MARKOV CHAINS AND ANT COLONY OPTIMIZATION

WAYNE YANG

Abstract. In this paper we introduce the notion of a Markov chain and some of its basic properties. We apply Markov chains to the analysis of Ant Colony Optimization algorithms. In particular, we look at the Max-Min Ant System for finding shortest path and use the mixing times of Markov chains to provide upper-bounds on running time estimation.

Contents

1. Introduction to Markov Chains

Markov chains are a large class of stochastic processes characterized by the **Markov property** which simply means that the next state that the process takes depends only on the current state and none of the previous states. Formally, we can express this description for processes in discrete time in the following way:

Definition 1.1. A stochastic process is said to have the **Markov property** if:

$$\mathbb{P}\{X_{t+1} = x \,|\, X_0 = x_0, ..., X_t = x_t\} = \mathbb{P}\{X_{t+1} = x \,|\, X_t = x_t\}. \tag{1.2}$$

Such processes are called **Markov Chains**.

This definition implicitly restricts Ω to be countable because in the uncountable setting, the probability measure of a single point is necessarily 0. Markov chains over uncountable state-spaces require a more nuanced definition.

We often make the assumption that the transition probabilities do not change over times. If this is the case, then we can describe the conditional probability that $X_{t+1} = x$ given that $X_t = y$ as a function strictly in x and y. This gives rise to the following definition:

Date: August 2013.

Definition 1.3. A Markov chain is said to be **time-homogeneous** if for all $t > 0$:

$$\mathbb{P}\{X_{t+1} = x | X_t = y\} = p(y, x) \text{ for some } p : \Omega \times \Omega \to [0, 1] .$$

We can interpret p as the probability of jumping from state y to state x in a single time-step. For the rest of the paper, all Markov chains will be time-homogeneous.

Definition 1.4. For a Markov chain $\{X_t\}$ with state space Ω, the **transition matrix** is an $|\Omega| \times |\Omega|$ matrix P (possibly infinite) such that each row index and each column index corresponds to a state and the entries are given by:

$$P_{x,y} = p(x, y),$$

where $p(x, y)$ is the transition probability of jumping from state x to state y.

We note that $0 \leq P_{x,y} \leq 1$ since they are probabilities and that $\sum_{y \in \Omega} P_{x,y} = 1$ since the chain takes some state in Ω with probability 1 and each state is a disjoint event.

Definition 1.5. We define the **t-step probability** denoted $p_t(x, y)$ to be the probability that $X_t = y$ given that $X_0 = x$.

Lemma 1.6. *For a Markov chain $\{X_t\}$ with state space Ω and transition matrix P, the (x, y) entry of the matrix P^t is precisely $p_t(x, y)$.*

Proof. For $t = 1$, the statement is true by definition. So assume that the lemma holds for P^t, then consider:

$$\begin{aligned} p_{t+1}(x, y) = & \sum_{w \in \Omega} P^t(x, w) p(w, y) \\ = & \sum_{w \in \Omega} P^t_{x,w} P_{w,y} \\ = & P^{t+1}_{x,y}. \end{aligned}$$

The first equality follows from the total law of probability and the second equality follows from the inductive hypothesis. □

From this point forward, we will use $P^t(x, \cdot)$ to denote the distribution on Ω of the t^{th} step of a chain starting at x.

2. Ergodicity

So far we have a very broad definition of Markov chain, but we are interested in a certain class of Markov chains with nice properties. Therefore, we give the following definitions:

Definition 2.1. A state $x \in \Omega$ is said to **communicate** with a state $y \in \Omega$ if and only if there exist $s, t > 0$ such that:

$$\begin{aligned} \mathbb{P}\{X_t = y \,|\, X_0 = x\} > 0 \\ \mathbb{P}\{X_s = x \,|\, X_0 = y\} > 0. \end{aligned}$$

We denote this as $x \leftrightarrow y$.

Definition 2.2. A **communication class** for a Markov chain is set of states $W \subseteq \Omega$ such that:

$$\begin{aligned} x &\leftrightarrow y, \quad \forall x, y \in W \\ w &\not\leftrightarrow z \quad \forall w \in W,\ \forall z \in \Omega \setminus W. \end{aligned}$$

Definition 2.3. A Markov chain is **irreducible** if the sample space Ω is a communication class.

It is easy to check that $\leftrightarrow$ forms a well defined equivalence relation (and thus a partition) on Ω. Knowing this fact, we can interpret an irreducible Markov chain as one which only has the simple partition.

Definition 2.4. The **period** of a state $x \in \Omega$ is defined to be:

$$L(x) = gcd\{t \,|\, \mathbb{P}\{X_t = x \,|\, X_0 = x\} > 0\}.$$

We say x is **aperiodic** if and only if $L(x) = 1$.

Definition 2.5. A Markov chain $\{X_t\}$ is said to be **aperiodic** if and only if every state $x \in \Omega$ is aperiodic.

Lemma 2.6. *Let $\{X_t\}$ be a Markov chain over state space Ω, then for any communication class $W \subseteq \Omega$,*

$$L(x) = L(y), \ \ \forall x, y \in W$$

Proof. Since x, y are in the same communication class, $x \leftrightarrow y$ and we have:

$$\exists r, s \text{ such that } P^r(x, y) > 0, \ P^s(y, x) > 0$$

Which implies,

$$P^{r+s}(x, x) \geq P^r(x, y)P^s(y, x) > 0.$$

Therefore,

$$L(x) \mid r + s.$$

For any $t > 0$ such that $P^t(y, y) > 0$, we have that:

$$P^{r+s+t}(x, x) \geq P^r(x, y)P^t(y, y)P^s(y, x) > 0.$$

Which means that,

$$L(x) \mid r + s + t \ \Rightarrow L(x) \mid t \ \Rightarrow L(y) \mid L(x).$$

Then, by a symmetric argument, $L(x) \mid L(y)$. Thus,

$$L(x) = L(y).$$

□

In particular, Lemma 2.6 guarantees that if we find that one state in an irreducible Markov chain is aperiodic, the whole chain must be.

Definition 2.7. A state $x \in \Omega$ is said to be **recurrent** if and only if:

$$\mathbb{P}\{X_t = x \text{ for infinitely many } t > 0\} = 1.$$

A Markov chain is recurrent if every state is recurrent.

Definition 2.8. A state $x \in \Omega$ is said to be **transient** if and only if it is not recurrent. A Markov chain is transient if every state is transient.

Recurrence / transience simply describes whether or not it is possible to leave a state and never return. However, the above definition is often difficult to directly demonstrate which leads us to an alternative characterization using the notion of return times.

Definition 2.9. The **return time** of a state x is the random variable T_x which is characterized by:

$$T_x = \inf_t \{t \mid X_t = x, X_0 = x\}.$$

Theorem 2.10. *Let $\{X_t\}$ be a Markov chain with transition matrix P over a state space Ω. Then for any state $x \in \Omega$, the following are equivalent:*

***i.*:** *x is recurrent.*
***ii.*:** $\mathbb{P}\{T_x < \infty\} = 1$.
***iii.*:** $\sum_{t=0}^{\infty} P^t(x,x) = \infty$.

Proof. (*i.* $\iff$ *ii.*)

Suppose that x is recurrent. By definition, the chain must hit state x infinitely often. In particular, the chain must hit the state x at least twice. This guarantees the existence of non-negative integers $s < t$ such that:

$$X_s = x \text{ and } X_t = x.$$

By the Markov property and our assumption that the chain is time-homogeneous,

$$\mathbb{P}\{X_t = x, X_s = x\} = \mathbb{P}\{X_{t-s} = x, X_0 = x\}.$$

Which holds for all $s < t$. Therefore,

$$\mathbb{P}\{T_x < \infty\} = \mathbb{P}\{t - s < \infty\} = 1.$$

Conversely, suppose that T_x is always finite. Then, suppose by contradiction that x is a transient state.

$$\exists\, t > 0 \text{ such that } X_t = x \text{ and } \forall\ s > t,\ X_s \neq x.$$

Then by the Markov Property,

$$\mathbb{P}\{X_t = x, X_s \neq x\ \forall\, s < t\} = \mathbb{P}\{X_0 = x, X_n \neq x\ \forall\, n > 0\}.$$

However, if this occurs, it must be true that:

$$T_x = \inf_t \{t \mid X_t = x, X_0 = x\} = \infty.$$

But this contradicts our assumption that the return time is finite.

(*ii.* $\iff$ *iii.*)

We define the number of visits to x for a chain starting at x to be:

$$V_x = \sum_{t=0}^{\infty} 1_{\{X_t = x, X_0 = x\}}.$$

Then we consider:

$$\begin{aligned}\sum_{t=0}^{\infty} P^t(x,x) = &\sum_{t=0}^{\infty} \mathbb{P}\{X_t = x, X_0 = x\} \\ = &\sum_{t=0}^{\infty} \mathbb{E}1_{\{X_t=x, X_0=x\}} \\ = &\mathbb{E}\sum_{t=0}^{\infty} 1_{\{X_t=x, X_0=x\}} \\ = &\mathbb{E}V_x.\end{aligned}$$

But V_x is infinite if and only if T_x is finite. □

The next lemma demonstrates that recurrence is communication class property:

Lemma 2.11. *Suppose $\{X_t\}$ is a Markov chain with state space Ω, suppose $W \subseteq \Omega$ is a communication class. Then given any state x in W:*

$$x \in W \text{ is transient} \iff y \in W \text{ is transient } \forall y \in W$$

Proof. Suppose that $x \in W$ is transient and fix $y \in \Omega$. Since W is a communication class, there exist positive integers s, t such that:

$$P^s(x,y) > 0 \text{ and } P^t(y,x) > 0$$

Then we observe that for all $r \geq 0$:

$$P^{s+r+t}(x,x) \geq P^s(x,y)P^r(y,y)P^t(y,x)$$

This is true because the left side of the inequality is the probability of the event $A = \{X_0 = x, X_{s+r+t} = x\}$ and the right side is probability of the event $B = \{X_0 = x, X_s = y, X_{s+r} = y, X_{s+r+t} = x\}$. Then it is clear that $B \subseteq A$ and thus we get the above inequality.

