Capacity of Multi-Channel Wireless Networks with Random Channel Assignment: A Tight Bound

Technical Report (October 2006)

Vartika Bhandari

Dept. of Computer Science, and

Coordinated Science Laboratory

University of Illinois at Urbana-Champaign

vbhandar@uiuc.edu

Nitin H. Vaidya

Dept. of Electrical and Computer Eng., and

Coordinated Science Laboratory

University of Illinois at Urbana-Champaign

nhv@uiuc.edu

I. INTRODUCTION

The issue of transport capacity of a randomly deployed wireless network under random (c, f) channel assignment was considered by us in [1]. We showed in [1] that when the number of available channels is $c = O(\log n)$, and each node has a single interface assigned a random f subset of channels, the capacity is $\Omega(W\sqrt{\frac{f}{cn\log n}})$ and $O(W\sqrt{\frac{P_{rnd}}{n\log n}})$, and conjectured that optimal capacity was $\Theta(W\sqrt{\frac{P_{rnd}}{n\log n}})$. We now present a lower bound construction that yields capacity $\Omega(W\sqrt{\frac{P_{rnd}}{n\log n}})$ whenever $f > 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$. Thus for values of c and f that satisfy $f > 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$, the optimal capacity under random (c, f) assignment is proved to be $\Theta(W\sqrt{\frac{P_{rnd}}{n\log n}})$. We conjecture that this would be the case for all $2 \leq f \leq c$ (for any given $c = O(\log n)$).

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **OCT 2006**	2. REPORT TYPE	3. DATES COVERED **00-00-2006 to 00-00-2006**

4. TITLE AND SUBTITLE **Capacity of Multi-Channel Wireless Networks with Random Channel Assignment: A Tight Bound**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **University of Illinois at Urbana-Champaign,Department of Computer Science,Coordinated Science Laboratory,Urbana,IL,61801**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**30**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

II. NOTATION AND TERMINOLOGY

We use the following asymptotic notation:

- $f(n) = O(g(n))$ means that $\exists c, N_o$, such that

  $f(n) \leq cg(n)$ for $n > N_o$

- $f(n) = o(g(n))$ means that $\lim_{n \to \infty} \frac{f(n)}{g(n)} = 0$

- $f(n) = \omega(g(n))$ means that $g(n) = o(f(n))$

- $f(n) = \Omega(g(n))$ means that $g(n) = O(f(n))$

- $f(n) = \Theta(g(n))$ means that $\exists c_1, c_2, N_o$, such that

  $c_1 g(n) \leq f(n) \leq c_2 g(n)$ for $n > N_o$

When $f(n) = O(g(n))$, any function $h(n) = O(f(n))$ is also $O(g(n))$. We often refer to such a situation as $h(n) = O(f(n)) \implies O(g(n))$.

As in [2], we say that the per flow network capacity is $\lambda(n)$ if each flow in the network can be guaranteed a throughput of at least λ. Whenever we use log without explicitly specifying the base, we imply the *natural* logarithm.

Some Useful Results

Theorem 1: (Vapnik-Chervonenkis Theorem) Let S be a set with finite VC dimension $VCdim(\mathcal{S})$. Let $\{X_i\}$ be i.i.d. random variables with distribution P. Then for $\varepsilon, \delta > 0$:

$$Pr\left(\sup_{D \in \mathcal{S}} |\frac{1}{N} \sum_{i=1}^{N} I_{X_i \in D} - P(D)| \leq \varepsilon\right) > 1 - \delta$$

$$\text{whenever } N > \max\left(\frac{8VCdim(\mathcal{S})}{\varepsilon} \log_2 \frac{16e}{\varepsilon}, \frac{4}{\varepsilon} \log_2 \frac{2}{\delta}\right)$$

Theorem 2: (Chernoff Bound [3]) Let $X_1, ..., X_n$ be independent Poisson trials, where $Pr[X_i = 1] = p_i$. Let $X = \sum_{i=1}^{n} X_i$. Then, for any $\beta > 0$:

$$Pr[X \geq (1+\beta)E[X]] < \left(\frac{e^\beta}{(1+\beta)^{(1+\beta)}}\right)^{E[X]} \tag{1}$$

Theorem 3: (Chernoff Upper Tail Bound [3]) Let $X_1, ..., X_n$ be independent Poisson trials, where $Pr[X_i = 1] = p_i$. Let $X = \sum_{i=1}^{n} X_i$. Then, for $0 < \beta \leq 1$:

$$Pr[X \geq (1+\beta)E[X]] \leq \exp(-\frac{\beta^2}{3}E[X]) \tag{2}$$

Theorem 4: (Chernoff Lower Tail Bound [3]) Let $X_1, ..., X_n$ be independent Poisson trials, where $Pr[X_i = 1] = p_i$. Let $X = \sum_{i=1}^{n} X_i$. Then, for $0 < \beta < 1$:

$$Pr[X \leq (1-\beta)E[X]] \leq \exp(-\frac{\beta^2}{2}E[X]) \tag{3}$$

Lemma 1: [3] When n balls are thrown independently and uniformly at random into n bins, then $Pr[$ any bin has $> \frac{3\log n}{\log\log n}$ balls $] \leq \frac{1}{n}$ for sufficiently large n.

Lemma 2: If m balls are thrown into b bins independently and uniformly at random, then $Pr[$ any bin has $> \frac{2m}{b}$ balls $] \leq b \cdot exp(-\frac{m}{3b})$.

Proof: For bin i, let $I_{i1}, I_{i2}, ..., I_{im}$ be indicator variables indicating whether ball j $(j = 1, 2, ..., m)$ fell into bin i. Then, $Pr[I_{ik} = 1] = \frac{1}{b}$. Let $X_i = \sum I_{ik}$. Then $E[X_i] = \frac{m}{b}$. By application of the Chernoff bound from Theorem 3 (setting $\beta = 1$), we then obtain that $Pr[X > \frac{2m}{b}] \leq Pr[X \geq \frac{2m}{b}] \leq exp(-\frac{m}{3b})$. Thereafter, application of the union bound yields that $Pr[$ any bin has $> \frac{2m}{b}$ balls $] \leq b \cdot exp(-\frac{m}{3b})$. ∎

Lemma 3: Suppose we are given a unit toroidal region with n nodes located uniformly at random, and the region is sub-divided into axis-parallel square cells of area $a(n)$ each. If $a(n) = \frac{100\alpha(n)\log n}{n}, 1 \leq \alpha(n) \leq \frac{n}{100\log n}$, then each cell has at least $(100\alpha(n) - 50)\log n$, and at most $(100\alpha(n) + 50)\log n$ nodes, with high probability.

