Math-Stat-491-Fall2014-Notes-IV

Hariharan Narayanan

November 19, 2014

1 Introduction

We will be closely following the book
"Essentials of Stochastic Processes", 2nd Edition, by Richard Durrett,
for the topic 'Finite Discrete time Markov Chains' (FDTM). This note is for giving a sketch of the important proofs. The proofs have a value beyond what is proved - they are an introduction to standard probabilistic techniques.

2 Transient and recurrent

We have so far proved a number of equivalent characteristic properties of transient and recurrent states.

Theorem 2.1. *A state x in a finite Markov chain is transient iff it satisfies the following equivalent conditions:*

1. *The probability ρ_{xx} of returning to x, is less than* 1.
2. *Starting at x the expected number of returns to x, $E_x(N_x)$ is finite.*
3. *In the Markov chain graph there is a directed path from x to some recurrent node y but no return directed path.*
4. *For some state y, the probability of reaching in finite time from x to y is nonzero but the probability of reaching in finite time from y to x is less than* 1.
5. *If $\pi(\cdot)$, is a stationary distribution of the Markov chain $\pi(x) = 0$. (Note: $\pi(x) = 0$ in every stationary distribution, for a transient state. If there is more than one stationary distribution, every recurrent state will have value zero in some of them, but not in all.)*

Sketch of proof:
(1) is the definition. (2) is equivalent to (1) because $E_x(N_x) = \rho_{xx}/(1-\rho_{xx})$.
$(2 \rightarrow 3)$
Let C be the collection of all states reachable from x by a directed path. We then have

$$\Sigma_{y\in C} E_x(N_y) = \Sigma_{y\in C}\Sigma_n^{\infty} p^n(x,y) = \Sigma_n^{\infty}\Sigma_{y\in C} p^n(x,y) = \Sigma_n^{\infty}(1) = \infty.$$

This means at least one state in C is a recurrent state, say z. Any state that you can reach from a recurrent state is itself recurrent. So we cannot reach x from z.

$(3 \to 4)$

Clear.

$(4 \to 1)$

$1-\rho_{xx} =$ (The probability of not returning) $\geq$ (probability associated with a path from x to y) multiplied by $(1-\rho_{yx})$, the probability of not returning to x from y, (which is nonzero). So $\rho_{xx} < 1$.

$(1 \leftrightarrow 5)$

Let $\pi(\cdot)$ be a stationary distribution of the Markov chain. For any state x, $E_y(N_x) = \Sigma_n^\infty p^n(y,x)$. If x is transient we have this expression reducing to a finite value. Then we have

$$\infty > \Sigma_y \pi(y) E_y(N_x) = \Sigma_y \pi(y) \Sigma_n^\infty p^n(y,x) = \Sigma_y \Sigma_n^\infty \pi(y) p^n(y,x) = \Sigma_n^\infty \pi(x).$$

But the extreme right of the above expression is finite only if $\pi(x) = 0$.
If x is recurrent, we can construct a stationary distribution where the probabilites are nonzero only in the irreducible component in which x is present. We then have

$$\infty = \Sigma_y \pi(y) E_y(N_x) = \Sigma_y \pi(y) \Sigma_n^\infty p^n(y,x) = \Sigma_y \Sigma_n^\infty \pi(y) p^n(y,x) = \Sigma_n^\infty \pi(x).$$

This implies $\pi(x) > 0$. Thus when x is not transient there is atleast one stationary distribution which takes nonzero value on x.

Theorem 2.2. *A state x in a Markov chain is recurrent iff it satisfies the following equivalent conditions:*

1. *The probability ρ_{xx} of returning to x, is equal to* 1.
2. *Starting at x the expected number of returns to x, $E_x(N_x)$ is infinite.*
3. *In the Markov chain graph, if there is a directed path from x to some node y there is always a return directed path.*
4. *If for some state y, the probability of reaching in finite time from x to y is nonzero then the probability of reaching in finite time from y to x is nonzero.*
5. *If for some state y, the probability of reaching in finite time from x to y is nonzero then this probability is* 1*, and the probability of reaching from y to x is also* 1.
6. *For atleast one stationary distribution $\pi(\cdot)$, of the Markov chain $\pi(x) > 0$. (If the Markov chain has the unique stationary distribution $\pi(\cdot)$, then $\pi(x) > 0$).*

3 Finite, irreducible, aperiodic Markov chains

The period of a Markov chain at a state x is the gcd of all the loop lengths at x. The Markov chain in Figure 1 has period 2 at each of its states.