Summing over every $r > 0$, we get that:

$$\begin{aligned}&\sum_{r=0}^{\infty} P^{s+r+t}(x,x) \geq \sum_{r=0}^{\infty} P^s(x,y)P^r(y,y)P^t(y,x) \\ \Rightarrow &\sum_{r=0}^{\infty} P^r(y,y) \leq \frac{1}{P^s(x,y)P^t(y,x)} \sum_{r=0}^{\infty} P^{s+r+t}(x,x) \\ &\leq \frac{1}{P^s(x,y)P^t(y,x)} \sum_{k=0}^{\infty} P^k(x,x) < \infty\end{aligned}$$

The last inequality holds because we assume x to be transient and therefore the sum must be finite by Theorem 2.10. Then, y must be transient as well by Theorem 2.10. □

It immediately follows that any state x is recurrent if and only if every other state in the communication class of x is recurrent. In particular, we notice that if a Markov chain is irreducible, then the whole system must either be transient or recurrent.

When dealing with Markov chains with a countably infinite state-space, we must make one further distinction within recurrent states.

Definition 2.12. Let $\{X_t\}$ be a Markov chain with transition matrix P over a countably infinite state space Ω. A state $x \in \Omega$ is said to be **positive recurrent** if and only if it is both recurrent and $\mathbb{E}T_x < \infty$.

A state $x \in \Omega$ is said to be **null recurrent** if and only if it is both recurrent and $\mathbb{E}T_x = \infty$.

Lemma 2.13. *Let $\{X_t\}$ be a Markov chain with transition matrix P over state space Ω. For a communication class $W \subset \Omega$, the following holds:*

$$x \in W \text{ is positive recurrent } \Rightarrow y \in W \text{ is positive recurrent, } \forall\, y \in W.$$

Proof. We first define for any $x, y \in \Omega$:

$$T_{x,y} = \inf_t \{t \mid X_t = y, X_0 = x\}.$$

Let T_y denote the return time to y, then we observe:

$$T_y \leq T_{y,x} + T_{x,y}.$$

Define the event A to be the event that the chain starting at x hits y before returning to x. Then we can split the expected return time of x in the following way:

$$\mathbb{E}T_x = \mathbb{E}(T_x|A)\mathbb{P}\{A\} + \mathbb{E}(T_x|A^c)(1 - \mathbb{P}\{A\}).$$

But if the event A occurs, then the return time for x can be split into the time it takes to reach y before going back, ie:

$$\mathbb{E}(T_x|A) = \mathbb{E}(T_{x,y} + T_{y,x}).$$

We know that $\mathbb{P}\{A\}$ and $\mathbb{E}T_x$ are both finite and positive alues, which gives that:

$$\mathbb{E}T_y \leq \mathbb{E}(T_{y,x} + T_{x,y}) \leq \frac{\mathbb{E}T_x}{\mathbb{P}\{A\}} < \infty.$$

So y must be positive recurrent.

□

The sort of Markov chains we are interested in are precisely the ones which are irreducible, aperiodic and positive recurrent. Markov chains which satisfy these properties have a special name. Therefore, we provide the following definition.

Definition 2.14. A Markov chain $\{X_t\}$ is said to be **ergodic** if it is irreducible, aperiodic, and positive recurrent.

3. The Stationary Distribution

The reason that ergodicity is a useful property is that the distribution of the t^{th} time step of an ergodic Markov chain will stabilize over time. First, we define what such a Markov chain will converge to and develop several important results.

Definition 3.1. For a Markov chain $\{X_t\}$ over a state space Ω with transition matrix P, a distribution π on Ω (represented as a vector of probabilities) is called **stationary** if it satisfies:

$$\pi P = \pi.$$

Or equivalently:

$$\pi(x) = \sum_{y \in \Omega} \pi(y) p(y, x)\ \forall\, x \in \Omega.$$

A Markov chain with a stationary distribution has at least two interesting properties. First, if the initial state of the chain is distributed according to π, then every subsequent step is also distributed according to π. Second, we will show that under mild conditions, the distribution of X_t will converge to the stationary distribution under any initial distribution. We will closely follow a proof which can be found in Levin, Peres, and Wilmer [2].

Proposition 3.2. *For an ergodic Markov chain $\{X_t\}$ over the state space Ω with transition matrix P, there exists a stationary distribution π on Ω.*

Proof. Fix an arbitrary state $z \in \Omega$. Since $\{X_t\}$ is ergodic, it is recurrent and will always have a finite return time to the state z. So we consider the average number of times the chain visits other states before returning to z by examining the chain when it starts at $X_0 = z$. Then, we define:

$$\begin{aligned}\tilde{\pi}(y) = &\quad \mathbb{E}\{\#\text{ of visits to } y \text{ before returning to } z\} \\ = &\quad \sum_{t=0}^{\infty} \mathbb{P}\{X_t = y, T_z > t\}.\end{aligned}$$

While not yet a distribution a priori, we claim that $\tilde{\pi}$ does satisfy stationarity:

$$\begin{aligned}\sum_{x\in\Omega} \tilde{\pi}(x)p(x,y) = &\quad \sum_{x\in\Omega}\sum_{t=0}^{\infty} \mathbb{P}\{X_t = x, T_z > t\}p(x,y) \\ = &\quad \sum_{x\in\Omega}\sum_{t=0}^{\infty} \mathbb{P}\{X_t = x, T_z \geq t+1, X_{t+1} = y\} \\ = &\quad \sum_{t=0}^{\infty} \mathbb{P}\{X_{t+1} = y, T_z \geq t+1\}.\end{aligned}$$

The last equality holds by switching the order of summation and realizing that the events $\{X_t = x, T_z \geq t+1, X_{t+1} = y\}$ and $\{X_t = x', T_z \geq t+1, X_{t+1}\}$ are disjoint whenever $x \neq x'$. Furthermore, we have that:

$$\bigcup_{x\in\Omega} \{X_t = x, T_z \geq t+1, X_{t+1} = y\} = \{X_{t+1} = y, T_z \geq t+1\}.$$

We can re-index the summation starting at 1 and consider the two possible (and disjoint) cases for the event in question. Either $X_t = y$ and $T_z > t$ or $X_t = y$ and $T_z = t$. Then we have that:

$$\begin{aligned}\sum_{t=0}^{\infty} \mathbb{P}\{X_{t+1} = y, T_z \geq t+1\} = &\quad \sum_{t=1}^{\infty} \mathbb{P}\{X_t = y, T_z > t\} + \sum_{t=1} \{X_t = y, T_z = t\} \\ = &\quad \tilde{\pi}(y) - \mathbb{P}\{X_0 = y, T_z > 0\} + \sum_{t=1} \{X_t = y, T_z = t\} \\ = &\quad \tilde{\pi}(y).\end{aligned}$$

Equality holds because either $y = z$ or $y \neq z$. If $y = z$, then both of the terms on the right are 1 by definition. If $y \neq z$, then both are zero.

Then, in order to go from $\tilde{\pi}$ to a valid distribution, we observe that the sum of the number visits to other states before returning to z is exactly the return time of z. Hence, we have the following equality:

$$\sum_{x\in\Omega}\tilde{\pi}(x) = \mathbb{E}T_z.$$

Since we assume z to be positive recurrent, $\mathbb{E}T_z$ must be finite and we can normalize $\tilde{\pi}(x)$ by $\mathbb{E}T_z$ to get a distribution π such that for all $x \in \Omega$:

$$\pi(x) = \frac{\tilde{\pi}(x)}{\mathbb{E}T_z}.$$

□

Proposition 3.3. *If an ergodic Markov chain $\{X_t\}$ over state space Ω has a stationary distribution π, then $\pi(x) > 0$ for all $x \in \Omega$.*

Proof. Suppose by contradiction that there is some $x \in \Omega$ such that $\pi(x) = 0$. Then by the definition of stationarity we have:

$$\pi(x) = \sum_{y\in\Omega}\pi(y)p(y,x).$$

Therefore,

$$\pi(y) = 0 \ \forall y \in \Omega \text{ st } p(y,x) > 0.$$

Then we can construct a sequence of sets $A_1 \subset A_2 \subset A_3...$ in the following way:

$$\begin{aligned}
A_1 = &\quad \{x\}\\
A_2 = &\quad A_1 \cup \{y \,|\, p(y,x) > 0 \,\forall x \in A_1\}\\
A_3 = &\quad A_2 \cup \{z \,|\, p(z,y) > 0 \,\forall y \in A_2\}\\
&\quad \vdots\\
A_n = &\quad A_{n-1} \cup \{u \,|\, p(u,z) > 0 \,\forall z \in A_{n-1}\}.\\
&\quad \vdots
\end{aligned}$$

By induction, it's clear that $\pi(x) = 0$ for all $x \in A_n$ for any $n > 0$. Then since $\{X_t\}$ is irreducible there is a path from any $y \in \Omega$ to x, so the construction of A_n guarantees that:

$$\bigcup_{n=1}^{\infty} A_n = \Omega$$
$$\Rightarrow \pi(x) = 0 \,\forall x \in \Omega$$

Which contradicts the assumption that π is a valid distribution over Ω.