Proof: We know that the set of axis-parallel squares has VC-dimension 3. In our construction, we have a set of axis-parallel square cells $\mathcal{S}$ such that the cells all have area $a(n) = \frac{100\alpha\log n}{n}$. Then considering the n random variables X_i denoting node positions, $Pr[X_i \in D(D \in \mathcal{S})] = \frac{100\alpha\log n}{n}$. Then, from

3

the VC-theorem (Theorem 1):

$$Pr\left(\sup_{D \in \mathcal{S}} \left| \frac{\text{No. of nodes in } D}{n} - \frac{100\alpha(n)\log n}{n} \right| \leq \varepsilon(n)\right) > 1 - \delta(n)$$

$$\text{whenever } n > \max\left(\frac{24}{\varepsilon}\log_2\frac{16e}{\varepsilon}, \frac{4}{\varepsilon}\log_2\frac{2}{\delta}\right)$$

This is satisfied when $\varepsilon(n) = \delta(n) = \frac{50\log n}{n}$. Thus, with probability at least $1 - \frac{50\log n}{n}$, the population $Pop(D)$ of cell D satisfies:

$$\frac{(100\alpha(n) - 50)\log n}{n} \leq Pop(D) \leq \frac{(100\alpha(n) + 50)\log n}{n} \tag{4}$$

■

Lemma 4: Let us consider the set of all circles of radius R and area $A(n) = \pi R^2$ on the unit toroid. If $A(n) = \frac{100\alpha(n)\log n}{n}, 1 \leq \alpha(n) \leq \frac{n}{100\log n}$, then each circle has at least $(100\alpha(n) - 50)\log n$, and at most $(100\alpha(n) + 50)\log n$ nodes, with high probability.

Proof: The set of all circles of radius R in the plane has VC-dimension 3. Thereafter by the same argument as in the proof of Lemma 3, the result proceeds. ■

Lemma 5: If n pairs of points (P_i, Q_i) are chosen uniformly at random in the unit area network, the resultant set of straight-line formed by each pair $L_i = P_i Q_i$ satisfies the condition that no cell has more than $n\sqrt{a(n)}$ lines passing through it.

Proof: Given the lines L_i are i.i.d., the proof of Lemma 3 in [4] can be applied to prove this result. ■

Theorem 5: (Hall's Marriage Theorem [5]) Let $S = \{S_1, S_2, \ldots S_n\}$ be a finite collection of finite sets. There exists a system of distinct representatives of S if and only if the following condition holds for any $T \subseteq S$: $|\cup T| \geq |T|$

Lemma 6: The number of subsets of size k chosen from a set of m elements is given by $\binom{m}{k} \leq \left(\frac{me}{k}\right)^k$.

4

III. RANDOM (c, f) ASSIGNMENT

In this assignment model, a node is assigned a subset of f channels uniformly at random from the set of all possible channel subsets of size f. Thus the probability that two nodes share at least one channel is given by $p_{rnd} = 1 - (1 - \frac{f}{c})(1 - \frac{f}{c-1})...(1 - \frac{f}{c-f+1})$.

Lemma 7: For $c = O(\log n)$, and $1 < f \leq c$, the following holds:

$$\frac{cp_{rnd}}{f} \leq \min\{\frac{c}{f}, 2f\} \tag{5}$$

Proof: Since $p_{rnd} \leq 1$, we obtain that $\frac{cp_{rnd}}{f} \leq \frac{c}{f}$.

If $f \geq \sqrt{\frac{c}{2}}$, then $\frac{cp_{rnd}}{f} \leq \sqrt{2c} \leq 2f$ follows from the observation that $p_{rnd} \leq 1$. Hence, we focus on the case $f < \sqrt{\frac{c}{2}}$.

$$
\begin{aligned}
1 - p_{comm} &= (1 - \frac{f}{c})(1 - \frac{f}{c-1})...(1 - \frac{f}{c-f+1}) \\
&\geq (1 - \frac{f}{c-f+1})^f > (1 - \frac{2f}{c})^f \geq 1 - \frac{2f^2}{c} \\
&\therefore p_{rnd} \leq \frac{2f^2}{c} \\
&\therefore \frac{cp_{rnd}}{f} \leq 2f
\end{aligned}
\tag{6}
$$

Thus, $\frac{cp_{rnd}}{f} \leq \min\{\frac{c}{f}, 2f\}$. ∎

Lemma 8: $\min\{\frac{c}{f}, 2f\} \leq \sqrt{2c}$

Proof: For a given c, we have $2 \leq f \leq c$. Thus, given c, $\frac{c}{f}$ is a monotonically decreasing function of f, while $2f$ is a monotonically increasing function of f. $\frac{c}{f} = 2f = \sqrt{2c}$ at $f = \sqrt{\frac{c}{2}}$. For $f \leq \sqrt{\frac{c}{2}}$, $\min\{\frac{c}{f}, 2f\} = 2f \leq \sqrt{2c}$, and for $f > \sqrt{\frac{c}{2}}$, $\min\{\frac{c}{f}, 2f\} = \frac{c}{f} \leq \sqrt{2c}$. Thus $\min\{\frac{c}{f}, 2f\} \leq \sqrt{2c}$. ∎

A. Sufficient Condition for Connectivity

This theorem has been stated and proved in [1]. However, we repeat it here in the interests of clarity.

Theorem 6: With random (c, f) assignment, if $\pi r^2(n) = \frac{800\pi \log n}{p_{rnd} n}$, then:

$$Pr[\text{ network is connected }] \rightarrow 1$$

Proof: We present a construction based on a notion of per-node backbones. Consider a subdivision of the toroidal unit area into square cells of area $a(n) = \frac{100 \log n}{p_{rnd} n}$. Then by setting $\alpha(n) = \frac{1}{p_{rnd}}$ in Lemma 3 there are at least $\frac{50 \log n}{p_{rnd}}$ nodes in each cell with high probability. Set $r(n) = \sqrt{8a(n)}$. Then a node in any given cell has all nodes in adjacent cells within its range. Within each cell, choose $\frac{2 \log n}{p_{rnd}}$ nodes uniformly at random, and set them apart as *transition facilitators* (the meaning of this term shall become clear later). This leaves at least $\frac{48 \log n}{p_{rnd}}$ nodes in each cell that can act as *backbone candidates*.

Consider any node in any given cell. The probability that it can communicate to any other random node in its range is p_{rnd}. Then the probability that in an adjacent cell, there is no backbone candidate node with which it can communicate is less than $(1 - p_{rnd})^{\frac{48 \log n}{p_{rnd}}} \leq \frac{1}{e^{48 \log n}} = \frac{1}{n^{48}}$. The probability that a given node cannot communicate with any node in some adjacent cell is thus at most $\frac{8}{n^{48}}$ (as there are upto 8 adjacent cells per node). By applying the union bound over all n nodes, the probability that at least one node is unable to communicate with any backbone candidate node in at least one of its adjacent cells is at most $\frac{8}{n^{47}}$.

