EE 550: Notes on Capacity and Connectivity in Large Networks

Michael J. Neely
University of Southern California
http://www-bcf.usc.edu/~mjneely

I. Introduction

These notes discuss capacity and connectivity in large networks where many wireless devices are randomly located in a unit area. Cases of *static networks* and *mobile networks* are both of interest. Fundamental capacity scaling laws for static networks are developed in [1]. The work [2] shows that capacity scaling laws are improved when mobility is introduced, hence, *mobility increases capacity*. Unfortunately, that increase in capacity can only be achieved with a corresponding increase in network delay. The work [3] introduces a simple cell partitioned model for which exact network capacity can be computed.

A. Cell partitioned networks

Consider a network of wireless devices that are randomly placed in a unit area. Geographically divide the area into C smaller regions called *cells*, where C is a given positive integer. The shape of each cell is arbitrary: It is only assumed that cell regions are disjoint and each cell has the same area. Thus, each cell has area $1/C$, so that the sum area is 1. Two example cell structures are shown in Fig 1. Here, the cells are used only to represent different geographic locations of the network. In particular, we do not assume there is a basestation in each cell. The wireless devices must communicate with each other using some (possibly ad hoc) communication technique.

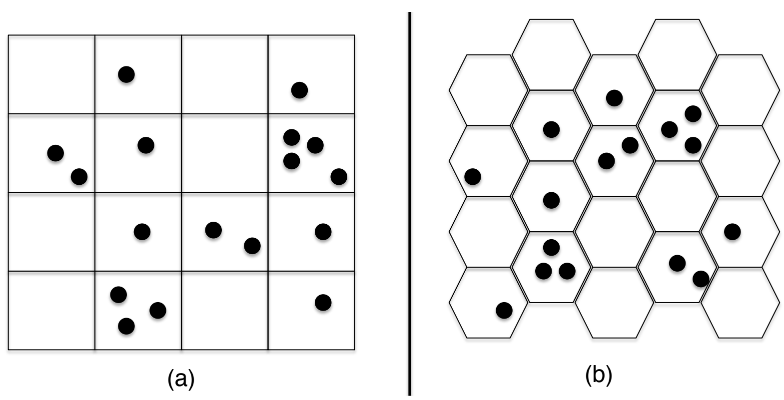

Fig. 1. Wireless devices randomly placed in a cell partitioned network with: (a) a square cell structure; (b) a hexagon cell structure.

B. Two standard random placement models

There are two standard models for random wireless device placement. The first is when we assume there are N devices, where N is a given positive integer, and these are independently and uniformly placed over the network. The second is the *Poisson point process* model where the number of devices in each disjoint cell is an independent Poisson random variable. For convenience, the first model shall be discussed in the context of *mobile networks* and the second shall be discussed in the context of *static networks*.

II. Large mobile networks

Consider a cell partitioned network with C cells and N wireless users, where N and C are given positive integers. Define:

$$C = \text{Number of cells}$$
$$N = \text{Number of users}$$
$$d = \frac{N}{C} = \text{user/cell density}$$

Suppose the system operates over time slots $t \in \{0, 1, 2, \ldots\}$. Users stay in their current cell for the full duration of a time slot, but can move to a new cell at the end of each slot. It is assumed that user mobility is *independent* (so the current location of user i is independent of the current location of user j whenever $i \neq j$). Further, it is assumed that the mobility process for each user has a *steady state distribution* that is uniform over all C cells. For convenience, assume the mobility process starts out in steady state on the first slot (slot $t = 0$). Hence, for user 1 on slot 0, we have:

$$P[\text{user 1 is in cell } i \text{ on slot } 0] = 1/C \quad \forall i \in \{1, \ldots, C\}$$

The assumption that the system starts in steady state at time slot 0 means that the probabilities associated with this slot are also the steady state probabilities that describe long term average behavior in the system.