□

For the next lemma, we can generalize the idea of a return time for a single point to a subset $S \subseteq \Omega$ by defining:

$$T_S = \inf_t\{t \mid X_t \in S \mid X_0 \in S\}.$$

Lemma 3.4. *(Kac) If an irreducible Markov chain $\{X_t\}$ with transition matrix P and state space Ω has a stationary distribution, then for any subset $S \subseteq \Omega$, we have that:*

$$\sum_{x\in S}\pi(x)\mathbb{E}T_S = 1.$$

Proof. Define another Markov chain $\{Y_t\}$ over Ω with transition matrix $\tilde{P}$ where the entries in $\tilde{P}$ are given by:

$$\tilde{P}(x,y) = \frac{\pi(y)P(y,x)}{\pi(x)}.$$

First, we claim that π is also stationary for $\{Y_t\}$ since we have:

$$\begin{aligned}\sum_{y\in\Omega}\pi(y)\tilde{P}(y,x) = &\sum_{y\in\Omega}\pi(y)\frac{\pi(x)P(x,y)}{\pi(y)}\\ = &\sum_{y\in\Omega}\pi(x)P(x,y)\\ = &\pi(x).\end{aligned}$$

Furthermore, if we use the fact that $P(x,y) = \pi(x)\tilde{P}(y,x)/\pi(y)$, we have that:

$$\begin{aligned}\mathbb{P}\{X_0 = x_0, ..., X_t = x_t\} = &\pi(x_0)P(x_0,x_1)\cdot ... \cdot P(x_{t-1},x_t)\\ = &\pi(x_t)\tilde{P}(x_t,x_{t-1})\cdot ... \cdot \tilde{P}(x_1,x_0)\\ = &\mathbb{P}\{Y_0 = x_t, ..., Y_t = x_0\}.\end{aligned}$$

Then we claim that both $\{X_t\}$ and $\{Y_t\}$ are recurrent. To see this, we define for any $x \in S$.

$$\begin{aligned}\alpha(t) = &\mathbb{P}\{X_t = x, X_s \neq x \text{ for } s > t\}\\ = &\mathbb{P}\{X_t = x\}\mathbb{P}\{T_x = \infty\}\\ = &\pi(x)\mathbb{P}\{T_x = \infty\}.\end{aligned}$$

For each natural number t, there is a clear association between $\alpha(t)$ and the event $\{X_t = x, X_s \neq x \text{ for } s > t\}$. Furthermore, the events associated with $\alpha(r)$ and $\alpha(s)$ are mutually exclusive whenever $r \neq s$. Therefore,

$$\sum_{t=0}^{\infty}\alpha(t) \leq 1.$$

It's clear that $\alpha(t)$ does not actually depend on t and thus must equal zero for all t. However, by Proposition 3.3, $\pi(x) > 0$ and thus $\mathbb{P}\{T_x = \infty\} = 0$ and $\{X_t\}$ must be recurrent. A similar argument can be used to show that $\{Y_t\}$ is also recurrent.

Next, we consider the identity:

$$\pi(z_0)P(z_0,z_1)\cdot ... \cdot P(z_{t-1},z_t) = \pi(z_t)\tilde{P}(z_t,z_{t-1})\cdot ... \cdot \tilde{P}(z_1,z_0).$$

If we set $z_0 = x$ for some $x \in S$, $z_t = y$ for some $y \in \Omega$, we have that:

$$\sum_{z_1,...,z_{t-1}\notin S}\pi(x)P(x,z_1)\cdot ... \cdot P(z_{t-1},y) = \sum_{z_1,...,z_{t-1}\notin S}\pi(y)\tilde{P}(z_t,z_{t-1})\cdot ... \cdot \tilde{P}(z_1,x).$$

This simplifies to:

$$\pi(x)\mathbb{P}\{T_s \geq t, X_t = y\} = \pi(y)\mathbb{P}\{\tilde{T}_S = t, Y_t = x\},$$

where $\tilde{T}_S = \inf_t\{t > 0 \mid Y_t \in S\}$.

Then, if we sum this expression over $x \in S$ and $y \in \Omega$, we get:

$$\sum_{x\in S}\pi(x)\sum_{t=1}^{\infty}\sum_{y\in\Omega}\mathbb{P}\{T_s \geq t, X_t = y\} = \sum_{y\in\Omega}\sum_{t=0}^{\infty}\sum_{x\in\Omega}\mathbb{P}\{\tilde{T}_S = t, Y_t = x\}.$$

$$\Rightarrow \sum_{x \in S} \pi(x) \sum_{t=1}^{\infty} \mathbb{P}\{T_S \geq t\} = 1.$$

But T_S is a non-negative random variable and the fact that $\{Y_t\}$ is recurrent.

$$\sum_{x \in S} \pi(x) \mathbb{E}T_S = \mathbb{P}\{\tilde{T}_S < \infty\} = 1.$$

□

Proposition 3.5. *If an irreducible and aperiodic Markov chain $\{X_t\}$ has a stationary distribution π, it must be positive recurrent (and therefore ergodic).*

Proof. By the Kac Lemma, we have that:

$$\pi(x)\mathbb{E}T_x = 1.$$

By Proposition 3.3, $\pi(x) > 0$, so:

$$\mathbb{E}T_x = \frac{1}{\pi(x)} < \infty.$$

□

4. Convergence to Stationarity

In order to get any sort of convergence result, we need to define precisely what we mean by distance between two distributions over the same state space. Intuitively, we want a number that describes how different the probabilities of attaining a state x are for the two distributions.

Definition 4.1. Given two distributions μ, ν on a sample space Ω, we define the **total variation distance** to be:

$$\|\mu - \nu\|_{TV} = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|.$$

One of the tools we will use to control total variation distance is known as coupling, which links together two different random variables as a joint distribution.

Definition 4.2. Suppose that μ and ν are two arbitrary distributions on the state space Ω. A **coupling** of μ and ν is a pair of random variables (X, Y) such that X and Y are jointly distributed over $\Omega \times \Omega$ and satisfy:

$$\begin{aligned} \mathbb{P}\{X = x\} = &\quad \mu(x) \ \forall\, x \in \Omega \\ \mathbb{P}\{Y = y\} = &\quad \nu(y) \ \forall\, y \in \Omega. \end{aligned}$$

The relation between X and Y can be anything as long as it's a valid joint distribution and the marginal distributions of X and Y are μ and ν, respectively. For example, allowing X and Y to be independent is still a valid coupling. On the other hand, if $\mu = \nu$, then $X = Y$ is also a valid coupling. However, these constructions are often not the most informative and useful couplings fall somewhere between these two extremes.

Proposition 4.3. *Let μ and ν be two distributions on Ω. Then the following equality holds:*

$$\|\mu - \nu\|_{TV} = \inf\{\mathbb{P}\{X \neq Y\} \mid (X, Y) \text{ is a coupling of } \mu \text{ and } \nu\}.$$

Proof. Consider any coupling (X, Y) of μ and ν and any event $A \subset \Omega$:

$$\begin{aligned} \mu(A) - \nu(A) = &\ \mathbb{P}\{X \in A\} - \mathbb{P}\{Y \in A\} \\ \leq &\ \mathbb{P}\{X \in A, Y \notin A\} \\ \leq &\ \mathbb{P}\{X \neq Y\}. \end{aligned}$$

$$\Rightarrow \ \|\mu - \nu\|_{TV} \leq \inf\{\mathbb{P}\{X \neq Y\} \mid (X, Y) \text{ is a coupling of } \mu \text{ and } \nu\}.$$

In order to prove equality, we directly construct the coupling where equality holds. This coupling is called the **optimal coupling**.

Define:

$$\begin{aligned} p = &\ \sum_{x \in \Omega} \min\{\mu(x), \nu(x)\} \\ = &\ \sum_{\substack{x \in \Omega \\ \mu(x) < \nu(x)}} \mu(x) + \sum_{\substack{x \in \Omega \\ \mu(x) \geq \nu(x)}} \nu(x). \end{aligned}$$

Consider the following manipulation of the definition of total variation distance:

$$\begin{aligned} \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(y)| = &\ \frac{1}{2} \sum_{\substack{x \in \Omega \\ \mu(x) \geq \nu(x)}} (\mu(x) - \nu(x)) + \frac{1}{2} \sum_{\substack{x \in \Omega \\ \mu(x) < \nu(x)}} (\nu(x) - \mu(x)) \\ = &\ \sum_{\substack{x \in \Omega \\ \mu(x) \geq \nu(x)}} \mu(x) - \nu(x). \end{aligned}$$

The fact the last inequality holds can be seen by defining the event $A = \{x \mid \mu(x) \geq \nu(x)\}$ and observing that:

$$\sum_{\substack{x \in \Omega \\ \mu(x) \geq \nu(x)}} (\mu(x) - \nu(x)) = \mu(B) - \nu(B)$$

$$\sum_{\substack{x \in \Omega \\ \mu(x) < \nu(x)}} (\nu(x) - \mu(x)) = \nu(B^c) - \mu(B^c).$$

Then subtracting them, we get equality.