We associate with each node x a set of nodes $\mathcal{B}(x)$ called the primary backbone for x. $\mathcal{B}(x)$ is constituted as follows. Throughout the procedure, cells that are already covered by the under-construction backbone are referred to as *filled* cells. x is by default a member of $\mathcal{B}(x)$, and its cell is the first *filled* cell. From each adjacent cell, amongst all backbone candidate nodes sharing at least one common channel with x, one is chosen uniformly at random is added to $\mathcal{B}(x)$. Thereafter, from each cell bordering a filled cell, of all nodes sharing at least one common channel with some node already in $\mathcal{B}(x)$, one is chosen uniformly at random, and is added to $\mathcal{B}(x)$; the cell gets added to the set of filled cells. This process continues iteratively, till there is one node from every cell in $\mathcal{B}(x)$. From our earlier observations, $\mathcal{B}(x)$ eventually covers all cells with probability at least $1 - \frac{8}{n^{47}}$. Now consider any pair of nodes x and y. If $\mathcal{B}(x) \cap \mathcal{B}(y) \neq \phi$ the two are obviously connected, as one can proceed from x on $\mathcal{B}(x)$ towards one of the intersection nodes, and thence to y on $\mathcal{B}(y)$, and vice-versa. Suppose, the two backbones are disjoint. Then x and y are still connected if there is some cell such that the member of $\mathcal{B}(x)$ in that cell (let us call it q_x) can communicate with the member of $\mathcal{B}(y)$ in that cell (let us call it q_y), either directly, or through a third node. q_x and q_y can communicate directly with probability 1 if they share a common channel. Thus the case of interest is one in which no cell has q_x and q_y sharing a channel.

If they do not share a common channel, we consider the event that there exists a third node amongst the *transition facilitators* in the cell through whom they can communicate. Note that, for two given backbones

6

$\mathcal{B}(x)$ and $\mathcal{B}(y)$, the probability that in a network cell, given q_x and q_y that do not share a channel, they can both communicate with a third node z that did not participate in backbone formation and is known to lie in the same cell, is independent across cells. Therefore, the overall probability can be lower-bounded by obtaining for one cell the probability of q_x and q_y communicating via a third node z, given they have no common channel, considering that each cell has at least $\frac{2\log n}{p_{rnd}}$ possibilities for z, and treating it as independent across cells. We elaborate this further.

Let q_x have the set of channels $C(q_x) = \{c_{x_1},...,c_{x_f}\}$, and q_y have the set of channels $C(q_y) = \{c_{y_1},...,c_{y_f}\}$, such that $C(q_x) \cap C(q_y) = \phi$. Consider a third node z amongst the transition facilitators in the same cell as q_x and q_y. We desire z to have at least one channel common with both $C(q_x)$ and $C(q_y)$. Then let us merely consider the possibility that z enumerates its f channels in some order, and then inspects the first two channels, checking the first one for membership in $C(q_x)$, and checking the second one for membership in $C(q_y)$. This probability is $\left(\frac{f}{c}\right)\left(\frac{f}{c-1}\right) > \frac{f^2}{c^2}$. Thus q_x and q_y can communicate through z with probability $p_z > \frac{f^2}{c^2} = \Omega(\frac{1}{\log^2 n})$. There are $\frac{2\log n}{p_{rnd}}$ possibilities for z within that cell, and all the possible z nodes have i.i.d channel assignments. Thus, the probability that q_x and q_y cannot communicate through any z in the cell is at most $(1-p_z)^{\frac{2\log n}{p_{rnd}}}$, and the probability they can indeed do so is $p_{xy} > 1 - (1-p_z)^{\frac{2\log n}{p_{rnd}}}$.

Thus, the probability that this happens in none of the $\frac{1}{a(n)} = \frac{p_{rnd}n}{100\log n}$ cells is at most $(1-p_{xy})^{\frac{p_{rnd}n}{100\log n}} < (1-p_z)^{\frac{2\log n}{p_{rnd}}\frac{p_{rnd}n}{100\log n}} < (1-\frac{f^2}{c^2})^{\frac{2\log n}{p_{rnd}}\frac{p_{rnd}n}{100\log n}} \rightarrow e^{-\Omega(\frac{n}{\log^2 n})}$ (recall that $c = O(\log n)$). Applying union bound over all $\binom{n}{2} < \frac{n^2}{2}$ node pairs, the probability that some pair of nodes are not connected is at most $\frac{n^2 e^{-\Omega(\frac{n}{\log^2 n})}}{2} < \frac{1}{2}e^{-\Omega(\frac{n}{\log^2 n})+2\log n} \rightarrow 0$. Thus the probability of a connected network converges to 1. ∎

IV. LOWER BOUND ON CAPACITY

We present a constructive proof that achieves $\Omega(W\sqrt{\frac{p_{rnd}}{n\log n}})$ when $f > 10(1+\log\frac{192}{5}+\log\frac{c}{f})$. Thus, e.g., this bound would hold for all $f > 10(1+\log\frac{192}{5}+\log c)$.

Once again we use a square cell construction. The surface of the unit torus is divided into square cells of area $a(n)$ each, and the transmission range is set to $\sqrt{8a(n)}$, thereby ensuring that any node in a given cell is within range of any other node in any adjoining cell. Since we utilize the *Protocol Model* [2], a node C can potentially interfere with an ongoing transmission from node A to node B, only if $BC \leq (1+\Delta)r(n)$. Thus, a transmission in a given cell can only be affected by transmissions with cells within a distance $(2+\Delta)r(n)$ from it. Since Δ is independent of n, the number of cells that interfere with

7

a given cell is only some constant (say β).

We choose $a(n) = \frac{250\max\{\log n, c\}}{p_{rnd}n} = \Theta(\frac{\log n}{p_{rnd}n})$ (since $c = O(\log n)$).

Then the following holds:

Lemma 9: Each cell has at least $\frac{4na(n)}{5} = \frac{200\max\{\log n, c\}}{p_{rnd}}$ and at most $\frac{6na(n)}{5} = \frac{300\max\{\log n, c\}}{p_{rnd}}$ nodes w.h.p.

Proof: We have chosen $a(n) = \frac{250\max\{\log n, c\}}{p_{rnd}n}$. Thus $a(n) \geq \frac{100\log n}{p_{rnd}n}$. Then if $c \leq \log n$, we can set $\alpha = \frac{2.5}{p_{rnd}} > 1$ in Lemma 3, and when $c > \log n$, i.e., $c = \gamma \log n (\gamma > 1)$ (recall that $c = O(\log n)$), we can set $\alpha = \frac{2.5\gamma}{p_{rnd}} > 1$, to obtain that the following holds with probability at least $1 - \frac{50\log n}{n}$ for all cells $\mathcal{D}$:

$$\frac{250\max\{\log n, c\}}{p_{rnd}} - 50\log n \leq Pop(\mathcal{D}) \leq \frac{250\max\{\log n, c\}}{p_{rnd}} + 50\log n$$

Thereafter noting that $\frac{250\max\{\log n, c\}}{p_{rnd}} - 50\log n \geq \frac{200\max\{\log n, c\}}{p_{rnd}}$, and $\frac{250\max\{\log n, c\}}{p_{rnd}} + 50\log n \leq \frac{300\max\{\log n, c\}}{p_{rnd}}$, completes the proof. ∎

Corollary 1: Each cell has at least $\frac{200\log n}{p_{rnd}}$ nodes w.h.p.