When the Markov chain is *aperiodic* at a particular state x, the gcd of all loop lengths at x is 1. A consequence of this is that whenever $n > N$, for a large enough N, we will have a loop of length n at x.

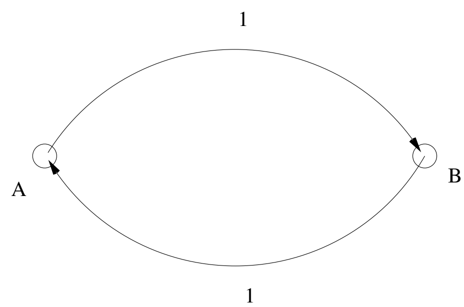

Figure 1: Markov chain with period 2

When the Markov chain is irreducible and finite, the gcd of loop lengths at any state x, is the same as that at any other state y.
For a finite irreducible Markov chain, directed paths exist from any state to any other. This fact combined with the above is sufficient to show that whenever $n > N$, for a large enough N, we will have a path length n from x to any other state y.
By using Chapman- Kolmogorov Theorem, it follows that $p^n(i, j) > 0$, whenever $n > N$, for a sufficiently large N. This means that the matrix P^n is fully positive whenever $n > N$, for a sufficiently large N.

4 Limit Theorems

We have already seen that every irreducible Markov chain has all its states recurrent and further has a unique stationary distribution. We remind the reader that $\rho_{xy} \equiv Pr_x\{T_y < \infty\} \equiv$ the probability of reaching from x to y in finite time. The following is a useful characterization of irreducible Markov chains.

Lemma 4.1. *A Markov chain is irreducible*

- *iff* $\rho_{xy} > 0$, *for every pair of states* x, y
- *iff* $\rho_{xy} = 1$, *for every pair of states* x, y.

Proof of the Lemma:
The first part is a restatement of the definition of an irreducible Markov chain.
For the second, the 'if' part is clear.
We now prove the 'only if' part. All the states in the irreducible Markov chain are recurrent. Therefore $\rho_{yy} = 1$, for every state y. Further there is a directed path from x to y for each pair of states x, y. Suppose there is a path of length n from y to x. Starting from y the probability of not returning to y is greater than or equal to $p^n(y, x) \times (1 - \rho_{xy})$. But we know this probability is zero since y is recurrent. We conclude that $\rho_{xy} = 1$.

This completes the proof of the lemma.

One of the significant results about aperiodic, irreducible Markov chains is that starting from any initial distribution and running the Markov chain we will eventually converge to the (unique) stationary distribution. Formally,

Theorem 4.1. *Let the Markov chain be irreducible and aperiodic. Let $\pi'(\cdot)$ be an initial distribution of the states. Then ${\pi'}^T P^n$ converges to the unique stationary distribution as $n \to \infty$.*

However if the Markov chain has a period greater than 1, the limit does not exist. A good example is the Markov chain in Figure 1. It is clear that this chain has the transition matrix

$$P = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{1}$$

Suppose we take $\pi^T = (0, 1)$. Then $\pi^T P^n$ would yield $(1, 0)$ for n odd and $(0, 1)$ for n even and therefore there would be no convergence as $n \to \infty$.

Another important result about irreducible Markov chains is that if you run such a Markov chain indefinitely, the frequency of encountering the different states (i.e., the number of times a state is encountered divided by the total time n) converges to the unique stationary distribution. This happens even if the chain has a period greater than 1. For simplicity we prove a weaker version - expected value of the quantity in question converges rather than the quantity itself- of this result for the restricted case of aperiodic irreducible Markov chains.
Let the Markov chain be irreducible and aperiodic and let $N_n(y)$ denote the number of visits to y at times $\leq n$.

Theorem 4.2. *$E(N_n(y))/n$ converges to $\pi(y)$, as $n \to \infty$. If the initial state is picked according to the stationary distribution $\pi(\cdot)$, then $E_\pi(N_n(y))/n = \pi(y)$.*

An immediate consequence is the following.