Returning back to the construction of p, the last argument now gives that:

$$p = 1 - \|\mu - \nu\|_{TV}.$$

The coupling works in the following way:

(1) Flip a coin with probability of landing on heads given by p.
(2) If the coin is heads, choose a value Z from Ω where each $x \in \Omega$ has probability of being chosen given by:

$$\gamma_{same} = \frac{\min\{\mu(x), \nu(x)\}}{p}$$

Then set $X = Y = Z$.
(3) If the coin is tails, then we choose x and Y according the following probabilities:

$$\gamma_X(x) = \mathbb{P}\{X = x\} = \begin{cases} \frac{\mu(x) - \nu(x)}{\|\mu - \nu\|_{TV}} \text{ if } \mu(x) > \nu(x) \\ 0 \text{ otherwise} \end{cases}$$

$$\gamma_Y(y)\mathbb{P}\{Y=y\} = \begin{cases} \frac{\nu(x)-\mu(x)}{\|\mu-\nu\|_{TV}} \text{ if } \nu(x) > \mu(x) \\ 0 \text{ otherwise} \end{cases}$$

Then it's clear that:

$$p\gamma_{same} + (1-p)\gamma_X = \mu$$
$$p\gamma_{same} + (1-p)\gamma_Y = \nu.$$

If the coin lands on tails, then $X \neq Y$ by construction. Hence:

$$\mathbb{P}\{X \neq Y\} = 1 - p = 1 - (1 - \|\mu - \nu\|_{TV}) = \|\mu - \nu\|_{TV}.$$

□

Definition 4.4. A **coupling of Markov chains** with transition matrix P is a sequence of pairs of random variables $\{(X_t, Y_t)\}_{t=0}^{\infty}$ such that:

***i*:** $\{X_t\}$ and $\{Y_t\}$ are both Markov chains are transition matrix P over the state space Ω (with possibily different initial distributions).

***ii*:** (X_t, Y_t) is a coupling for each $t \geq 0$.

We can modify any coupling of Markov chains such that the X_t and Y_t behave the same after any point in which the two are in the same state at the same time. In mathematical terms, we want:

$$\text{if } X_s = Y_s, \text{ then } X_t = Y_t \, \forall \, t \geq s. \tag{4.5}$$

Theorem 4.6. *Let $\{(X_t, Y_t)\}$ be any coupling of Markov chains, each with transition matrix P and state space Ω, such that equation 2.29 holds. Let T_{couple} denote the time in which the two chains meet. i.e.*

$$T_{couple} = \min\{t \mid X_t = Y_t\}.$$

Then, if we consider the case where $X_0 = x$ and $Y_0 = y$:

$$\|P^t(x,\cdot) - P^t(y,\cdot)\|_{TV} \leq \mathbb{P}\{T_{couple} > t\}.$$

Proof. Since X_t and Y_t are marginally distributed as the single Markov chain with transition matrix P over Ω, we have that:

$$P^t(x,z) = \mathbb{P}\{X_t = z\} \text{ and } P^t(y,z) = \{Y_t = z\}.$$

Since the coupling is assumed to satisfy equation 4.5, the event that $\{T_{couple} > t\}$ occurs if and only if $X_t \neq Y_t$ (otherwise they have previously met or meet at t). Then Proposition 4.3 implies that:

$$\|P^t(x,\cdot) - P^t(y,\cdot)\|_{TV} \leq \mathbb{P}\{X_t \neq Y_t\} = \mathbb{P}\{T_{couple} > t\}.$$

□

Now we have all the tools necessary to prove convergence of ergodic Markov chains to their unique stationary distributions.

Theorem 4.7. *Let $\{X_t\}$ be an ergodic markov chain with transition matrix P over state space Ω. Then there exists a unique stationary distribution π on Ω such that for any initial distribution μ:*

$$\lim_{t\to\infty} \|\mu P^t - \pi\|_{TV} = 0.$$

Proof. Since $\{X_t\}$ is ergodic, there must exist some stationary distribution π by Proposition 3.2.

Consider the Markov chain $\{(X_t, Y_t)\}$ over $\Omega \times \Omega$ where Y_t is an independent copy of X_t and transition probabilities are given by:

$$\tilde{P}((x,y),(z,w)) = P(x,z)P(y,w).$$

First we note that the chain is aperiodic and irreducible in each coordinate. Since the coordinates are independent, the whole chain is aperiodic and irreducible.

Now, we define the distribution $\pi \times \pi$ over $\Omega \times \Omega$ with probability mass functions given by:

$$(\pi \times \pi)((x,y)) = \pi(x)\pi(y),\ \forall (x,y) \in \Omega \times \Omega.$$

Now we claim that $\pi \times \pi$ is stationary for $\{(X_t, Y_t)\}$ since:

$$\begin{aligned}
(\pi)\tilde{P}(z,w) = & \sum_{(x,y)\in\Omega\times\Omega} (\pi \times \pi)(x,y)\tilde{P}((x,y),(z,w)) \\
= & \sum_{(x,y)\in\Omega\times\Omega} \pi(x)\pi(y)P(x,z)P(y,w) \\
= & \sum_{x\in\Omega} \pi(x)P(x,z) \sum_{y\in\Omega} \pi(y)P(y,w) \\
= & \ \pi(z)\pi(w) \\
= & \ (\pi \times \pi)(z,w).
\end{aligned}$$

Which is exactly the definition of stationarity. Then by Proposition 3.5, $\{(X_t, Y_t)\}$ is positive recurrent. So if we define for any $x \in \Omega$:

$$T_x = \min\{t > 0 \mid X_t = Y_t = x\}.$$

We have that:

$$\mathbb{P}\{T_x < \infty\} = 1.$$

Furthermore, we have that $\{(X_t, Y_t)\}$ is a coupling of Markov chains and we can modify it such that they run together as soon as they meet (as in equation 4.5). So if we initialize (X_t, Y_t) by the distribution $\mu \times \pi$, we have by Theorem 4.6 that:

$$\|\mu P^t - \pi\|_{TV} \leq \mathbb{P}\{X_t \neq Y_t\} = \mathbb{P}\{T_{couple} > t\}.$$

But since $\{(X_t, Y_t)\}$ is positive recurrent, we have that:

$$\mathbb{P}\{T_{couple} < \infty\} = 1 \ \Rightarrow \lim_{t\to\infty} \mathbb{P}\{T_{couple} > t\} = 0.$$

Therefore:

$$\lim_{t\to\infty} \|\mu P^t - \pi\|_{TV} = 0.$$

Then since this is true for any initial distribution μ, we have that π is unique.

□

5. Mixing Time

Theorem 5.1. *Let $\{X_t\}$ be an ergodic Markov chain with transition matrix P over a countable state space Ω. Let μ and ν be any two distributions over Ω. Denote the distribution over Ω of X_1 with initial distribution μ as μP and the distribution of X_1. Then, we have that:*

$$\|\mu - \nu\|_{TV} \geq \|\mu P - \nu P\|_{TV}.$$

Proof.

$$\begin{aligned}
\|\mu P - \nu P\| = &\quad \frac{1}{2}\sum_{x\in\Omega}|\mu P(x) - \nu P(x)| \\
= &\quad \frac{1}{2}\sum_{x\in\Omega}\left|\sum_{y\in\Omega}\mu(y)P(y,x) - \sum_{y\in\Omega}\nu(y)P(y,x)\right| \\
\leq &\quad \frac{1}{2}\sum_{x\in\Omega}\sum_{y\in\Omega}|\mu(y)P(y,x) - \nu(y)P(y,x)| \\
= &\quad \frac{1}{2}\sum_{y\in\Omega}|\mu(y) - \nu(y)|\sum_{x\in\Omega}P(y,x) \\
= &\quad \frac{1}{2}\sum_{y\in\Omega}|\mu(y) - \nu(y)| = \|\mu - \nu\|_{TV}.
\end{aligned}$$

□

Corollary 5.2. *Let $\{X_t\}$, P, and Ω as above and let π denote the stationary distribution of $\{X_t\}$. Let μ be an arbitrary initial distribution and let μP^t denote the distribution of X_t. Then:*

$$\|\mu P^t - \pi\|_{TV} \geq \|\mu P^{t+1} - \pi\|_{TV}.$$

Proof. Since μP^t is a valid distribution on Ω, the definition of a Markov chain means that we can treat it as a valid inital distribution and consider the distribution of step $t+1$. Then the result is a direct consequence of Theorem 5.1 and the definition of a stationary distribution.

□

The theorem and the corollary tell us that the total variation distance of the distribution of the Markov chain from the stationary distribution is non-increasing. Furthermore, we already know that the chain asymptotically converges to stationarity. The question now becomes, how long does it take before a Markov chain becomes "close enough" to the stationary distribution? First, in order to get rid of the dependence of the total variation distance of a Markov chain from stationarity on initial distribution, we define the following expressions as functions of t:

$$\begin{aligned}
d(t) = &\quad \max_{\mu}\|\mu P^t - \pi\|_{TV} \\
\overline{d}(t) = &\quad \max_{\mu,\nu}\|\mu P^t - \nu P^t\|_{TV}.
\end{aligned}$$

If we can bound $d(t)$ and $\overline{d}(t)$, we can be assured that the total variation distance of the chain is bounded regardless of the initial distribution. However, it is often easier to work with alternative characterizations.

Lemma 5.3.

$$\begin{aligned}
d(t) = &\quad \max_{x\in\Omega}\|P^t(x,\cdot) - \pi(\cdot)\|_{TV} \\
\overline{d}(t) = &\quad \max_{x,y}\|P^t(x,\cdot) - P^t(y,\cdot)\|_{TV}.
\end{aligned}$$

Proof. For an arbitrary initial distribution μ, we have that:

$$\begin{aligned}
\|\mu P^t - \pi\|_{TV} = & \frac{1}{2}\sum_{x\in\Omega} \|\mu P^t(x) - \pi(x)| \\
= & \frac{1}{2}\sum_{y\in\Omega} |\mu(y)| \sum_{x\in\Omega} |P^t(y,x) - \pi(x)| \\
= & \sum_{y\in\Omega} \mu(y) \cdot \frac{1}{2}\sum_{x\in\Omega} |P^t(y,x) - \pi(x)| \\
= & \sum_{y\in\Omega} \mu(y) \|P^t(x,\cdot) - \pi\|_{TV} \\
= & \|P^t(x,\cdot) - \pi\|_{TV} \\
\leq & \max_{x\in\Omega} \|P^t(x,\cdot) - \pi\|_{TV}.
\end{aligned}$$

But since the distribution which assigns x probability 1 and 0 everywhere else is a valid distribution, the other side of the inequality immediately follows.