The constructions in the rest of this paper work on assumption that Lemma 9 holds. Thus most of the results in the rest of the paper are implicitly conditioned on this lemma.

We also state the following facts:

$$\frac{f}{c} \leq p_{rnd} \leq 1 \tag{7}$$

For large n, since $c = O(\log n)$, and $2 \leq f \leq c$:

$$na(n) = \frac{250\max\{\log n, c\}}{p_{rnd}} = O(\log^2 n)$$
$$\frac{n\sqrt{a(n)}}{c} = \frac{1}{c}\sqrt{\frac{250n\max\{\log n, c\}}{p_{rnd}}} = \Omega(\sqrt{\frac{n}{\log n}}) \tag{8}$$
$$\therefore f(n) = O(na(n)) \implies f(n) = O(\frac{n\sqrt{a(n)}}{c})$$

8

$$\frac{1}{\sqrt{a(n)}} = \sqrt{\frac{p_{rnd}n}{250\max\{\log n, c\}}} = O(\frac{n}{\log n})$$

$$\frac{n\sqrt{a(n)}}{c} = \frac{1}{c}\sqrt{\frac{250n\max\{\log n, c\}}{p_{rnd}}} = \Omega(\sqrt{\frac{n}{\log n}}) \tag{9}$$

$$\therefore f(n) = O(\frac{1}{\sqrt{a(n)}}) \implies f(n) = O(\frac{n\sqrt{a(n)}}{c})$$

We now define a term M_u where $M_u = \lceil \frac{9na(n)}{25} \rceil = \lceil \frac{90f\max\{\log n, c\}}{cp_{rnd}} \rceil$ and show that the following holds:

Lemma 10: If there are at least $\frac{200\max\{\log n, c\}}{p_{rnd}}$ nodes in every cell, of which we choose $\frac{180\max\{\log n, c\}}{p_{rnd}}$ nodes uniformly at random to examine, then, in each cell, amongst those $\frac{180\max\{\log n, c\}}{p_{rnd}}$ nodes, at least $c - \lfloor \frac{f}{4} \rfloor$ channels have at least M_u nodes capable of switching on them, w.h.p.

Proof: Consider any single cell D. Let us denote by $\mathcal{E}$ the set of $\frac{180\max\{\log n, c\}}{p_{rnd}}$ nodes lying in cell D that are chosen uniformly at random for examination. Denote by I_{ji} the indicator variable that is 1 if a node j can switch on channel i and 0 else. $Pr[I_{ji} = 1] = \frac{f}{c}$ and $X_i = \sum_{j \in \mathcal{E}} I_{ji}$ is the number of nodes in $\mathcal{E}$ capable of switching on channel i. Then $E[X_i] = \frac{f}{c}\frac{180\max\{\log n, c\}}{p_{rnd}} = 2M_u$. In light of Lemma 7, this leads to the following equations:

$$E[X_i] = \frac{180f\max\{\log n, c\}}{cp_{rnd}} \tag{10}$$

$$E[X_i] \geq \frac{180\max\{\log n, c\}}{\min\{2f, \frac{c}{f}\}} \geq \frac{90\max\{\log n, c\}}{f} \tag{11}$$

$$E[X_i] \geq 180f \text{ from Eqn. 10 noting that } p_{rnd} \leq 1 \tag{12}$$

$$E[X_i] \geq \frac{180\max\{\log n, c\}}{\min\{2f, \frac{c}{f}\}} \geq \frac{180\max\{\log n, c\}}{\sqrt{2c}} > 90\max\{\frac{\log n}{\sqrt{c}}, \sqrt{c}\} \geq 90\sqrt{\log n} \text{ (from Lemma 8)} \tag{13}$$

Note that from the following equations, it also proceeds that $M_u \geq \lceil \max\{\frac{45max\{\log n, c\}}{f}, 90f, 45\sqrt{\log n}\} \rceil$.

Let I_i' denote an indicator variable which is 1 if $X_i < \frac{E[X_i]}{2}$, and 0 else. Then from the Chernoff bound in Theorem 4, $Pr[I_i' = 1] = Pr[X_i < \frac{E[X_i]}{2}] \leq Pr[X_i \leq \frac{E[X_i]}{2}] \leq \exp(-\frac{E[X_i]}{8})$. Besides, the I_i''s are negatively correlated, as each node can only have f channels assigned to it, and thus, in the given cell D, having some channel (say c_i) assigned to a large number of nodes can only decrease the presence of another channel (say c_j).

9

Then if $X = \sum_{i=1}^{c} I_i'$, $E[X] \leq c\exp(-\frac{E[X_i]}{8}) \leq \exp(-\frac{E[X_i]}{8} + O(\log\log n)) \leq \exp(-\frac{3E[X_i]}{25})$ for large n (since $E[X_i] = \Omega(\sqrt{\log n})$ from Eqn. 13). Due to the negative correlation of I_i''s, we can still apply the Chernoff bound (this is a well-known fact, e.g., see [6]). By setting $(1+\beta)E[X] = \frac{f}{4}$ in Theorem 2 (note that $E[X] \leq \exp(-\frac{3E[X_i]}{25}) \leq \exp(-\frac{3}{25}(180f)) < \frac{f}{4}$, yielding $\beta > 0$), we obtain by appropriate substitutions at each step, the following:

$$
\begin{aligned}
Pr[X \geq \lceil\tfrac{f}{4}\rceil] \leq Pr[X \geq \tfrac{f}{4}] &\leq \left(\frac{e^\beta}{(1+\beta)^{(1+\beta)}}\right)^{E[X]} \\
&< \left(\frac{e}{(1+\beta)}\right)^{(1+\beta)E[X]} = \left(\frac{4eE[X]}{f}\right)^{\frac{f}{4}} \\
&\leq \left(\frac{4e\exp(-\frac{3}{25}\frac{90\max\{\log n,c\}}{f})}{f}\right)^{\frac{f}{4}} \quad \text{from Eqn. 11} \\
&\leq \left(\frac{4e\exp(-\frac{270\max\{\log n,c\}}{25f})}{f}\right)^{\frac{f}{4}} = \frac{\exp\left(\frac{-270\max\{\log n,c\}}{100}\right)}{(\frac{f}{4e})^{\frac{f}{4}}} \\
&\leq \frac{\exp(-2.7\max\{\log n,c\})}{(\frac{1}{2e})^{\frac{f}{4}}} \leq \frac{\exp(-2.7\max\{\log n,c\})}{(\frac{1}{e^2})^{\frac{f}{4}}} \\
&\leq \exp(-2.7\max\{\log n,c\})\exp(\tfrac{f}{2}) \quad \text{since } 2 \leq f \leq c \\
&\leq \exp(-2\max\{\log n,c\}) \leq \frac{1}{n^2}
\end{aligned}
\tag{14}
$$