Theorem 4.3. *Let $f : \{states\ of\ Markov\ chain\} \to \Re$. Then $E(f(y)N_n(y))/n$ converges to $f(y)\pi(y)$, as $n \to \infty$.*

Remark
Stronger versions of these theorems are actually true. $N_n(y)/n$ also converges to $\pi(y)$, and $f(y)N_n(y)/n$ converges to $f(y)\pi(y)$. Proofs of these stronger theorems need sophisticated techniques and exploit the fact that times between successive visits to a state y by the Markov chain are *iid* random variables.

5 Proofs

5.1 Proofs of Theorems 4.2 and 4.3

We will assume Theorem 4.1 to begin with and prove Theorems 4.2 and 4.3. We need the following lemma for which we omit the proof.

Lemma 5.1. *Let $a_0, a_1, \cdots$, be a sequence that converges to a. Then $\lim_{n\to\infty}(1/n)\Sigma_{k=0}^{n-1} a_k$ also equals a.*

Let the Markov chain be aperiodic and irreducible. We have

$$\lim_{n\to\infty} {\pi'}^T(P^n) = \pi^T,$$

where π'^T is an arbitrary distribution on the states and π^T, is the unique stationary distribution. In particular we could have taken π'^T as the vectors $(1, 0, \cdots, 0), \cdots, (0, \cdots, 1)$. Therefore we conclude that every row of P^n tends to π^T, or, equivalently,

$$\lim_{n\to\infty} p_{ij}^n = \pi(j).$$

From the lemma we know that
$\lim_{n\to\infty}(1/n)\Sigma_{k=0}^{n-1} p_{ij}^k = \lim_{n\to\infty} p_{ij}^n = \pi(j)$.
Now p_{ij}^n is the probability of visiting the state j at the n^{th} step starting at state i at time 0. Therefore the mean fraction of visits $E(N_n(j))/n$ to state j in the interval $[0, n-1]$ is $(1/n)\Sigma_{k=0}^{n-1} p_{ij}^k$. The long run mean weighted fraction of time that the process spends in state j is $\lim_{n\to\infty}(E(N_n(j))/n) = \lim_{n\to\infty}(1/n)\Sigma_{k=0}^{n-1} p_{ij}^k = \pi(j)$.
If we pick the initial state according to the stationary distribution $\pi(\cdot)$, the subsequent state distributions remain $\pi(\cdot)$. The expected value of $N_n(y) = \pi(y)$, when $n = 1$. So we have $E_\pi(N_n(y)) = n\pi(y)$, by using the linearity of expectation. (Let $X_i'(y) = 1$, if $X_i = y$, and equal to zero if $X_i \neq y$. When the initial distribution is picked according to $\pi(\cdot)$, $Pr\{X_i'(y) = 1\}$ is $\pi(y)$. Now $N_n(y) = X_1'(y) + \cdots + X_n'(y)$. So $E_\pi N_n(y) = E_\pi(X_1'(y) + \cdots + X_n'(y)) = E_\pi(X_1'(y)) + \cdots + E_\pi(X_n'(y)) = n\pi(y)$.)

Next let us suppose that each visit to state j incurs a cost $f(j)$. Let $M_n(j) \equiv f(j) \times N_n(j)$ denote the 'weighted' number of visits to state j in the interval $[0, n-1]$. Therefore the mean weighted fraction of visits $E(M_n(j))/n$ to state j in the interval $[0, n-1]$ is $(1/n)\Sigma_{k=0}^{n-1} f(j)p_{ij}^k$. The long run mean weighted fraction of visits to state j is $\lim_{n\to\infty} E(M_n(j))/n = \lim_{n\to\infty}(1/n)\Sigma_{k=0}^{n-1} f(j)p_{ij}^k = f(j)\pi(j)$.

This completes the proof of Theorems 4.2 and 4.3.

5.2 Proof of Theorem 4.1

We will first sketch the proof.
Suppose we have two identical versions of the Markov chain in question, pick the initial states according to different distributions (in particular, we could have started at different initial states) and run them. After some time T the two Markov chains might hit the same state. After that the probabilities that they hit any particular state at any time n have to be the same. Therefore the difference in the probability distributions at time n is connected to the probability that the two Markov chains have not yet hit the same state at a time preceding n. We reach this far through the first two 'claims' below.