A. *Simulating a mobility process with uniform location distribution*

The simplest mobility process is the *i.i.d. mobility model* that assumes each user independently selects a new cell from one of the C cells at the end of every time slot. It is clear that on every slot t, the location distribution of each user is uniform over all C cells. Under this model, the square structure or hexagonal structure of the cells (as in Fig. 1) is irrelevant. However, this mobility model may not be realistic, since it allows a user to jump from one corner of the network to an opposite corner on just one timeslot. That is, under the i.i.d. mobility model, user locations have *no memory* from one slot to the next. This can be used as a simplified analytical model and/or as an extreme case of rapid mobility.

A more realistic model is a *Markov random walk model* where users can decide to either stay in their current cell, or move to a neighboring cell. For the square cell structure of Fig. 1a, each interior cell has 4 neighbors (North, South, East, West), and so there are 4 choices for movement. For the hexagonal cell structure of Fig. 1b, each interior cell has 6 neighbors and so there are 6 choices of movement. In both cases, the edge cells have a fewer number of neighbors. However, it is convenient to assume these edge cells have *the same number of movement options* (4 for the square structure, 6 for the hexagon structure), but that users stay in the same cell if they decide to move in an infeasible direction. In particular, under the square cell structure, if a user is already on the West edge of the network and decides to move West, it stays in its same location. A random walk model in this scenario can use a parameter $\beta \in (0, 1]$ to describe the probability of an attempted move. Specifically:

- Every slot $t \in \{0, 1, 2, \ldots\}$, each user independently decides to either stay in its current cell (with probability $1 - \beta$) or to attempt to move to a neighboring cell. It is assumed that all cells have the same number of "move" options, some of which may be infeasible.
- If a user decides to attempt to move to a neighboring cell, it independently decides to move in one of the (fixed number) of directions, equally likely over all directions. If the user decides to move in a feasible direction, it moves to that cell and is located there on slot $t+1$. If it decides to move in an infeasible direction (such as moving West when it is already on the West edge), it stays in its current cell.

Using the theory of *reversibility* for discrete time Markov chains, it can be shown that this Markov random walk results in a steady state location that is uniform over all C cells.[1]

B. *Interference model*

Assume that data is segmented into *packets* and each packet requires 1 time slot to transmit. Suppose that the transmission radius of the device is controlled so that, every timeslot $t \in \{0, 1, 2, \ldots\}$, each user can transmit only to other users in its current cell. Thus, inter-cell communication is not allowed.

To avoid interference between users in the same cell, it is assumed that at most one user can transmit a packet per cell per slot. If there are multiple users in a given cell, the choice of *which* user gets to transmit, and *which* packet it transmits, is done by some control algorithm that coordinates behavior in each cell.

For simplicity, it is assumed that there is no inter-cell interference. Hence, transmission decisions can be made separately over each cell, and transmissions in different cells do not interfere or collide. By controlling the transmission power, it is reasonable to assume that transmissions in two cells that are far away from each other do not interfere with each other. However, transmissions between neighboring cells would naturally interfere. Hence, the assumption of no inter-cell interference would require each neighboring cell to use an orthogonal signaling scheme, such as using different frequency bands in each cell. This is related to a *coloring* problem: For example, it is known that we can use only three colors to color each cell of the hexagonal structure of Fig. 1b while ensuring that no adjacent cells have the same color.

Under this model, a transmission can take place in a given cell $i \in \{1, \ldots, C\}$ if and only if there are at least two users in that cell.

[1]This shall be illuminated later in the course, but for the context of *continuous time Markov chains* rather than discrete time Markov chains. For now, students are free to simulate the above random walk to verify that steady state is uniform over all cells.