Applying union bound over all $\frac{1}{a(n)} \leq n$ cells in the network, the probability that this happens in any cell is at most $\frac{1}{n}$. Thus, with probability at least $1 - \frac{1}{n}$, $X < \lceil\frac{f}{4}\rceil$, i.e., $X \leq \lfloor\frac{f}{4}\rfloor$ (since X is an integer), and hence each cell has at least $c - \lfloor\frac{f}{4}\rfloor$ channels with $X_i \geq \frac{E[X_i]}{2}$ *candidate* nodes capable of switching on them. Thus, by our definition of X, each cell has at least $c - \lfloor\frac{f}{4}\rfloor$ channels with $X_i \geq \lceil\frac{E[X_i]}{2}\rceil$ *candidate* nodes capable of switching on them (since X_i is also an integer). From Eqns. 10, 11, 12 and 13, and the definition of M_u, we know that $M_u = \lceil\frac{E[X_i]}{2}\rceil$. Thus, the lemma is proved. ∎

A. Routing

Recall that we use the traffic model of [2], where each source S first chooses a pseudo-destination D', and then selects the node D nearest to it as the actual destination. In [2], the route SD'D was followed, whereby the flow traversed cells intersected by the straight line SD', and then took an extra last hop if required. The following lemmas (some also stated in [1]) for SD'D routing are applicable here:

Lemma 11: No node is the destination of more than $O(na(n))$ flows.

Proof: While we had presented a brief proof outline for this lemma in [1], we present a more detailed proof here. Consider that a flow's pseudo-destination falls in a certain cell D. Consider a circle of radius $\sqrt{a(n)}$, and hence area $a(n)$ centered around this pseudo-destination. Then, this circle falls entirely within cell D and the 8 cells adjacent to cell D, and from Lemma 4, all such circles contain $\Theta(na(n))$ nodes. In the worst-case, one of these nodes could potentially be the source node for that flow. However, the circle still has more than one node other than the flow's source. Thus, the flow will select as its destination, some node within this circle. Hence a flow can only be assigned a destination that lies in the same cell or 8 cells adjacent to the pseudo-destination's cell. Thus, it proceeds that a node can only be destination for flows whose pseudo-destination lies within its own cell, or one of the 8 cells adjacent to it. From Lemma 3, 9 cells of area $a(n)$ each will contain $\Theta(na(n))$ pseudo-destinations. Thus no node is destination of more than $O(na(n))$ flows. ∎

Lemma 12: For large n, at least one node is a destination for $\Omega(\log n)$ flows with a probability at least $\frac{1}{e}(1 - \frac{1}{e})(1 - \delta)$, where $\delta > 0$ is an arbitrarily small constant.

Proof: The necessary condition for connectivity in [7] (Theorem 2.1 of [7]) is established by proving that if we consider $R(n)$ such that $\pi R^2(n) = \frac{\log n + b(n)}{n}$, where $\limsup b(n) = b < \infty$, then with positive probability, there exists at least one node x which is isolated, i.e., there is no other node within distance $R(n)$ of x. In the context of [7], this was utilized by interpreting $R(n)$ as transmission range, and thus obtaining a lower bound for connectivity. However, we now exploit that result in a different manner to prove our lemma as follows: Choose $R(n) = \frac{1}{\pi}\sqrt{\frac{\log n + 1}{n}}$, i.e., $b(n) = b = 1$. Note that in this proof, $R(n)$ is *not* the transmission range; it is merely a chosen distance value. Then by invoking Theorem 2.1 from [7], there exists a node A such that there is no other node within a distance $R(n)$ from it, with probability p where $\liminf p \geq e^{-b}(1 - e^{-b}) = \frac{1}{e}(1 - \frac{1}{e})$. In fact, from the proof of Theorem 2.1 in [7], it proceeds that $p \geq (1 - \varepsilon)\frac{1}{e}(1 - \frac{1}{e})$, for any $\varepsilon > 0$, and sufficiently large n. Call this event $\mathcal{E}_1$.

Thus, given event $\mathcal{E}_1$ has occurred and such a node A exists, if we consider the Voronoi tessellation generated by the n nodes, then the Voronoi polygon of A has area at least $\pi(\frac{R(n)}{2})^2 = \frac{\pi R^2(n)}{4} = \frac{\log n + 1}{4n}$. Note that this tessellation constitutes a spatial partition of the network area. Also, it immediately proceeds from the traffic model, that if a flow's pseudo-destination falls within the polygon of node x, then x is selected as that flow's destination, unless x is itself the source of that flow (since a generator (node) is

always the nearest generator to points within its own polygon). Also recall that pseudo-destinations are chosen uniformly at random. Let $X_i, 1 \leq i \leq n$ be indicator variables such that $X_i = 1$ if A is flow i's destination, and 0 else. Then $Pr[X_i = 1] = 0$ if A is the source of flow i (and there is exactly one such i). For all other values of i, $Pr[X_i = 1 | \mathcal{E}_1] \geq \frac{\log n + 1}{4n}$, since A is selected as flow i's destination if either (1) flow i's pseudo-destination falls in A's Voronoi polygon (the probability of this event is given by the area of A's Voronoi polygon, and is thus at least $\frac{\log n + 1}{4n}$, or (2) if flow i's pseudo-destination falls within the polygon of its own source, and A is the next-nearest node (we ignore this probability, as we only require a lower bound). Let $X = \sum X_i$. Thus $E[X | \mathcal{E}_1] \geq (1 - \frac{1}{n}) \frac{\log n + 1}{4} \geq \frac{\log n}{4}$ for large n. The X_i's are i.i.d., and thus application of the Chernoff bound from Theorem 4, with $\beta = \frac{1}{2}$ yields that:

$$Pr[X \leq \frac{\log n}{8} | \mathcal{E}_1] \leq Pr[X \leq \frac{E[X]}{2} | \mathcal{E}_1] \leq \exp(-\frac{E[X]}{8}) \leq \exp(-\frac{\log n}{32}) = \frac{1}{n^{\frac{1}{32}}} \tag{15}$$

Denote by $\mathcal{E}_2$ the event that some node indeed is destination to at least $\frac{\log n}{8}$ flows. Then it proceeds from Eqn. 15 that $Pr[\mathcal{E}_2 | \mathcal{E}_1] \geq 1 - \frac{1}{n^{\frac{1}{32}}}$. Also, $Pr[\mathcal{E}_2] \geq Pr[\mathcal{E}_1] Pr[\mathcal{E}_2 | \mathcal{E}_1]$. Hence at least one node is a destination for $\Omega(\log n)$ flows with a probability at least $(1 - \varepsilon)e^{-b}(1 - e^{-b})(1 - \frac{1}{n^{\frac{1}{32}}}) \geq \frac{1}{e}(1 - \frac{1}{e})(1 - \delta)$ for any chosen $\delta > \varepsilon$, and sufficiently large n. ∎

Lemma 13: The number of SD'D routes that traverse any cell is $O(n\sqrt{a(n)})$.