Subsequently, we show that if the original Markov chain is aperiodic and irreducible, then the probability that the two Markov chains have not yet hit the same state, at a time preceding n, tends to zero as n tends to ∞. This is done through a trick.

We build a 'product Markov chain' whose states are ordered pairs (x, y) of states of the original Markov chain and whose transition probability from (x_1, y_1) to (x_2, y_2) is $p(x_1, x_2) \times p(x_2, y_2)$, where $p(i, j)$ is the transition probability from i to j in the original Markov chain. We show that this product Markov chain is irreducible and aperiodic if the original Markov chain is irreducible and aperiodic.

Now starting from two different states x_1, x_2 in the two identical copies is the same as starting from (x_1, x_2) in the product chain. The product does have states of the kind (x, x) and, because of irreducibility, the probability of reaching from (x_1, x_2) to (x, x) in finite time is 1. This is the same as saying that the probability that the two identical Markov chains have not hit the same state before time n tends to zero as n tends to ∞.

Claim 1

Let X_n, Y_n be two identical Markov chains. Let T be the earliest time when $X_n = y_n$. Then

$$Pr\{X_n = y, T \leq n\} = Pr\{Y_n = y, T \leq n\}.$$

Proof of Claim 1 :

We have

$Pr\{X_n = y, T \leq n\}$

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{X_n = y, X_m = x, T = m\}$

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{X_n = y \mid X_m = x, T = m\} Pr\{X_m = x, T = m\}$

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{X_n = y \mid X_m = x\} Pr\{X_m = x, T = m\}$, by strong Markov property,

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{Y_n = y \mid Y_m = x\} Pr\{Y_m = x, T = m\}$

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{Y_n = y \mid Y_m = x, T = m\} Pr\{Y_m = x, T = m\}$

$= \Sigma_{m=1}^{n} \Sigma_x Pr\{Y_n = y, Y_m = x, T = m\}$

$= Pr\{Y_n = y, T \leq n\}$.

This proves Claim 1.

Claim 2

$\Sigma_y |Pr\{X_n = y\} - Pr\{Y_n = y\}| \leq 2Pr\{T > n\}$.

Proof of Claim 2 :

Since the distributions X_n, Y_n agree on $\{T \leq n\}$, we have for each y,

$$|Pr\{X_n = y\} - Pr\{Y_n = y\}| \leq Pr\{X_n = y, T > n\} + Pr\{Y_n = y, T > n\}.$$

Summing this over all y gives

$$\Sigma_y |Pr\{X_n = y\} - Pr\{Y_n = y\}| \leq 2Pr\{T > n\}.$$

This proves Claim 2.

Let p, p', denote the transition probabilities for the original Markov chain and product Markov chains, respectively.

Claim 3

If the original Markov chain is irreducible and aperiodic the product Markov chain is irreducible and aperiodic, i.e., for any pair of states $(x_1, x_2), (y_1, y_2)$, there is a directed path in the Markov chain graph from (x_1, x_2) to (y_1, y_2), i.e.,${p'}^n((x_1, x_2), (y_1, y_2))$, is nonzero if $n \geq N$, for sufficiently large N.

Proof of Claim 3 :

Let $(x_1, y_1), (x_2, y_2)$ be states in the product chain. In the original Markov chain whenever $n \geq N$, where

N is sufficiently large we have $p^n(i,j) > 0$, for every pair of states i,j (including (i,i)). Thus there is a path of length n from x_1 to x_2 as well as a path of length n from y_1 to y_2. In the product chain there is therefore a path of length n whenever $n \geq N$, from (x_1, y_1) to (x_2, y_2). If the states in the n length paths in the original Markov chain are $z_1 = x_1, \cdots, z_n = x_2$ and $z_1 = y_1, \cdots, z_n = y_2$, in the product chain the states in the n length paths are $z'_1 = (x_1, y_1), \cdots, z'_n = (x_2, y_2)$. Thus the product chain is irreducible and aperiodic.
This proves Claim 3.

Claim 4
Let the two Markov chains start at states x_1, y_1 respectively. Let T be the first time the two Markov chains are at the same state. Then $Pr\{T < \infty\} = 1$, i.e., $Pr\{T > n\} \to 0$, as $n \to \infty$.

