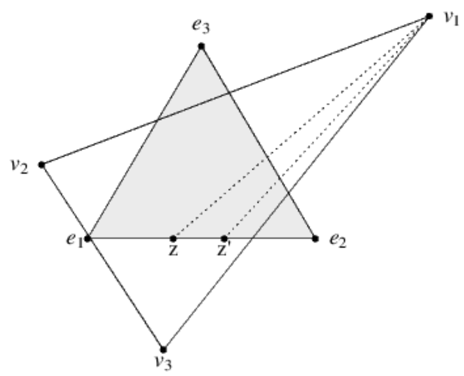

Figure 3: n=4

points on different sides of A, we can think of the distance between them as the minimum distance by travelling along the surface of A, which broadens the sense in which f_j is a contraction.

We could also take some $\gamma < 1$ such that ϕ on A^V contracts by the factor γ. This allows us to consider contraction after multiple mappings. We will say a sequence σ in $\{1, 2, ..., n\}$ is an allowed sequence if $\sigma_{k+1} \neq \sigma_k$ for any k, and σ_1 is not i for some predetermined i. We then let $f_\sigma^{(t)}(z) = f_{\sigma_t}(f_{\sigma_{t-1}}(...f_{\sigma_1}(z)...))$.

Lemma 3.3. *There is some constant $\kappa > 0$ such that, for any short enough interval $[u, w] = \{tu + (1-t)w | t \in [0, 1]\} \subset A_i^V$ for any i, and anly allowed sequence σ,*

$$|f_\sigma^{(t)}(w) - f_\sigma^{(t)}(u)| \leq \gamma^t |w - u| \tag{3}$$

and

$$e^{-\kappa|w-u|} \frac{|v-u|}{|w-u|} \leq \frac{|f_\sigma^{(t)}(v) - f_\sigma^{(t)}(u)|}{|f_\sigma^{(t)}(w) - f_\sigma^{(t)}(u)|} \leq e^{\kappa|w-u|} \frac{|v-u|}{|w-u|} \tag{4}$$

for $t = 1, 2, ...$ and any $v \in (u, w)$.

Proof. Because $f_j(A_i^V) \subset A_j^V$ for any pair $i \neq j$, the first equation results immediately from Lemma 3.2.

For the second inequality, for points a, b on the same A_j^V, define $g : [a, b] \to [0, \gamma]$ by $g(ta + (1-t)b) = u$, where u is the number such that $f(ta + (1-t)b) = (uf(a) + (1-u)f(b)) \frac{|f(a)-f(b)|}{|a-b|}$. Expanding g around a with Taylor's theorem, for $c \in (a, b)$ we get

$$\frac{g(c) - g(a)}{g(b) - g(a)} = \frac{|c-a|}{|b-a|}(1 + \eta)$$

where $\eta = O(b - a)$.

Repeating this, we get

$$\frac{|f_\sigma^{(t)}(v) - f_\sigma^{(t)}(u)|}{|f_\sigma^{(t)}(w) - f_\sigma^{(t)}(u)|} = \frac{|c-a|}{|b-a|} \prod_{m=1}^{t} (1 + \eta_m)$$

with $|\eta_m| < \kappa |w - u| \gamma^n$. Here κ can be chosen independent of u, v, w, and σ, since g is $\mathcal{C}^2$ on the compact set $\cup A_i^V$. $\square$

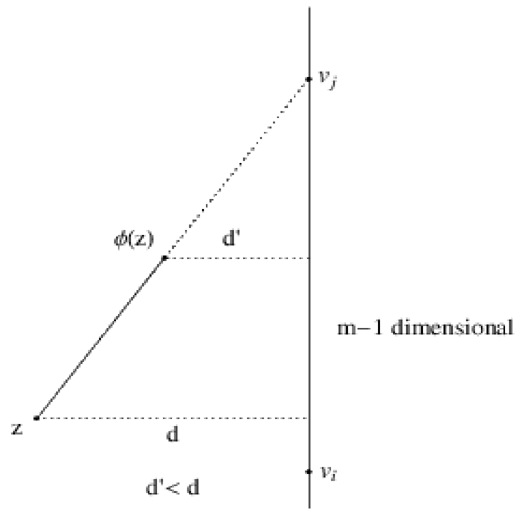

Figure 4: Case of Theorem 3.4

We now give a characterization of the behavior of trajectories.

Theorem 3.4. *All points z on a periodic orbit lie in A^V for any dimension number of nodes n. After a trajectory visits all n nodes, it is inside A^V.*

Proof. We show that no point outside A^V lies on a periodic orbit. The set A^V is closed under the forward map, so any periodic orbit with a point outside A^V must lie entirely outside of A^V. Suppose that we have such an orbit, and that it goes to m nodes, $m < n$.

Take the $m-1$-plane containing the nodes involved in the cycle. No points in the orbit lie in this $m-1$-plane. But whenever we map a point in the orbit to the next point in the orbit, the distance of this point to the $m-1$-plane is less than the distance of our previous point. See Figure 4. But after going through the orbit, we arrive at our original point, a contradiction. Thus we can't have a cycle outside of A^V that involves fewer than n nodes.

There is still the case when we have a periodic orbit that goes to all n-sides. We will show, by induction on n, that any trajectory that visits n nodes will be inside A^V when sent to the nth node.