C. Number of transmissions

We say that a *transmission opportunity* arises in cell i on slot t if that cell has two or more users on slot t.[2] For simplicity, we focus attention on slot $t = 0$ (since the system is in steady state at time $t = 0$, the analysis for slot 0 will also hold at all times $t \in \{0, 1, 2, \ldots\}$). Let N_i be the number of users in cell i on slot 0, defined for all $i \in \{1, \ldots, C\}$. Recall that, on slot 0, each user is independently and uniformly located over each of the C cells. Hence, for all $i \in \{1, \ldots, C\}$ we have:

$$P[N_i = k] = \binom{N}{k} (1/C)^k (1 - 1/C)^{N-k} \quad , \forall k \in \{0, 1, \ldots, N\}$$

For each $i \in \{1, \ldots, C\}$, define T_i as an indicator function that is 1 if cell i has a transmission opportunity on slot $t = 0$ (and 0 else):

$$T_i = \begin{cases} 1 & \text{if } N_i \geq 2 \\ 0 & \text{otherwise} \end{cases}$$

Let T be the total number of transmission opportunities on slot $t = 0$:

$$T = \sum_{i=1}^{C} T_i$$

Hence

$$\mathbb{E}[T] = \sum_{i=1}^{C} \mathbb{E}[T_i] = \sum_{i=1}^{C} P[T_i = 1] = CP[T_1 = 1]$$

where the final equality holds because all users are equally likely to be in each cell, so $P[T_i = 1] = P[T_1 = 1]$ for all $i \in \{1, \ldots, C\}$. Define $p = P[T_1 = 1]$. Then,

$$\begin{aligned} p &= P[T_1 = 1] \\ &= P[N_1 \geq 2] \\ &= 1 - (1 - 1/C)^N - N(1/C)(1 - 1/C)^{N-1} \\ &= 1 - (1 - d/N)^N - d(1 - d/N)^{N-1} \end{aligned}$$

Hence

$$\boxed{\mathbb{E}[T] = Cp = C\left[1 - (1 - d/N)^N - d(1 - d/N)^{N-1}\right]}$$

Define $\theta_N(d)$ as the *expected number of transmission opportunties per user* as a function of parameters N, d:

$$\theta_N(d) = \frac{\mathbb{E}[T]}{N} = \frac{1 - (1 - d/N)^N - d(1 - d/N)^{N-1}}{d}$$

Now consider increasing both N and C to infinity while keeping the ratio $d = N/C$ fixed. Let $\theta(d)$ be the limiting value:

$$\boxed{\theta(d) = \lim_{N \to \infty} \theta_N(d) = \lim_{N \to \infty} \frac{1 - (1 - d/N)^N - d(1 - d/N)^{N-1}}{d} = \frac{1 - e^{-d} - de^{-d}}{d}}$$

D. Optimizing the user/cell density

It is useful to optimize the user/cell density parameter d to maximize the number of expected transmission opportunities per user per slot. This reduces to maximizing the value of $\theta(d)$ over all $d > 0$. A plot of the $\theta(d)$ function versus d is given in Fig. 2. Notice that $\theta(d)$ tends to zero as $d \to 0$ and also as $d \to \infty$. This is intuitive: When d is very small, the number of users is much smaller than the number of cells. Thus, there are many cells that are empty and many that have only 1 user, so it is unlikely for a cell to have a transmission opportunity. In the opposite extreme when d is large, there are many users packed into each cell, at most 1 of which can transmit. Thus, each particular user must compete with many other users in its cell, and so the expected number of transmission per user is very small. The optimal value of d^* is

$$\begin{aligned} d^* &\approx 1.79328 \\ \theta^* &\approx 0.298426 \end{aligned}$$

The intuition is that, in this case, it is optimal to design the cell size so that there are approximately two users (on average) in each cell.

[2] One situation where all transmission opportunities are taken is as follows: Consider a situation where users move and opportunistically grab packets from popular files of other users. For simplicity, assume that whenever two users meet, each one has a packet that the other user wants. For example, user 1 may want a popular youtube video, while user 2 may have downloaded it yesterday and still has it cached. Assume the problem of choosing which user in a cell transmits to which other user in that cell is decided in some particular way. Then a packet is transmitted in a cell whenever that cell has at least 2 users.