Proof: The proof for this lemma is based on a proof in [4]. Cosider a cell $\mathcal{D}$. From Lemma 5 (which proceeds from a lemma in [4]) we know that the number of SD' straight-lines traversing any single cell are $O(n\sqrt{a(n)})$. We must now consider the number of routes whose last D'D hop may enter this cell $\mathcal{D}$. If D is in the same cell as D', there is no extra hop. Let us now consider the case that D' lies in one of the 8 adjacent cells, but D lies in the cell $\mathcal{D}$ (note that D cannot lie in cell $\mathcal{D}$, if D' does not lie in $\mathcal{D}$ or its adjacent cells, as is evident from the proof of Lemma 11). The number of flows for which D' lies in one of the 8 cells adjacent to $\mathcal{D}$ is $O(na(n))$ w.h.p., from Lemma 3. Also from Eqn. 8, and the fact that $c > 1$, we know that $O(na(n) \implies O(n\sqrt{a(n)})$. Thus the total number of traversing routes is $O(n\sqrt{a(n)})$. ∎

Lemma 14: Given only straight-line routing (no detour; and no additional last hop), the number of flows that enter any cell on their i-th hop is at most $\lfloor \frac{5na(n)}{4} \rfloor$ w.h.p., for any i.

Proof: Recall that in our model (which is the same as in [2]) each source S chooses a pseudo-destination D', and the node nearest it as destination. The route follows the straight line SD', and may then require an additional last hop. Let us consider just the straight-line part SD'. Thus all the n SD'

lines are i.i.d. Denote by X_i^k the indicator variable which is 1 if the flow k enters a cell $\mathcal{D}$ on its i-th hop. Then, as observed in [4] (proof of Lemma 3), for i.i.d. straight lines, the X_i^k's are identically distributed, and X_i^k and X_j^l are independent for $k \neq l$. However for a given flow k, at most one of the X_i^k's can be 1 as a flow only traverses a cell once. Then $Pr[X_i^k = 1] = a(n) = \frac{250 \max\{\log n, c\}}{p_{rnd} n}$, and as the X_i^k's correspond to different flows, they are all independent.

Let $X_i = \sum_{k=1}^{n} X_i^k$. Then $E[X_i] = na(n)$. Also, for a certain i, the X_i^k's are independent [4]. Then by application of the Chernoff bound from Thereom 3 (with $\beta = \frac{1}{4}$):

$$
\begin{aligned}
Pr[X_i \geq \frac{5E[X_i]}{4}] &\leq \exp(-\frac{E[X_i]}{48}) \\
\therefore Pr[X_i &> \frac{1250 \max\{\log n, c\}}{4p_{rnd}}] \\
&\leq \exp(-\frac{250 \max\{\log n, c\}}{48 p_{rnd}}) < \frac{1}{n^5}
\end{aligned}
\tag{16}
$$

The maximum value that i can take is $\frac{2}{\sqrt{a(n)}} = \sqrt{\frac{2np_{rnd}}{250 \max\{\log n, c\}}} < n$. Also the number of cells is $\frac{1}{a(n)} \leq n$. Then by application of union bound over all i, and all cells $\mathcal{D}$, the probability that $X_i \geq \frac{5E[X_i]}{4}$ is less than $\frac{1}{n^3}$, and thus the number of flows that enter any cell on any hop is less than $\frac{5na(n)}{4} = \frac{1250 \max\{\log n, c\}}{4p_{rnd}}$ with probability at least $1 - \frac{1}{n^3}$. Resultantly, since X_i is an integer, we can say that it is at most $\lfloor \frac{5na(n)}{4} \rfloor$ w.h.p. ∎

Having stated and proved these lemmas, we now describe our routing and link-scheduling strategy further.

Similar to the construction for connectivity in Section 6, we construct a backbone for each node.

Initially, from each cell, we choose $\frac{180 \max\{\log n, c\}}{p_{rnd}}$ nodes uniformly at random as backbone candidates. The remaining nodes (which are at least $\frac{20 \max\{\log n, c\}}{p_{rnd}}$ in number) are deemed *transition facilitators*.

A channel i is deemed *proper* in cell $\mathcal{D}$ if it occurs in at least M_u backbone candidate nodes in $\mathcal{D}$. Then from Lemma 10, the number of *proper* channels in a cell is $c' \geq c - \lfloor \frac{f}{4} \rfloor \geq c - \lfloor \frac{c}{4} \rfloor \geq \lceil \frac{3c}{4} \rceil \geq \frac{3c}{4}$. Besides, we can show the following:

13

Lemma 15: [1] Consider any cell $\mathcal{D}$, with c' *proper* channels. Let $\mathcal{A}$ be the set of all nodes lying in the 8 adjacent cells $\mathcal{D}(k), 1 \leq k \leq 8$. Let $\mathcal{C}(\mathcal{B})$ denote the union of the available set(s) of proper channels (w.r.t. cell $\mathcal{D}$) of all nodes in set $\mathcal{B} \subseteq \mathcal{A}$. Then for all $\mathcal{B} \subseteq \mathcal{A}$ such that $|\mathcal{B}| = \lceil \frac{fna(n)}{4c} \rceil$, the following holds: If $f > 10(1 + \log \frac{192}{5} + \log \frac{c}{f})$:

$$|\mathcal{C}(\mathcal{B})| \geq \lceil \frac{3c}{8} \rceil$$

Proof: Recall that $c' \geq c - \lfloor \frac{f}{4} \rfloor \geq c - \lfloor \frac{c}{4} \rfloor \geq \lceil \frac{3c}{4} \rceil \geq \frac{3c}{4}$. Also, we are considering the case $f > 10(1 + \log \frac{192}{5} + \log \frac{c}{f})$. Recall from Lemma 9 that no cell has more than $\frac{6na(n)}{5} = \frac{300 \max\{\log n, c\}}{p_{rnd}}$ nodes w.h.p. The total number of nodes in $\mathcal{A} = \overset{8}{\underset{k=1}{\cup}} D(k)$ is at most $\frac{48na(n)}{5}$.

From Lemma 6, the number of subsets of the specified cardinality is thus at most $\binom{\frac{48na(n)}{5}}{\lceil \frac{fna(n)}{4c} \rceil} \leq \left(\frac{\frac{48ena(n)}{5}}{\lceil \frac{fna(n)}{4c} \rceil} \right)^{\lceil \frac{fna(n)}{4c} \rceil} \leq \left(\frac{\frac{48ena(n)}{5}}{\frac{fna(n)}{4c}} \right)^{\frac{fna(n)}{4c}+1} \leq \left(\frac{192ec}{5f} \right)^{\frac{fna(n)}{4c}+1} = \exp((1 + \log \frac{192}{5} + \log \frac{c}{f})(\frac{fna(n)}{4c} + 1)) < \exp(\frac{f}{10}(\frac{fna(n)}{4c}) + 1) = \exp(\frac{f^2na(n)}{40c} + \frac{f}{10}) < \exp(\frac{f^2na(n)}{40c}) \leq \exp(\frac{11f^2na(n)}{400c})$ (recall that $f > 10(1 + \log \frac{192}{5} + \log \frac{c}{f})$, and also that, from Eqn. 12, $\frac{f}{10} \leq \frac{2M_u}{1800}$, which yields $\frac{fna(n)}{2500c} < \frac{f^2na(n)}{400c}$).