The other expression can be derived by a similar argument. □

Lemma 5.4. *$\overline{d}(t)$ is submultiplicative, ie:*

$$\overline{d}(s+t) \leq \overline{d}(s)\overline{d}(t).$$

Proof. Fix $x, y \in \Omega$ and let (X_s, Y_s) be the optimal coupling of $P^s(x,\cdot)$ and $P^s(y,\cdot)$. Therefore, we have that:

$$\|P^s(x,\cdot) - P^s(y,\cdot)\|_{TV} = \mathbb{P}\{X_s \neq Y_s\}.$$

If we consider an entry of the $s+t$ step transition matrix, we have that:

$$\begin{aligned}
P^{s+t}(x,w) = & \sum_{z\in\Omega} P^s(x,z)P^t(z,w) \\
= & \sum_{z\in\Omega} \mathbb{P}\{X_s = z\}P^t(z,w) \\
= & \mathbb{E}P^t(X_s,w).
\end{aligned}$$

Similarly, $P^s(y,w) = \mathbb{E}P^t(Y_s,w)$ and we get that:

$$\begin{aligned}
& P^{s+t}(x,w) - P^{s+t}(y,w) = \mathbb{E}(P^t(X_s,w) - P^t(Y_s,w)). \\
\Rightarrow \|P^{s+t}(x,\cdot) - P^{s+t}(y,\cdot)\|_{TV} = & \frac{1}{2}\sum_{w\in\Omega} |\mathbb{E}(P^t(X_s,w) - P^t(Y_s,w))| \\
\leq & \mathbb{E}\left(\frac{1}{2}\sum_{w\in\Omega} |P^t(X_s,w) - P^t(Y_s,w)|\right) \\
= & \mathbb{E}\|P^t(X_s,\cdot) - P^t(Y_s,\cdot\|_{TV} \\
\leq & \mathbb{E}\overline{d}(t)1_{\{X_s\neq Y_s\}} \\
= & \overline{d}(t)\mathbb{P}\{X_s \neq Y_s\} \\
= & \overline{d}(t)\|P^t(x,\cdot) - P^t(y,\cdot)\|_{TV}.
\end{aligned}$$

This holds for arbitrary x, y initial states, so maximizing the total variation distance over x, y finishes the proof.

□

Definition 5.5. For an ergodic Markov chain $\{X_t\}$ over the state-space Ω with stationary distribution π, the **mixing time** of $\{X_t\}$ is given by:

$$t_{mix}(\varepsilon) = \min_t\{t \,:\, d(t) < \varepsilon\}.$$

A common convention is to set $\varepsilon = \frac{1}{4}$ and denote:

$$t_{mix} = t_{mix}\left(\frac{1}{4}\right).$$

Proposition 5.6.

$$d(lt_{mix}(\varepsilon)) \leq (2\varepsilon)^l \text{ for any non-negative integer } l.$$

Proof.

$$d(lt_{mix}(\varepsilon)) \leq \overline{d}(lt_{mix}(\varepsilon)) \leq (\overline{d}(lt_{mix}(\varepsilon)))^l \leq (2\varepsilon)^l.$$

The first inequality is obvious. The second inequality follows from Lemma 5.4. Lastly, the third inequality comes from the fact that total variation distance is inherently bounded. □

So choosing $\varepsilon = \frac{1}{4}$ gives a nice bound:

$$d(t_{mix}) \leq 2^{-l}.$$

In other words, the total variation distance is exponentially decreasing with steps of size t_{mix}. If we can compute, estimate, or even bound the mixing time, then we can get a pretty good picture of the long-term behavior of the chain (assuming we know about the stationary distribution).

Lastly, we can bound $t_{mix}(\varepsilon)$ for arbitrary $\varepsilon > 0$ in terms of the t_{mix} for $\frac{1}{4}$ as shown in the following lemma.

Lemma 5.7.

$$t_{mix}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil t_{mix}.$$

Proof.

$$\begin{aligned} d(\lceil \log_2 \varepsilon^{-1} \rceil t_{mix}) \leq & \quad \overline{d}(t_{mix})^{\lceil \log_2 \varepsilon^{-1} \rceil} \\ \leq & \quad 2^{-\lceil \log_2 \varepsilon^{-1} \rceil} \\ \leq & \quad 2^{-\log_2 \varepsilon^{-1}} = \varepsilon. \end{aligned}$$

Therefore,

$$t_{mix}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil t_{mix}.$$

□

6. Ant Colony Optimization

Ant Colony Optimization (ACO) is a heuristic inspired by the foraging behavior of actual ants to solve problems that can be reduced to the question of finding an optimal path through a graph. In nature, individual ants randomly wander away from the colony searching for food while laying pheromone down. As more ants leave the colony, they are drawn towards the pheromone trails laid down by the previous ants, thus reinforcing certain paths. Furthermore, pheromones evaporate over time so longer paths will become less attractive over time since they take longer to traverse (and therefore have more time to evaporate).

Algorithms of this sort are attractive for several reasons. First of all, though the algorithms have yet to be understood fully at a theoretical level, there has been empirical evidence of their efficacy in a variety of applications [7]. Furthermore, actual ants are capable of navigating complex environments despite the fact that individual ants have very limited cognitive abilities. The goal of ACO is to replicate the emergent problem solving abilities of a collective group of simple agents. Lastly, the randomized nature of ACO guarantees that the search for an optimum will never strictly converge to a local optimum that is not globally optimal.

7. MMAS Algorithm

We will work with a version of ant colony optimization known as the iteration-best MAX-MIN Ant System (MMAS) applied to a shortest path problem with a single destination [7]. The precise statement of the problem is the following:

Suppose that we have a directed acyclic graph $G(V, E)$. We will call a connected sequence of vertices and edges from a start vertex u_0 to a target vertex v_0 a path. For any vertex $u \in V$, we will denote $deg(u)$ to be the number of edges pointing away from u. We will denote Δ as the maximum degree over all the vertices and $diam(G)$ will be the maximum number of edges in any path. Denote the set of paths from a vertex u to v_0 by $\mathcal{S}_u$. Furthermore, we define:

$$\mathcal{S} = \bigcup_{u \in V} S_u.$$

The distance traversed over a given path $l = (u_0, ..., v_0)$ is given by some function $f : S \to \mathbb{R}$. The goal is to find a solution $x^* \in \mathcal{S}_{u_0}$ such that:

$$f(x^*) = \min_x \{f(x) \mid x \in \mathcal{S}_{u_0}\}.$$

Although f maps from the set of all possible paths starting at all vertices, the actual search space is only the set of paths starting at u_0.

To make things less cluttered, we will assume that the target v_0 can be reached from any vertex u.

Fix $\tau_{min} \in (0, 1)$. Then define an initial function:

$$\phi_0 : E \to [\tau_{min}, 1 - \tau_{min}]$$

This function gives a vector of length $|E|$:

$$(\phi_0(e))_{e \in E}$$

We will refer to this vector as the **pheromone vector**. As the algorithm runs, the pheromone vector will be updated at every iteration by way of updating the function $\phi_t : E \to [\tau_{min}, 1 - \tau_{min}]$ for every time step t.

Then at each iteration, λ independent random walks are initiated from every vertex $u \in V$ towards the destination v_0. If a random walker is currently at vertex u, then there is a finite set of possible outgoing edges from u, say $\{e_1, ..., e_K\}$. Then the next step that the walker takes will be chosen from these edges with the following probabilities:

$$\mathbb{P}\{\text{walker chooses } e_j \mid \text{ walker is at } u\} = \frac{\phi(e_j)}{\sum_{i=1}^{K} \phi(e_i)}.$$

If the walker chooses an edge e_j, then the walker moves to the corresponding vertex. This process continues until the walker reaches v_0. At the end of the walk, the path that was traversed is stored. We will refer to these random walkers as **ants**.

Example 7.1. Consider the following graph:

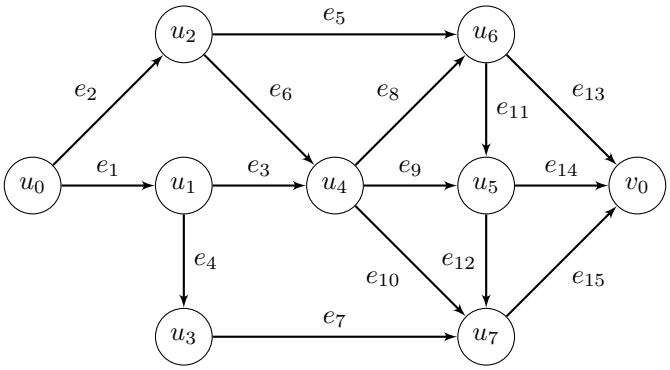

Here u_0 is the start vertex and v_0 is the target. Suppose we pick the vector ϕ_0 such that $\phi(e_i) = 1 - \tau_{min}$ for all $i = 1, ..., 15$. If we start an ant at u_0, then the first step is either along edge e_1 to the vertex u_1 or along e_2 to the vertex u_2. Based on our choice of ϕ, the probability of an ant moving to u_1 is given by:

$$\mathbb{P}\{\text{ant picks } e_1 \mid \text{ ant is at } u_0\} = \frac{\phi_0(e_1)}{\phi_0(e_1) + \phi_0(e_2)} = \frac{1}{2}.$$

If an ant is at u_4, then the probability of choosing u_5 is $\frac{1}{3}$ because the appropriate edge to reach u_5 is e_9. The conditional probability of choosing e_9 while at u_4 is computed by:

$$\mathbb{P}\{\text{ant picks } e_9 \mid \text{ ant is at } u_4\} = \frac{\phi_0(e_9)}{\phi_0(e_8) + \phi_0(e_9) + \phi_0(e_{10})} = \frac{1}{3}$$

After all the ants starting at any vertex u complete their walks, the paths are evaluated according to the function f and the best path is stored. In precise terms, we can denote the path of the j^{th} ant starting from the vertex u as $l_u^{(j)}$, and we store:

$$l_u^* = \underset{l_u^{(j)}}{\operatorname{argmin}} f(l_u^{(j)}).$$

This process is done for every vertex $u \in V$. For ants starting at the intended start vertex u_0, the best path of the iteration is stored separately and compared against the best path of all previous iterations with respect to f.