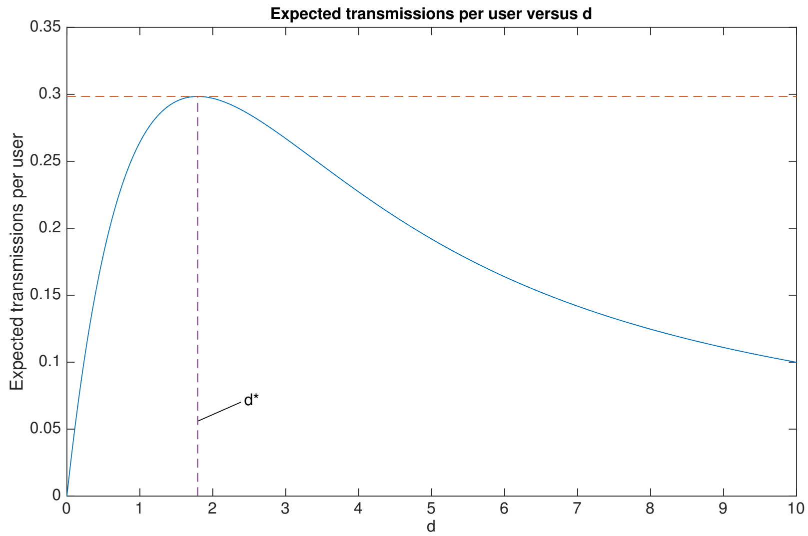

Fig. 2. A plot of $\theta(d)$ versus d.

E. Relation to network capacity for mobile networks

The value $\mathbb{E}\left[T\right]/N$ computed in the previous subsection turns out to be important for characterizing the fundamental *multi-hop network capacity* in this mobile network scenario. In particular, assume each user wants to send an *infinite stream* of packets to some other user with a streaming rate λ packets/slot. For simplicity, assume N is even and suppose users are paired so that

$$\begin{aligned} 1 &\leftrightarrow 2 \\ 3 &\leftrightarrow 4 \\ &\cdots \\ (N-1) &\leftrightarrow N \end{aligned}$$

so that user 1 wants to send a stream of packets to user 2; user 2 wants to send a stream of packets to user 1; user 3 wants to send a stream of packets to user 4; user 4 wants to send a stream of packets to user 3, and so on. The packets can be relayed from user to user over multi-hop paths before reaching their destinations. Every slot t, decisions must be made in each cell about which user gets to transmit, who it transmits to, and which particular packet it sends. This implicitly includes some type of *routing protocol* to decide how packets are routed to relays, which relays are used, and how many relays are used for each path. Some scheduling and routing algorithms may be able to support larger λ rates than others. The work in [3] shows that the *network capacity* μ, being the supremum value of λ that can be supported (considering all possible scheduling and routing algorithms) is given by:

$$\mu = \frac{p+q}{2}$$

where p is the probability that a cell has at least two users on a slot, and q is the probability that the cell has at least one source-destination pair on a slot:

$$\begin{aligned} p &= 1 - (1-d/N)^N - d(1-d/N)^{N-1} \\ q &= 1 - (1-(1/C)^2)^{N/2} \end{aligned}$$

If N and C are increased to ∞ while maintaining the ratio $N/C = d$, the resulting limiting capacity is:

$$\boxed{\mu(d) = \frac{1 - e^{-d} - de^{-d}}{2d}}$$

The capacity $\mu(d)$ is optimized at the same density as before: $d^* = 1.79328$. With this optimized density, the capacity is $\mu^* = 0.1492$. The value 0.1492 is remarkably close to a numerically optimized value of throughput in a 2-hop relay algorithm given in [2] which uses a closest-neighbor transmission scheme together with a more complex (and more realistic) signal-to-interference-plus-noise interference model. This suggests that, when density is optimized, the simplified cell partitioned network is an accurate model for a signal-to-interference-plus-noise model in a network with closest-neighbor transmission selection.