The base case, $n = 3$, was shown in Lemma 4.4 in [1], and is not difficult to see by drawing a picture.

Now assume the statement for $n-1$, and take a trajectory that hits all n-nodes. Consider the first $n-1$ sides hit by the orbit. The corresponding $n-1$ of the v_i form an $n-1$-simplex, and if the orbit goes inside the projection of this simplex out from the n-th node, it will stay in this projection and will then go inside A^V when projected towards the n-th node.

Consider this process on the first $n-2$ of the v_i. By the induction hypothesis, the orbit ends up in the projection of the $n-2$-simplex with respect to the nth node. When it goes towards the $n-1$st node, it then stays within the projection of the $n-1$-simplex with respect to the nth node. It then stays within this projection until it goes to the nth node, entering A^V. □

A consequence of the above proof is the following.

Corollary 3.5. *For any starting point $z(0)$, there is some $t_0 > 0$ such that $z(t) \in C(\gamma)$ for all $t \geq t_0$.*

Proof. If a point ever enters A^V, then it will stay in there. If not, the argument in the proof of Theorem 3.4 shows that the orbit hits at most $n-1$ sides. It also shows that the orbit gets arbitrarily close to the simplex formed by these at most $n-1$ nodes corresponding to these sides, and thus gets arbitrarily close to the n-simplex formed by the nodes. It then eventually enters the contracting region. $\square$

The next lemma and its proof are unchanged from the 3-node case given in [1]. A node-cycle for a trajectory is a sequence in $\{1, ..., n\}$ corresponding to the order in which the trajectory visits the nodes.

Lemma 3.6. *For any eventually m-periodic trajectory $z(t)$, $t = 0, 1, 2, ...$:*

1. *there exists an orbit $u_1, ..., u_m$ of period m onto which the trajectory $z(t)$ converges;*
2. *trajectories $z'(t)$ with the same node-cycle converge onto the same orbit;*
3. *there can be at most one orbit having a given node-cycle.*

Proof. 1. Corollary 3.5 shows that there is some t_0 such that for $t \geq t_0$ we have $z(t) \in \mathcal{C}(\gamma)$. (LEMMA 4.2) implies

$$|z((t+1)m+i) - z(tm+1)| \leq \gamma^m |z(tm+i) - z((t-1)m+i)|,$$

and thus $u_i = \lim_{t\to\infty} z(tm+i)$ exists. For 2 and 3, two trajectories with the same node-cycle must converge onto the same orbit since otherwise the two distinct orbits would not satisfy the contraction property. $\square$

4 Further Research

For the $n = 3$ case, [1] proves the following:

Theorem 4.1. *For almost all decision points $d_i \in A$, $i = 1, 2, 3$, the triangle process Z has finitely many periodic orbits. For such sets of decision points, all trajectories $z(t)$ are eventually periodic and each converges onto one of these orbits as $t \to \infty$.*

However, for $n > 3$, we were unable to prove the above. Let $\mathcal{P} := \{\phi^{(-t)}(d_i)|i = 1, ..., n; t = 0, 1, 2, ...\}$ be the set of preimages of the decision points. When $n = 3$, if $\mathcal{P}$ is finite, then $\mathcal{P}$ partitions A into a finite number of intervals whose interior will never hit a decision point. Each interval will then converge to a single periodic orbit, with endpoints converging to these also, giving finitely many orbits. They prove that for almost all decision points, $\mathcal{P}$ is finite.

For an intuitive idea of why, one can use Theorem 3.4 to show that the set of points with an infinite number of preimages has measure zero. This does not gives the proof, because which points have infinitely many preimages depends on the decision points, but gives some idea for why the theorem makes sense.

However, the situation is more complicated for $n > 3$. It's possible that all decision points have a finite number of preimages, but there are some points on the decision lines (for $n = 4$, an $n-2$-dimensional subset for n) that have infinitely many preimages. What if, for $n = 4$, there is a set of points with infinitely many preimages that forms a triangle around one of the d_i? Then any decision line from that d_i would have a point with infinitely many preimages.

Based on computer simulations for $n = 4$, we hypothesize that the above theorem holds, at least for $n = 4$, but were unable to make progess on it. We do have the following partial result:

Lemma 4.2. *If $\mathcal{R} := \{$preimages of points on decision lines$\}$ splits A into finitely many regions, then the n-process has a finite number of periodic orbit and each trajectory converges onto one of these.*

Proof. Since each region will never contain an image of a decision point, it will remain connected under repeated mappings by ϕ, each time going into another of the regions. Each point in such a region will eventually enter the contracting region, and will then contract. Since each node-cycle contains a single orbit and the other points in the region will share this node-cycle, they will converge to this orbit by Lemma 3.6. For points on preimages of decision lines, these will each follow one of these orbits. □

5 Acknowledgements

Thanks to Misha Guysinsky and Shilpak Banerjee for guidance in writing this paper and assisting development of the above ideas.

References

[1] Macphee, I. M.; Menshikov, M. V.; Popov, S.; and Volkov, S (2006). Periodicity in the Transient Regime of Exhaustive Polling Systems. Ann. Appl. Probab. Vol. 16, No. 4, 1816-1850.