The next step is the most crucial aspect of the algorithm: the pheromone update. Once we have l_u^* for each u, we can define the set of first edges of all the best paths in the iteration:

$$B = \{e \mid e \text{ is the first edge of } l_u^* \text{ for some } u \in V\}.$$

Then, we update each entry of the pheromone vector in the following way:

$$\phi_{t+1}(e) = \begin{cases} \min\{(1-\rho)\phi_t(e) + \rho, 1 - \tau_{min}\} \text{ if } e \in B \\ \max\{(1-\rho)\phi_t(e), \tau_{min}\} \text{ otherwise} \end{cases}$$

In the expression above, $\rho \in (0, 1)$ is a chosen parameter we will call the **evaporation factor**. The update rule first *evaporates* pheromone by reducing the

pheromone values for all the edges to simulate the actual evaporation of biological pheromones. Then, the first edges of the best paths are *reinforced* to simulate the process of actual ants leaving actual pheromones as they travel. Notice that the value of the pheromone vector on a particular edge can only be affected by the ants leaving from the vertex where the edge originates.

Example 7.2. Consider the graph from 7.1. Assume that ϕ_0 is initiated as before and $\rho \in (0,1)$. Suppose that the best path from u_0 is $l_u^* = (u_0, e_2, u_2, e_7, u_6, e_{13}, v_0)$. Then in the next iteration, the pheromone value on the edge e_2 will be given by:

$$\phi_1(e_2) = (1-\rho)(1-\tau_{min}) + \rho.$$

On the other hand, the pheromone value on the edge e_1 will be given by:

$$\phi_1(e_1) = (1-\rho)(1-\tau_{min}).$$

After the pheromone update is completed, the algorithm moves to the next iteration and the ants make their random walks with respect to the newly updated pheromone vector. Then the iterations continue until some stopping condition is met. The easiest example is simply fixing a number of iterations T.

Formally, the algorithm is given by:

for *iter from 1 to* T **do**
  for u *in* V **do**
    for j *from* 1 *to* λ **do**
      $i = 0$, $p_0^{(j)} \leftarrow u$, $V_1^{(j)} \leftarrow \{p \in V \mid (p_0^{(j)}, p) \in E\}$;
      while $p_i^{(j)} \neq v_0$ *and* $V_{i+1}^{(j)} \neq \emptyset$ **do**
        $i \leftarrow i+1$
        Choose a vertex $p_i^{(j)} \in V_i^{(j)}$ according to the probabilities:
        $\frac{\phi((p_{i-1}^{(j)}, p_i^{(j)}))}{\sum_{p \in V_i^{(j)}} \phi((p_{i-1}^{(j)}, p))}$
        $V_{i+1}^{(j)} \leftarrow \{p \in V \setminus \{p_0^{(j)}, ..., p_i^{(j)}\} \mid (p_i^{(j)}, p^{(j)}) \in E\}$
      end
      current_path $\leftarrow (p_0^{(j)}, ..., p_i^{(j)})$;
      if *f(current_path) < f(best_path)* **then**
        best_path $\leftarrow$ current_path;
      end
    end
    if $u == u_0$ **then**
      if *f(best_path) < f(stored_path)* **then**
        stored_path $\leftarrow$ best_path
      end
    end
  end
  update pheromones
end

The variable *stored_path* simply stores the best-so-far solution over all of the iterations from the desired start point u_0 to the destination v_0.

8. Convergence of MMAS

Since the algorithm puts a lower bound on the amount of pheromone on any given edge of the graph, there is a minimum positive probability of choosing an optimum solution at any given time-step t. Therefore the probability of never observing an optimal solution is zero and we will almost surely find an optimum.

However, this doesn't give any idea of precisely how long this process will take. Knowing that the algorithm will find the optimum in finite steps is unhelpful if the time it takes to reach those steps is longer than it takes to reach the heat death of the universe. Now, we will appeal to the stochastic process induced by the iterations of the algorithm to get a sense of the running time of the algorithm. In particular, if we denote the sequence of random pheromone vectors as $\{Y_t\}$, we will show that $\{Y_t\}$ is an ergodic Markov chain with a countable state space with a slight restriction to the initial pheromone vector. Furthermore, we will set upper bounds on the mixing time of $\{Y_t\}$ in a similar fashion to Sudholt [3].

We will assume the same problem as before with a finite directed acyclic graph $G(V, E)$ with start u_0 and destination v_0 with the function f under the same assumptions as before. We note that for any path $x \in S_{u_0}$, there is a unique set of edges associated with the vertices in the path. For convenience, we will say that an edge, e, connecting adjacent vertices, $u_1, u_2 \in x$ is in the path.

Now we provide some restrictions to the pheromone vectors which will be important in proving the ergodicity of the pheromone vectors. First, we define a vector $\phi^{(x)}$ for any path $x \in \mathcal{S}_{u_0}$ to be a pheromone vector such that:

$$\phi^{(x)}(e) = \begin{cases} 1 - \tau_{min} \text{ if } e \text{ is an edge in the path } x \\ \tau_{min} \quad \text{otherwise} \end{cases}$$

Then, we define:

$$C = \{\phi^{(x)} \mid x \in \mathcal{S}_{u_0}\}.$$

Lastly, we define:

$$\mathcal{V} = \{\phi \in \mathbb{R}^{|E|} \mid \mathbb{P}\{Y_t = \phi, Y_0 = \phi^{(x)}\} > 0 \text{ for some } t > 0 \text{ and for some } x \in S_{u_0}\}.$$

Observe that if we initialize the algorithm in the set $\mathcal{V}$, then we actually have a countably infinite number of possible pheromone vectors that the algorithm can actually achieve since there is only a finite number of new configurations available at each iteration. Now we need to show that Y_t satisfies the Markov property and demonstrate ergodicity.

Lemma 8.1. *Let $\{Y_t\}$ be a sequence of random variables over $\mathcal{V}$ such that Y_t is the vector of pheromone values at time t of the MMAS algorithm. Then $\{Y_t\}$ is a Markov chain.*

Proof. Returning to the definition, we want to show that:

$$\mathbb{P}\{Y_{t+1} = y \mid Y_0 = y_0, ..., Y_t = y_t\} = \mathbb{P}\{Y_{t+1} = y_{t+1} \mid Y_t = y_t\}.$$

But this is as simple as realizing that the random mechanism by which Y_{t+1} is generated if we are at time t is entirely determined by the behavior of the ants randomly walking across the graph. The random walks of the ants are in turn governed strictly by a function of the current pheromone values at time t. □

Lemma 8.2. *The pheromone vector of the MMAS algorithm is irreducible over $\mathcal{V}$.*

Proof. Consider two arbitrary pheromone vectors $\phi^{(1)}, \phi^{(2)} \in \mathcal{V}$. Then by definition we have that:

$$\exists\, t > 0 \text{ such that } P^t(\phi^*, \phi^{(2)}) > 0 \text{ for some } \phi^* \in C.$$

Therefore, we only need to prove that there is a positive probability that an arbitrary $\phi^{(1)}$ will reach any arbitrary $\phi^* \in C$. As discussed earlier, because there is a hard lower bound on the amount of pheromone on any given edge:

$$\exists\, p_{min} > 0 \text{ such that } p_{min} \leq \mathbb{P}\{e \text{ is re-inforced} \,|\, \phi\} \ \forall e \in E \text{ and } \forall\, \phi \in \mathcal{V}.$$

So for any given iteration, we note that the path x associated with ϕ^* has a positive probability of being reinforced while the other edges only evaporate. This guarantees a positive probability of reaching ϕ^* in t steps for some positive integer t. This is because consecutively updating the same edges induces a recursive function which is strictly increasing for the reinforced edges and strictly decreasing for the evaporating edges.

For the re-inforced edges, we observe that after the k^{th} consecutive re-inforcement on an edge e originally at $\phi_0(e)$ is given by:

$$\phi_k(e) = (1-\rho)^k \phi_0(e) + \sum_{i=1}^{k-1} (1-\rho)^i \rho.$$

We observe that:

$$\lim_{k\to\infty} \phi_k(e) = 1.$$

And we have already set a hard maximum of $1 - \tau_{min}$. This means that:

$$\exists\, K > 0 \text{ st } \ 1 - \tau_{min} \leq \phi_k(e) \leq 1, \ \forall\, k \geq K.$$

Then, for each possible value of $\phi_0(e)$ there exists such a K. In the worst case scenario, $\phi_0(e) = \tau_{min}$ but this K is still finite. Therefore after K steps, there is a positive probability that all the edges in x are pushed to the maximum.