It can be shown that the network capacity depends only on the steady state location distribution and hence is the same for any mobility model that has a uniform steady state. In the special case of the i.i.d. mobility model, a queueing, scheduling, and 2-hop relaying algorithm is developed in [3] that achieves average multi-hop delay exactly given by:

$$\mathbb{E}\left[Delay\right] = \frac{N-1-\lambda}{\mu-\lambda} \quad , \forall \lambda \in (0, \mu)$$

where $\mu = (p+q)/(2d)$. In particular, the algorithm has finite average delay whenever λ is strictly less than the network capacity μ. There are no known computations for exact delay under more complex mobility models.

III. Connectivity in Static Networks

Again consider a cell partitioned network with C cells. However, assume that the number of wireless users is not fixed, but is a random variable determined by sprinkling a random number of points over the network region according to a 2-dimensional *Poisson point process* (PPP).

A. 2-d Poisson point processes

A 2-dimensional Poisson point process with intensity parameter λ (points/area) is a random placement of points over an infinite 2-dimensional plane with the following two properties:

1) Fix any region of the plane with finite area a. Let N be the random number of points that lie in this region. Then N is a Poisson random variable with parameter λa, so that:

$$P[N = k] = \frac{(\lambda a)^k}{k!} e^{-\lambda a} \quad , \forall k \in \{0, 1, 2, \ldots\}$$

in particular

$$P[N = 0] = e^{-\lambda a}$$

2) Let m be a positive integer and consider any m *disjoint* regions of the plane with areas $a_1, \ldots, a_m$, respectively. Let $N_1, \ldots, N_m$ be the random number of points in each of these disjoint regions. Then $N_1, \ldots, N_m$ are mutually independent Poisson variables with parameters $\lambda a_1, \ldots, \lambda a_m$, respectively.

The reason that the Poisson distribution is used is that the sum of an independent number of Poisson random variables is itself a Poisson random variable. The Poisson point process naturally arises if we imagine chopping the plane into an infinite number of disjoint square regions with edge size δ (for some small $\delta > 0$), so the area of each square is δ^2, and then placing a single point independently in each square with probability $\lambda\delta^2$. The Poisson point process arises in the limit as $\delta \to 0$.

B. The number of users in a given cell

Assume the network region has area 1 and all C cells are disjoint with equal area $1/C$. Then the number of users in each cell $i \in \{1, \ldots, C\}$ is a Poisson random variable with parameter λ/C:

$$P[\text{There are } k \text{ users in cell } i] = \frac{(\lambda/C)^k}{k!} e^{-\lambda/C}$$

In particular:

$$P[\text{Cell } i \text{ is empty}] = e^{-\lambda/C}$$

The average number of users in a given cell is λ/C. Thus, the average number of users in the network is λ. In this context, we define d as the *expected* user/cell density:

$$d = \frac{\lambda}{C} = \text{expected user/cell density}$$

Recall that for the mobile network context, the optimal density was $d^* = 1.79328$. Due to the mobility, if a given user is alone in its cell on a particular slot, it will eventually encounter another user at a later time slot due to mobility. However, in the static network scenario it is crucial to avoid empty cells because that can create *network disconnectivity*. We shall find that, for static networks, the density d must be significantly larger than 1.79328. In fact, we shall find that, as the number of cells C is pushed to infinity, the density *should not be held constant, but should grow logarithmically in* C.

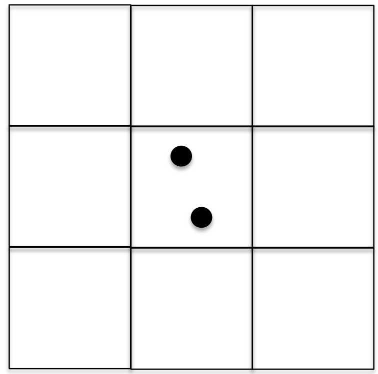

Fig. 3. A 3×3 block of cells that demonstrate a disconnection event. In this example, there are two users in the center block, surrounded by 8 empty neighbor cells. The two users cannot communicate with anyone outside of this block.