Consider one such subset $\mathcal{B}$ of specified cardinality. Denote by X_i the indicator variable which is 1 if channel i is not a member of $\mathcal{C}(\mathcal{B})$ and 0 else. Recall that each node has an i.i.d. random f subset of channels assigned to it. Then $Pr[X_i = 1] = (1 - \frac{f}{c})^{|\mathcal{B}|} = (1 - \frac{f}{c})^{\frac{fna(n)}{4c}} \leq e^{-\frac{f}{c}\frac{fna(n)}{4c}} = e^{-\frac{f^2na(n)}{4c^2}}$. Also, the X_i's are negatively correlated. Let $X = \sum_{i \text{ proper in } D} X_i$. Then $E[X] \leq c'e^{-\frac{f^2na(n)}{4c^2}}$. Setting $(1 + \beta)E[X] = \frac{c'}{2}$, one can see that $\beta = \frac{c'}{2E[X]} - 1 \geq \frac{c'}{2c'e^{-\frac{f^2na(n)}{4c^2}}} - 1 = \frac{e^{\frac{f^2na(n)}{4c^2}}}{2} - 1 \geq \frac{e^{\frac{125}{4}}}{2} - 1 > 0$ (recall that $na(n) = \frac{250 \max\{\log n, c\}}{p_{rnd}} \geq \frac{250c \max\{\log n, c\}}{2f^2} \geq \frac{125c^2}{f^2}$, from Lemma 7). Thus we can apply the Chernoff bound from Theorem 2 to obtain

---

[1] This can be viewed as a special variant of the Coupon Collector's problem [3], where there are c different types of coupons, each box has a random subset of f diferent coupons, and from a given population of boxes, we seek to ensure that any subset of boxes of a given cardinality will yield at least one each of at least $\lceil \frac{3c}{8} \rceil$ distinct coupons. While variants having multiple coupons per box have been considered in work on coding [8], they are not quite the same as ours.

that:

$$Pr[X \geq \frac{c'}{2}] < \left(\frac{e^\beta}{(1+\beta)^{(1+\beta)}}\right)^{E[X]} < \left(\frac{e}{(1+\beta)}\right)^{(1+\beta)E[X]}$$

$$= \left(\frac{2eE[X]}{c'}\right)^{\frac{c'}{2}} \leq \left(\frac{2ec'\exp(-\frac{f^2na(n)}{4c^2})}{c'}\right)^{\frac{c'}{2}}$$

$$\leq \left(2e\exp(-\frac{f^2na(n)}{4c^2})\right)^{\frac{c'}{2}} \leq (2e)^{\frac{c'}{2}}\left(\exp(-\frac{f^2na(n)}{4c^2})\right)^{\frac{c'}{2}} \qquad (17)$$

$$\leq (2e)^{\frac{c}{2}}\left(\exp(-\frac{f^2na(n)}{4c^2})\right)^{\frac{3c}{8}} \leq e^c\left(\exp(-\frac{3f^2na(n)}{32c})\right)$$

$$< \exp(-\frac{f^2na(n)}{16c})$$

$$(\because na(n) = \frac{250\max\{\log n, c\}}{p_{rnd}} \geq \frac{250c\max\{\log n, c\}}{2f^2}, \therefore c < \frac{f^2na(n)}{32c})$$

Taking union bound over all possible subsets $\mathcal{B}$, we get that the probability it happens for any such subset $\mathcal{B}$ is at most $\exp(\frac{11f^2na(n)}{400c}) \cdot \exp(-\frac{f^2na(n)}{16c}) = \exp(-\frac{7f^2na(n)}{200c}) = \exp(-\frac{35f^2\max\{\log n, c\}}{4cp_{rnd}}) \leq \exp(-\frac{35\max\{\log n, c\}}{8}) < \frac{1}{n^4}$. Another union bound over all $\frac{1}{a(n)} < n$ cells yields that the probability of $C(\mathcal{B})$ having less than $\frac{c'}{2}$ channels is at most $\frac{1}{n^3}$ over all cells D. Also, random variable X is an integer, and thus $X < \frac{c'}{2} \implies X \leq \lfloor\frac{c'}{2}\rfloor$. Thus $C(\mathcal{B}) \geq c' - \lfloor\frac{c'}{2}\rfloor = \lceil\frac{c'}{2}\rceil$. Finally observe that $\frac{c'}{2} \geq \frac{3c}{8}$. This completes the proof that $C(\mathcal{B}) \geq \lceil\frac{3c}{8}\rceil$. ∎

As mentioned earlier, the routing strategy is based on a per-node backbone structure similar to that used to prove the sufficient condition for connectivity. However, instead of constructing a full backbone for each node, only a partial backbone $\mathcal{B}_p(x)$ is constructed for each node x. $\mathcal{B}_p(x)$ only covers those cells which are traversed by flows for which x is either source or destination. A flow first proceeds along the route on the source backbone and will then attempt to switch onto the destination backbone.

We shall explain the backbone construction procedure in detail later. First we show how flows can be routed along the backbones.

Lemma 16: Suppose a flow has source x and destination y. Thus it is initially on $\mathcal{B}_p(x)$ and finally needs to be on $\mathcal{B}_p(y)$. Then after having traversed $\frac{c^2}{f^2}$ distinct hops (recall that $2 \leq f \leq c$), $c = O(\log n)$) , it will have found an opportunity to make the transition w.h.p. Moreover, if each flow gets to traverse

15

$\frac{c^2}{f^2}$ distinct hops (the $\frac{c^2}{f^2}$ hops for an individual flow need to be distinct; some of the flows may traverse common cells), then all n flows are able to transition w.h.p.