On the other hand, if we consider an edge e' originally at $\phi_0(e')$ that is consecutively evaporating for l steps, we can see that the pheromone value at the l^{th} consecutive evaporation is given by:

$$\phi_l(e') = \phi_0(e')(1-\rho)^l.$$

We observe that:

$$\lim_{l\to\infty} \phi_l(e') = 0.$$

And we have already set our minimum pheromone value to be τ_{min}. So if we consider the highest possible value for $\phi_0(e') = 1 - \tau_{min}$, we have that:

$$\exists L > 0 \text{ st } \phi_l(e') \leq \tau_{min}, \ \forall l \geq L.$$

Define the constant M by:

$$M = \max\{L, K\}.$$

After M iterations, there is a positive probability that the chain will be at ϕ^* starting from any vector in $\mathcal{V}$. Even if the chain reaches ϕ^* in less than M steps, there is a positive probability of remaining at ϕ^* because if the same path is reinforced, the bounds τ_{min} and $1 - \tau_{min}$ will force the pheromones to say the same.

□

Lemma 8.3. *The pheromone vector of the MMAS algorithm is aperiodic.*

Proof. For aperiodicity, we can consider the state in which an optimal path $x^* \in S$ has taken all maximal pheromone values and edges not in x^* have taken all minimal pheromone values at time t. Denote this state y^*. Then in time $t+1$, there is a positive probability that all edges of x^* are reinforced. But the bounds on pheromone values mean that this event leaves all edges of x^* at $1-\tau_{min}$ and all other edges at $\tau_{\min}$. Therefore, we have that:

$$gcd\{t \,|\, \mathbb{P}\{Y_t = y^* \,|\, Y_0 = y^*\}\} = 1.$$

Hence, the state y^* is aperiodic. Then since we know that Y_t is irreducible (at least based on our chosen initialization), we can use Lemma 2.6 to conclude that the whole Markov chain is aperiodic. □

Theorem 8.4. *The pheromone vector of the MMAS is ergodic and thus has a stationary distribution.*

Proof. Irreducibility and aperiodicity are shown in the previous two lemmas. Only positive recurrence remains.

Consider an arbitrary path $x^* \in S_{u_0}$ and consider the associated pheromone vector where every edge in x^* has pheromone value $1-\tau_{min}$ and τ_{min} everywhere else. Denote this vector as ϕ^*. Then for any other $\phi \in \mathcal{V}$, our goal is to bound the expected number of steps until ϕ^* is reached again.

First, reconstruct M as in the proof of Lemma 8.2, where M is the minimum number of consecutive iterations of reinforcing only the edges in x^* to reach ϕ^* from any state. Then we observe that the event $\{\phi^*$ is reached within M steps$\}$ has a fixed minimum probability p^* of occuring from any state.

If we consider the geometric random variable Z where a trial is M iterations with probability of success given by p^*, we see that that the expected time it takes to reach ϕ^* from any vector ϕ is bounded above by:

$$\mathbb{E}T_{\phi,\phi^*} \leq M * \mathbb{E}Z = \frac{M}{p^*}.$$

This means that the expected hitting time from any state ϕ is finite. Therefore, the expected return time must be finite and ϕ^* must be positive recurrent. By Lemma 2.13, Lemma 8.2, and Lemma 8.3, we have that the pheromone vector is ergodic. □

Now that we know there is a stationary distribution, the next step is to analyze the mixing time. To do this, we construct the following coupling of two pheromone vector Markov chains, $\{X_t\}$ and $\{Y_t\}$ with pheromone vectors $\phi^{(X)}$ and $\phi^{(Y)}$, respectively. For an arbitrary edge $e \in E$, we can denote the vertex where e originates as u_e. Then if we let $p_X(e)$ and $p_Y(e)$ denote the probability an ant at u_e selects the edge e under ϕ_X and ϕ_Y, then we will set for the j^{th} in each chain:

$$\mathbb{P}\{\text{both ants choose } e\} = \min\{p_X(e), p_Y(e)\}.$$

Then, with the remaining probability, the ants each choose an edge independently based on their respective pheromone vectors, ie:

$$\mathbb{P}\{\text{each ant chooses according to original construction}\} = 1 - \sum_{e} \min\{p_X(e), p_Y(e)\}.$$

If at any time all of the outgoing edges on a vertex u have taking all of the same pheromone values in both $\phi^{(X)}$ and $\phi^{(Y)}$, the ants in either chain will always make the same decision. This is true because if all the edges e outgoing from u have the same pheromone values, then $p_X(e) = p_Y(e)$ and there is zero probability that the ants make different decisions.

Now suppose without loss of generality that the graph G has ordered vertices $V = \{1, ..., N\}$ with respect to the destination N such that for any two vertices $i < j$, the longest path from i to N is at least the same length as the longest path from j to N. Define G_i to be the subgraph of G including the vertices $\{i, ..., N\}$. Once all of the pheromones are the same between X_t and Y_t for every edge in G_i, we say that X_t and Y_t are coupled in G_i.

Example 8.5. The following graph satisfies the ordering described above:

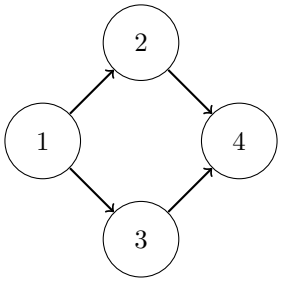

Observe that once X_t and Y_t couple in G_i, the j^{th} ant in each chain will always behave the same once in G_i. Furthermore, since pheromones in G_i are updated only by ants starting in G_i, the pheromone vectors also remained the same between the two chains. This means that we can consider the case where the coupling incrementally reaches the condition in equation 4.5 by first coupling in G_N trivially, then in G_{N-1} and so on.

In order to estimate the coupling time, we consider the time it takes to couple in G_i given that X_t and Y_t have already coupled in G_{i+1}. In particular, we realize that at any vertex u, if the same outgoing edge has taken the maximum pheromone value in both X_t and Y_t, it must be true that the behavior of both chains will be the same when passing through the vertex u. This is a consequence of the following lemma:

Lemma 8.6. *Consider a vertex u with an outgoing edge e satisfying $\phi(e) = 1 - \tau_{min}$, then for all other outgoing edges e', we must have:*

$$\phi(e') = \tau_{min}$$

Proof. We first prove the following inequality:

$$\sum_{e=(u,\cdot)\in E} \phi(e) \leq 1 + (deg(u) - 2)\tau_{min}. \tag{8.7}$$

For $deg(u) = 1$, there is nothing to prove. So first consider the case of $deg(u) = 2$ with outgoing edges e_1 and e_2. Then based on our initialization, we have that:

$$\phi(e_1) + \phi(e_2) \leq 1.$$

In the next iteration, we have that the sum of the updated pheromone values ϕ', is given by:

$$\phi'(e_1) + \phi'(e_2) = (\phi(e_1) + \phi(e_2))(1 - \rho) + \rho \leq (1 - \rho) + \rho = 1.$$

Therefore, if either $\phi(e_1)$ or $\phi(e_2)$ is at the maximum, the other must be at the minimum. This gives that:

$$\phi(e_1) + \phi(e_2) \leq 1 = 1 + (deg(u) - 2)\tau_{min}.$$

For general $deg(u)$, we observe that for any vector $\phi \in \mathcal{V}$, ϕ can be reached in finite updates from some boundary vector $\phi^* \in C$. Therefore, if we prove that the inequality holds for every possible sequence of pheromone updates initialized from any $\phi^* \in C$, we're done. For any such ϕ^*, we observe that the inequality necessarily holds since there is at most one outgoing edge at $1 - \tau_{min}$ and everything else is at τ_{min} for any given vertex.

For a vertex u, label the outgoing edges by order of decreasing pheromone values as $\phi(e_1), ..., \phi(e_{deg(u)})$. Then we maximize the sum if as many of the non-reinforced edges are at τ_{min} (since they won't decrease if they're at the lower boundary). This occurs when:

$$\phi(e_3) = \cdots = \phi(e_{deg(u)}) = \tau_{min}.$$

And either $\phi(e_1)$ or $\phi(e_2)$ is reinforced. But as we discussed, $\phi(e_1) + \phi(e_2) \leq 1$ at initialization. Then our earlier analysis tells us that the updated $\phi'(e_1)+\phi'(e_2) \leq 1$. Therefore, we have that the updated pheromone vector must satisfy:

$$\sum_{j=1}^{deg(u)} \phi(e_j) \leq 1 + (deg(u) - 2)\tau_{min}.$$

So if any outgoing edge e has the maximum pheromone value $1 - \tau_{min}$, it must be true that the sum of the pheromones on the other edges is less that $(deg(u)-1)\tau_{min}$. But there are $deg(u) - 1$ remaining edges and each must have at least τ_{min} as a pheromone value. Hence all the remaining edges must be at exactly τ_{min}.