C. Connectivity

Again suppose that at most one transmission can take place per cell, and that there is no inter-cell interference. However, now assume that nodes in a cell can transmit to other nodes in *the same cell or in an adjacent cell*. If one node wants to communicate with a node on the other side of the network, it must find a *path*. In this case, the path requires a sequence of nonempty cells, each being a neighbor of its predecessor, with the first cell containing the origin user and the last containing the destination user. For simplicity, assume the square network structure of Fig. 1a is used, and that a *neigboring cell* is one that shares an edge, or one diagonally located that shares a vertex. Thus, each cell in the interior of the network has 8 neighbors.

We say the network is *connected* if there is a path between any two nodes. The example in Fig. 1a shows a *connected network* since indeed all nodes can communicate with all other nodes via a possibly multi-hop path. An example of when the network is *disconnected* is when there exists a 3×3 block of cells within the network for which the following hold (see also Fig. 3):

- There is at least one user in the center cell of the 3×3 block.
- All 8 neighbors of the center cell are empty.
- There is at least one user located in some cell of the network that is not contained in this 3×3 block.

This situation leads to a disconnected network because the users in the center cell of the 3×3 block cannot communicate with users outside the block.

Suppose users are placed in a cell partitioned network according to a Poisson point process with intensity parameter λ. What is the probability that the network is connected? This probability is difficult to compute. However, notice that the network is certainly connected *if all cells contain at least one user*. Hence:

$$P[\text{Network is connected}] \geq P[\text{all cells are nonempty}]$$

The latter probability is easy to compute because the number of users in each cell are mutually independent random variables:

$$\boxed{P[\text{all cells are nonempty}] = (1 - P[\text{Cell 1 is empty}])^C = (1 - e^{-\lambda/C})^C}$$

D. Scaling and percolation

It is interesting to consider the limiting value of $P[\text{all cells are nonempty}]$ as the number of cells C is increased to ∞ and the average number of users λ is also increased to ∞. Suppose λ is given by

$$\lambda_C = aC\log(C) \tag{1}$$

where $a > 0$ is a given constant coefficient, and we write λ_C with a subscript to emphasize dependence on C. This assumes the average number of users λ_C grows like $O(C\log(C))$. Thus, the expected user/cell density, written d_C to emphasize dependence on C, is

$$d_C = \frac{\lambda_C}{C} = a\log(C)$$

and so the density is assumed to be $O(\log(C))$. The equality (1) implies:

$$e^{-\lambda_C/C} = \frac{1}{C^a}$$

Thus

$$P[\text{all cells are nonempty}] = (1 - \frac{1}{C^a})^C$$

Evaluating $\lim_{C\to\infty}(1-\frac{1}{C^a})^C$ reveals an interesting threshold phenomenon:

- Case 1: If $a=1$, then
$$\lim_{C\to\infty} P[\text{all cells are nonempty}] = 1/e$$
- Case 2: If $a>1$, then
$$\lim_{C\to\infty} P[\text{all cells are nonempty}] = 1$$
- Case 3: If $0<a<1$, then
$$\lim_{C\to\infty} P[\text{all cells are nonempty}] = 0$$

Note that Case 2 further implies: If $a>1$ then
$$\lim_{C\to\infty} P[\text{Network is connected}] = 1$$
Thus, to ensure the network is connected with probability approaching 1, it is sufficient to use the scaling $\lambda_C = aC\log(C)$ for some parameter $a>1$. The next subsection shows that if $a<1/8$ then
$$\lim_{C\to\infty} P[\text{Network is connected}] = 0$$
In particular, if the density scales less than logarithmically, the network connectivity probability approaches 0. If the density scales more than logarithmically, the network connectivity probability approaches 1. If the density scales logarithmically, then behavior depends crucially on the coefficient a.

This raises the question of what happens if $a\in[1/8,1]$? This seems to be an open question. It is likely that there is some threshold $a^*\in[1/8,1]$ such that the network connection probability goes to 1 if $a>a^*$, and goes to 0 if $a<a^*$. However, the value a^* is unknown.