Proof of Claim 4 :
For any state x_i, because of the irreducibility of the product Markov chain we have $Pr\{T' < \infty\} = 1$, where T' is the time to visit (x_i, x_i) from x_1, y_1 for the first time. By definition, $T \equiv min_{x_i}\{T'\}$. Clearly $Pr\{T < \infty\} = 1$. It follows therefore that if the two Markov chains start at states x_1, y_1 respectively, and if T be the first time the two Markov chains are at the same state, then $Pr\{T > n\} \to 0$, as $n \to \infty$.
This proves Claim 4.

By Claim 2, we have
$\Sigma_y |Pr\{X_n = y\} - Pr\{Y_n = y\}| \leq 2Pr\{T > n\}$. Now let us start the first Markov chain at some state x, and let the second state be picked according to the stationary distribution $\pi(\cdot)$. Subsequently the probability of the second chain being at y is $\pi(y)$. It follows that $\Sigma_y |p^n(x,y) - \pi(y)| \leq 2Pr\{T > n\} \to 0$ as $n \to \infty$. This proves Theorem 4.1.

6 Exit distributions and times

Consider a Markov chain with transient and recurrent states. We know that the nodes of the Markov chain graph can be partitioned into $T, R_1, \cdots, R_k$, where T is the set of transient states and the R_i, equivalence classes of recurrent states within each of which there are directed paths from any state to any other state. Two natural questions arise in connection with these chains:

- For a specified transient state x, what is the probability, if one starts from it, of hitting one of the R_i, say R_1?
  Example: In the 'gambler's ruin', if we start with k dollars, what is the probability of winning (i.e., reaching N) or losing (i.e., reaching 0)?

- For a specified transient state x, what is the expected time, if one starts from it, for hitting one of the R_i? (Here hitting implies reaching the first state in $\bigcup R_i$.)
  Example:In the 'gambler's ruin', if we start with k dollars, what is the expected time to reach N or 0, i.e., the expected time of quitting?

6.1 Probability of hitting R from x

For our purpose of computing the hitting probability from transient states, it is convenient to visualize Markov chains as having transient states T and absorbing states $R_1, \cdots, R_k$. This could be done by merging all the states in R_i into a single node R_i, with all the edges entering R_i carrying the same probability as before.

Let i be a state in T and the absorbing state of interest be R. Let P be the transition matrix. Let the probability of reaching R in n steps be denoted $p^n(i, R)$, and probability of reaching R be denoted $q(i, R)$, We then have

$$q(i, R) = \Sigma_{n=1}^{\infty} p^n(i, R),$$

$$p^n(i, R) = \Sigma_{j \in T} P(i, j) p^{n-1}(j, R),$$

$$q(i, R) = p(i, R) + \Sigma_{n=2}^{\infty} p^n(i, R) = p(i, R) + \Sigma_{n=2}^{\infty} \Sigma_{j \in T} P(i, j) p^{n-1}(j, R)$$

$$= p(i, R) + \Sigma_{j \in T} P(i, j) \Sigma_{n=2}^{\infty} p^{n-1}(j, R)$$

$$= p(i, R) + \Sigma_{j \in T} P(i, j) q(j, R).$$

We thus have the set of $|T|$ linear equations in $|T|$ unknowns $q(i, R), i \in T$

$$q(i, R) = p(i, R) + \Sigma_{j \in T} P(i, j) q(j, R).$$

The quantities which are known in these equations are the $p(i, R)$. The coefficient matrix has only nonnegative entries $P(i, j)$. We need to show that these equations have a unique solution for any specified set of values of $p(i, R)$.

Now we know that a set of linear equations $Ax = b$, with A square, has a unique solution iff A is invertible or, equivalently, the only solution to $Ax = 0$, is $x = 0$.