□

Based on the assumed ordering of the vertices $\{1, ..., N\}$, we have that the coupling time for G_i given that G_{i+1} has coupled is bounded above by the time it takes for X_t and Y_t to both reach the maximum pheromone value on the same outgoing edge from vertex i. Knowing this, we can derive the following result:

Theorem 8.8. *Let $\{(X_t, Y_t)\}$ be the coupled Markov chain of two pheromone vectors for an MMAS algorithm on a graph G with N vertices. Let T_i be the random variable for the time it takes for X_t and Y_t to reach the maximum pheromone value for the same outgoing edge from i, given that X_t and Y_t have coupled in G_{i+1}. Then:*

$$t_{mix} \leq O\left(\sum_{i=1}^{N-1} \mathbb{E}T_i\right) \text{ and } t_{mix} \leq O\left(\max\{\mathbb{E}T_i\} \cdot diam(G) \log_2 N\right).$$

Proof. From our previous analysis, we know that for the full coupling time, T_{couple}:

$$T_{couple} \leq \sum_{i=1}^{N-1} T_i.$$

$$\Rightarrow \ \mathbb{P}\{T_{couple} > t\} \leq \mathbb{P}\left\{\sum_{i=1}^{N-1} T_i > t\right\} \ \forall\, t > 0.$$

Then we apply Proposition 4.3 and Markov's Inequality to obtain:

$$d(t_{mix}) \leq \mathbb{P}\{T_{couple} > t_{mix}\} \leq \mathbb{P}\left\{\sum_{i=1}^{N-1} T_i > t_{mix}\right\} \leq \frac{1}{t_{mix}}\mathbb{E}\sum_{i=1}^{N-1} T_i.$$

$$\Rightarrow \ t_{mix} \leq \frac{1}{d(t_{mix})}\sum_{i=1}^{N-1} \mathbb{E}T_i.$$

This gives the first bound. For the second bound, we first define:

$$L_j = \{u \in V \mid \text{longest path from } u \text{ to } v_0 \text{ has } j \text{ edges}\}.$$

Then if all the vertices in $L_1, ..., L_{j-1}$ have coupled, the time until any given vertex $u \in L_j$ couples has a maximum expectation of $\max\{\mathbb{E}T_i\}$. If we apply Markov's Inequality, we see:

$$\mathbb{P}\{T_u > 2\max\{\mathbb{E}T_i\}\} \leq \frac{\mathbb{E}T_u}{2\max\{\mathbb{E}T_i\}} \leq \frac{1}{2}.$$

This simply means that the probability that the vertex $u \in L_j$ has not coupled after $2\max\{\mathbb{E}T_i\}$ iterations given that all of $L_1, ..., L_{j-1}$ has already coupled is less than $\frac{1}{2}$. Then after $\log_2 N + 1$ periods of $2\max\{\mathbb{E}T_i\}$ iterations, the probability that X_t and Y_t have not coupled at u is bounded by:

$$\mathbb{P}\{T_u > 2(\log_2 N + 1)\max\{\mathbb{E}T_i\}\} \leq 2^{-\log_2 N - 1} = \frac{1}{2N}.$$

Then by Boole's Inequality, we have that:

$$\mathbb{P}\left\{\bigcup_{u\in L_j}\{T_u > 2(\log_2 N + 1)\max\{\mathbb{E}T_i\}\}\right\} \leq \sum_{u\in L_j}\mathbb{P}\{T_u > 2\max(\log_2 N{+}1)\{\mathbb{E}T_i\}\} \leq \frac{1}{2}.$$

This means that probability that X_t and Y_t have coupled in L_j given that they are already coupled completely in $L_1, ..., L_{j-1}$ within $2(\log_2 N{+}1)\max\{\mathbb{E}T_i\}$ iterations is at least $\frac{1}{2}$. So if we let T_{L_j} denote the time until all of L_j is coupled given $L_1, ..., L_{j-1}$ have already done so, we have that:

$$\mathbb{E}T_{L_j} \leq 2 \cdot 2(\log_2 N + 1)\max\{\mathbb{E}T_i\}\}.$$

Then there are $diam(G)$ layers for which this holds, which means that:

$$\sum_{j=1}^{diam(G)} \mathbb{E}T_{L_j} \leq O(\max\{\mathbb{E}T_i\} \cdot diam(G) \cdot \log_2 N).$$

Since we know that:

$$T_{couple} \leq \sum_{j=1}^{diam(G)} T_{L_j}.$$

A similar argument to the one used for the first part of theorem gives:

$$t_{mix} \leq O(\max\{\mathbb{E}T_i\} \cdot diam(G) \cdot \log_2 N).$$

□

Up to now, we've only examined the random aspects of this algorithm, but the algorithm is designed to find fixed optima. We now estimate the optimization time of the algorithm in terms of the mixing time and stationary distribution. The key observation here is that the probability of an ant starting at u_0 to take the optimal path depends only on the pheromone vector after the last iteration's update. Therefore, we can consider the joint distribution of the pheromone vector at the start of time step t and the paths of the λ ants starting from vertex u_0 in time step t.

The sample space we are considering is $\mathcal{V} \times \mathcal{S}_{u_0}^{\lambda}$. Let $\{Y_t\}$ be the Markov chain of the pheromone vectors with stationary distribution π and define for each ant:

$$g(\phi) = \mathbb{P}\{\text{chosen path is optimal} \mid \phi\}.$$

Next, we observe that once the pheromone vector ϕ is fixed at the start of an iteration, the ants perform their random walks independently by construction. Therefore, the ants are conditionally independent given the pheromone vector.

Denote the event that at least one ant finds an optimum solution at time t by A_t. Since the ants are conditionally independent given ϕ, the number of ants starting at u_0 finds an optimum is a binomial random number with parameters λ and $g(\phi)$. Therefore, the marginal probability that at least one ant finds the optimal path at the t^{th} iteration is given by:

$$\mathbb{P}\{A_t\} = \sum_{\phi \in \mathcal{V}} \left(1 - (1 - g(t))^{\lambda}\right) \mathbb{P}\{Y_t = \phi\}.$$

In particular, if Y_t is drawn from the stationary distribution π, we have that:

$$\mathbb{P}\{A_t\} = \sum_{\phi \in \mathcal{V}} \left(1 - (1 - g(t))^{\lambda}\right) \pi(\phi).$$

With these results in mind, we can finally relate the actual optimization time to the inherent randomness of the algorithm. This relationship is described in the following theorem.

Theorem 8.9. *Let $\{Y_t\}$ be the ergodic Markov chain of the pheromone vector of the MMAS algorithm over state space $\mathcal{V}$ with stationary distribution π for a graph $G(V, E)$. Then let $p(\mu)$ denote the probability that an ant starting at u_0 reaches the destination v_0 via an optimal path when the distribution of the pheromones is μ. Then the expected optimization time of the algorithm is at most:*

$$t_{mix} \cdot O(p(\pi)^{-1} \log_2 \delta^{-1}),$$

where $\delta > 0$ is a constant dependent only on the stationary distribution.

Proof. Denote $g(\phi)$ to be the conditional probability that an ant chooses an optimal path given pheromone vector ϕ as above. Consider the set of all ε such that:

$$\left| \sum_{\substack{\phi \in \mathcal{V} \\ \pi(\phi) \geq \epsilon}} \left(1 - (1 - g(\phi))^{\lambda}\right) \pi(\phi) - p(\pi) \right| < \frac{p(\pi)}{2}. \tag{8.10}$$

Next, we define:

$$\delta = \sup\{\varepsilon \,|\, \varepsilon \text{ satisfies Equation 8.10}\}.$$

After $t^* := t_{mix}(\delta/2)$ iterations, we have that:

$$\max_{\phi} |\mathbb{P}\{Y_t = \phi\} - \pi(\phi)| \leq \delta.$$

So if we let A_t denote the event that an optimum is found at time step t, we get the following bound:

$$\mathbb{P}\{A_{t^*}\} \geq \sum_{\substack{\phi \in \mathcal{V} \\ \pi(\phi) \geq \delta}} \left(1 - (1 - g(\phi))^{\lambda}\right) \frac{\pi(\phi)}{2}.$$

This implies that the probability of finding an optimum is at least:

$$\mathbb{P}\{A_{t^*}\} \geq \frac{p(\pi)}{4}.$$

Then if an optimum is not found by t^*, we can reset the argument and wait another t^* steps. We note that since the convergence towards stationarity does not depend on initialization, we can treat whatever pheromone vector Y_{t^*} lands on as the start of an independent run of the Markov chain. Since we bound the probability of an optimum from below after each set of t^* iterations, we can bound the expected number of trials until optimization by:

$$4t^* p(\pi)^{-1} \leq 4t_{mix} \cdot p(\pi)^{-1} \lceil \log_2 \delta^{-1} \rceil.$$

□

The consequence of this theorem is that if we can estimate the mixing time and have some idea of the structure of the stationary distribution of the pheromone vector, we can begin to get a sense of the expected optimization time. Mixing time is only one piece of the puzzle and the analysis of the stationary distribution itself is a non-trivial task.

Acknowledgments. It is a pleasure to thank my mentor Antonio Auffinger for his advice and guidance in writing this paper. I would also like to thank Peter May for organizing the REU.

References

[1] Gregory Lawler. Introduction to Stochastic Processes. Chapman and Hall / CRC 2006.
[2] David A. Levin, Yuval Peres and Elizabeth L. Wilmer. Mixing Times and Markov Chains. http://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf.
[3] Dirk Sudholt. Using Markov-Chain Mixing Time Estimates for the Analysis of Ant Colony Optimization. FOGA'11, 2011.
[4] James Norris. Markov Chains. Cambridge Press 1998.
[5] Kyle Siegrist. Periodicity. http://www.math.uah.edu/stat/markov/Periodicity.html
[6] Thomas Stutzle and Marco Dorigo. A Short Convergence Proof for a Class of Ant Colony Optimization Algorithms. IEEE Transactions on Evolutionary Computation, Vol 6, No. 4, August 2002.
[7] Marco Dorigo, Mauro Biratarri, and Thomas Stutzle: Ant Colony Optimization. IEEE Computational Intelligence Magazine, November 2006.
[8] Thomas Stutzle and Holger Hoos. MAX-MIN Ant System. Future Generation Computer Systems 16, 2000.
[9] John Rice. Mathematical Statistics. 2007.