E. Disconnection probability

Recall that
$$\lambda = aC\log(C)$$
Assume $a<1/8$. To compute a bound on the disconnection probability, assume $C=(3s)^2$ for some positive integer s, so that the network region is a square grid with side length $3s$. Group the C cells into s^2 disjoint 3×3 blocks. For each block $i\in\{1,\ldots,s^2\}$, define the event E_i as the event that block i has a *disconnection event* of the type shown in Fig. 3, so that the middle cell of block i has at least one user and the 8 surrounding cells are empty. Hence
$$P[E_i] = (e^{-\lambda/C})^8(1-e^{-\lambda/C}) = e^{-8\lambda/C} - e^{-9\lambda/C} \quad, \forall i\in\{1,\ldots,s^2\}$$
If E_i^c denotes the complement of event E_i, then
$$P[E_i^c] = 1 - e^{-8\lambda/C} + e^{-9\lambda/C} = 1 - \frac{1}{C^{8a}} + \frac{1}{C^{9a}} \tag{2}$$
where the final equality uses $\lambda = aC\log(C)$. Define the (highly likely) event F that there are at least two nonempty cells. The complement event F^c is the (highly unlikely) event that either *all* cells are empty, or exactly $C-1$ cells are empty:
$$\begin{aligned}P[F^c] &= e^{-\lambda} + C(e^{-\lambda/C})^{C-1}(1-e^{-\lambda/C})\\ &= (1-C)e^{-\lambda} + Ce^{-\lambda\frac{(C-1)}{C}}\\ &= \frac{1-C}{C^{aC}} + \frac{C}{C^{a(C-1)}}\end{aligned}$$
Notice that for any $a>0$ we have:
$$\lim_{C\to\infty} P[F^c] = 0 \tag{3}$$
As described in Section III-C, if at least one event E_i occurs for some $i\in\{1,\ldots,s^2\}$, *and* if event F occurs (so that there are at least two nonempty cells), then the network *must* be disconnected. In the language of *probabilistic events* this means:
$$(\cup_{i=1}^{s^2}E_i)\cap F \subseteq \{\text{Network is disconnected}\}$$
Hence
$$P[(\cup_{i=1}^{s^2}E_i)\cap F] \leq P[\text{Network is disconnected}]$$

Thus

$$\begin{aligned}
P[\text{Network is disconnected}] &\geq P[(\cup_{i=1}^{s^2} E_i) \cap F] \\
&\overset{(a)}{\geq} P[\cup_{i=1}^{s^2} E_i] - P[F^c] \\
&= 1 - P[\cap_{i=1}^{s^2} E_i^c] - P[F^c] \\
&\overset{(b)}{=} 1 - (1 - \frac{1}{C^{8a}} + \frac{1}{C^{9a}})^{s^2} - P[F^c] \\
&\overset{(c)}{=} 1 - (1 - \frac{1}{C^{8a}} + \frac{1}{C^{9a}})^{C/9} - P[F^c]
\end{aligned}$$

where (a) holds by the probability inequality $P[A \cap B] \geq P[A] - P[B^c]$ for any events A, B; equality (b) holds by (2) together with the fact that the events $\{E_i^c\}_{i=1}^{s^2}$ are mutually independent; equality (c) holds because $C = 9s^2$. It can be shown that if $a < 1/8$ then

$$\lim_{C \to \infty} (1 - \frac{1}{C^{8a}} + \frac{1}{C^{9a}})^{C/9} = 0$$

Using this with (3) gives

$$\lim_{C \to \infty} P[\text{Network is disconnected}] \geq 1$$

Hence, if $a < 1/8$ we have

$$\lim_{C \to \infty} P[\text{Network is disconnected}] = 1$$

References

[1] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 388-404, March 2000.

[2] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. *IEEE/ACM Trans. on Networking*, vol. 10, no. 4, pp. 477-486, August 2002.

[3] M. J. Neely and E. Modiano. Capacity and delay tradeoffs for ad-hoc mobile networks. *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 1917-1937, June 2005.