Consider the equations

$$q(i, R) = \Sigma_{j \in T} P(i, j) q(j, R).$$

We have $\Sigma_{j \in T} P(i, j) \leq 1$, since the rows of P add up to 1 and the columns corresponding to T are not the full set (which includes R_i too). Let $k \in T$ be such that $q(k, R)$ has maximum magnitude and this value is not 0. Since $q(k, R) = \Sigma_{j \in T} P(k, j) q(j, R)$, the row sum $\Sigma_{j \in T} P(k, j) \leq 1$, and $q(k, R)$ has maximum magnitude, this can happen only if (1) $\Sigma_{j \in T} P(k, j) = 1$ and (2) if j is such that $P(k, j)$ is not zero, then $q(j, R)$ has the same sign and magnitude as $q(k, R)$. Repeating this argument we see that if there is a directed path from k to any state m in T, then we must have $q(m, R)$ of the same sign and magnitude as $q(k, R)$. The node k is connected by a directed path to one of the absorbing states,say R_i (see Theorem 2.1, part 3). Let the last node in this path which is also in T be s. In the matrix P, the entry $P(s, R_i) > 0$. So $\Sigma_{j \in T} P(s, j) < 1$. Now $q(s, R)$ has maximum magnitude but $\Sigma_{j \in T} P(s, j) < 1$. This contradiction can be avoided only if $q(s, R) = 0$.

We conclude that the set of equations

$$q(i, R) = p(i, R) + \Sigma_{j \in T} P(i, j) q(j, R),$$

has a unique solution whatever may the value of $p(i, R)$ be.

The above is the general method for calculating the probability of hitting R starting from a given state $i \in T$. But in special situations such as the gambler's ruin problem we can get explicit expressions for this probability.

6.2 Expected time of hitting R from x

The analysis of this subsection is valid for the case where the states of the Markov chain can be partitioned into T, R, where R is the single irreducible component. Equivalently, if the Markov chain is of the form $T, R_1, \cdots R_k$, this could be viewed as computing the expected time to reach some unspecified recurrent state from a given transient state. (By Theorem 2.1, from any transient state we can always reach atleast one recurrent state.) Formally if the Markov chain is of the form $T, R_1, \cdots R_k$, we could merge all the R_i into a single absorbing state R, leaving the probabilities of the edges entering R_i unchanged.

Let l_x be the random variable denoting the number of steps from $x \in T$ to hit R. From our discussion of the previous subsection, it is clear that $Pr\{l_x = n\} = p^n(x, R)$. Let $i \in T$. We have

$$p^n(i, R) = \Sigma_{j \in T} P(i, j) p^{n-1}(j, R), n > 1.$$

We note that for $n = 1$, the LHS is simply $p(i, R)$. Therefore to compute the expected value $E(l_i)$ we multiply both sides of this equation by n and sum from 2 to ∞, and then add the term $p(i, R)$. We have, for $n > 1$,

$$n \times p^n(i, R) = n \times (\Sigma_{j \in T} P(i, j) p^{n-1}(j, R)).$$

$$= (\Sigma_{j \in T} P(i, j) p^{n-1}(j, R)) + (n-1) \times (\Sigma_{j \in T} P(i, j) p^{n-1}(j, R)).$$

Summing from 2 to ∞, we get,

$$\Sigma_{n=2}^{\infty} n \times p^n(i, R) = (\Sigma_{j \in T} P(i, j) \Sigma_{n=2}^{\infty} p^{n-1}(j, R)) + (\Sigma_{j \in T} P(i, j) \Sigma_{n=2}^{\infty} (n-1) \times p^{n-1}(j, R))$$

$$= \Sigma_{j \in T} P(i, j) q(j, R) + \Sigma_{j \in T} P(i, j) E(l_j).$$

We remind the reader that $q(j, R)$ is the probability of hitting R starting from j. Now

$$E(l_i) = p(i, R) + \Sigma_{n=2}^{\infty} n \times p^n(i, R) = p(i, R) + \Sigma_{j \in T} P(i, j) q(j, R) + \Sigma_{j \in T} P(i, j) E(l_j),$$

$$= q(i, R) + \Sigma_{j \in T} P(i, j) E(l_j),$$

since we know $q(i, R) = p(i, R) + \Sigma_{j \in T} P(i, j) q(j, R)$. In the present situation we have only one (merged) recurrent state. It is therefore reachable from every transient state (see part 3 of Theorem 2.1). This means that for every transient state i, we must have $q(i, R) = 1$.

We thus have a set of $|T|$ equations in the $|T|$ unknowns $E(l_i)$. The coefficient matrix for the present set of linear equations is the same as the one we obtained for computing $q(i, R)$ in the previous subsection and we have already proved there that it is invertible.