Proof: Consider a flow traversing a sequence of cells $D_1, D_2, \ldots$. Then if the representative of $\mathcal{B}_p(x)$ (let us call it q_x) in D_i can communicate (directly or indirectly) with the representative of $\mathcal{B}_p(y)$ (let us call it q_y) in D_i, it is possible to switch directly from $\mathcal{B}_p(x)$ to $\mathcal{B}_p(y)$. If q_x and q_y share a channel this is trivial. If q_x and q_y do not share a channel, we consider the probability that the two can communicate via a third node from amongst the *transition facilitators* in D_i, i.e. there exists a transition facilitator z such that z shares at least one channel with q_x and one channel with q_y. In Section 6, we showed that q_x and q_y can communicate through a given z with probability $p_z > \frac{f^2}{c^2} = \Omega(\frac{1}{\log^2 n})$. Given our choice of cell area $a(n)$, and conditioned on the fact that each cell has $\frac{200 \max\{\log n, c\}}{p_{rnd}}$ nodes, of which $\frac{180 \max\{\log n, c\}}{p_{rnd}}$ are deemed *backbone candidates* and the rest are *transition facilitators*, there are at least $20\frac{\max\{\log n, c\}}{p_{rnd}} \geq \frac{20 \log n}{p_{rnd}}$ possibilities for z within that cell. All the possible z nodes have i.i.d. channel assignments. Thus, the probability that q_x and q_y cannot communicate through any z in the cell is at most $(1 - p_z)^{\frac{20 \log n}{p_{rnd}}}$, and the probability they communicate through some z is $p_{xy} > 1 - (1 - p_z)^{\frac{20 \log n}{p_{rnd}}}$.

Thus, the probability that this happens in none of the $\frac{c^2}{f^2}$ distinct cells is at most $(1 - p_{xy})^{\frac{c^2}{f^2}} < (1 - p_z)^{\frac{20c^2 \log n}{f^2 p_{rnd}}} < (1 - \frac{f^2}{c^2})^{\frac{20c^2 \log n}{f^2 p_{rnd}}} \rightarrow e^{-\frac{20 \log n}{p_{rnd}}} < \frac{1}{n^{20}}$. Applying union bound over all n flows, the probability that all flows are able to transition is at least $1 - \frac{1}{n^{19}}$. ∎

Thus, we require each route to have at least $\frac{c^2}{f^2}$ distinct hops. Resultantly, we cannot stipulate that *all* flows be routed along the (almost) straight-line path SD'D (Fig. 1). If SD'D is short, a detour may be required to ensure the minimum route-length. Such flows are said to be detour-routed.

Flow Transition Strategy: We stipulate that a non-detour-routed flow is initially in a *progress-on-source-backbone* mode, and keeps to the source backbone till there are only $\frac{c^2}{f^2}$ intermediate hops left to the destination. At this point, it enters a *ready-for-transition* mode, and actively seeks opportunities to make a transition to the destination backbone along the remaining hops. Once it has made the transition into the destination backbone, it proceeds towards the destination on that backbone along the remaining part of the route, and is thus guaranteed to reach the destination.

Thus, we stipulate that the (almost) straight-line SD'D path be followed if the straight-line route

16

Fig. 1.   Routing along a straight line

comprises $h \geq \frac{c^2}{f^2}$ hops. If S and D' (hence also D) lie close to each other, the hop-length of the straight line cell-to-cell path can be much smaller. In this case, a *detour* path SPD'D is chosen (Fig. 2), using a circle of radius $\frac{c^2}{f^2}r(n)$ in a manner similar to that for adjacent (c, f) assignment.

A detour-routed flow is always in *ready-for-transition* mode.

The need to perform *detour* routing for some source-destination pairs does not have any substantial effect on the average hop-length of routes or the relaying load on a cell, as we show further.

Lemma 17: The length of any route increases by at most $O(\log^2 n)$ hops.

Proof: The proof proceeds directly from the *detour* routing strategy. Recall that the area of a cell is $\frac{250 \max\{\log n, c\}}{p_{rnd} n}$, i.e., the side of each cell is $\Theta(\sqrt{\frac{\log n}{p_{rnd} n}})$ (more precisely it is $\frac{r(n)}{\sqrt{8}}$). The distance *SP* in Fig. 2 is at most $\frac{c^2}{f^2}r(n)$ (radius of circle), yielding at most $O(\frac{c^2}{f^2})$ hops, while *PD* is again *at most* $\Theta(\frac{c^2}{f^2})$ hops (diameter of circle). This increases route length by at most $O(\frac{c^2}{f^2}) = O(\log^2 n)$ hops (recall that $c = O(\log n)$). ∎

Fig. 2.   Illustration of detour routing

Lemma 18: If the number of flows in any cell is x in case of pure straight-line routing, it is at most $x + O(\frac{nc^4 r^2(n)}{f^4}) = x + O(\log^6 n)$ in case of detour routing.

Proof: Recall that $c = O(\log n)$. Since the detour occurs only up to a circle of radius $\frac{c^2}{f^2} r(n)$, the extra flows that may pass through a cell (compared to straight-line routing) are only those whose sources lie within a distance $\frac{c^2}{f^2} r(n)$ from some point in this cell. Thus all such possible sources fall within a circle of radius $(1 + \frac{c^2}{f^2}) r(n)$, and hence area $a_c(n) = \Theta(\frac{c^4 r^2(n)}{f^4})$. Then from Lemma 4 (with a suitable choice of $\alpha(n) \geq 1$), with high probability, any circle of this radius will have $O(na_c(n))$ nodes, and hence at most $O(na_c(n))$ sources. Thus the number of extra flows that traverse the cell due to detour routing is $O(na_c(n))$, and the total number of flows is $x + O(\frac{nc^4 r^2(n)}{f^4})$. Since $nr^2(n) = O(\frac{\log n}{p_{rnd}})$, and $p_{rnd} \geq \frac{f}{c}$, the total number of flows is $O(\frac{c^5 \log n}{f^5}) \implies x + O(\log^6 n)$ w.h.p.   ∎

Lemma 19: The number of flows traversing any cell is $O(n\sqrt{a(n)})$ even with detour routing.

Proof: From Lemma 13, we know that the number of flows passing through cell C with SD'D routing (without detours) is $O(n\sqrt{a(n)})$. Thus, from Lemma 18, the number of flows through a cell C, even after some flows are detour-routed, is at most $O(n\sqrt{a(n)}) + O(\log^6 n) \implies O(n\sqrt{a(n)})$ (since $a(n) = \Theta(\frac{\log n}{p_{rnd}})$).

18

Lemma 20: The number of flows traversing any cell in *ready-for-transition* mode is $O(\log^6 n)$ w.h.p.

Proof: First let us account for the SD' stretch of each flow, without considering the possible additional last hop. We account for it explicitly later in this proof.

By our construction, a non-detour routed flow enters the *ready-for-transition* mode only when it is $\frac{c^2}{f^2}$ hops away from its destination. All such flows must have their pseudo-destinations within a circle of radius $\Theta(\frac{c^2}{f^2}r(n))$ centered in the cell. The number of pseudo-destinations that lie within a circle of radius $\Theta(\frac{c^2}{f^2}r(n))$ from the cell is $\Theta(\frac{nc^4r^2(n)}{f^4}) \implies O(\frac{c^5}{f^5}\log n)$ w.h.p., (by observing that $p_{rnd} \geq \frac{f}{c}$, and using suitable choice of $\alpha(n) = O(\frac{c^5}{f^5})$ in Lemma 4). Also $c = O(\log n)$. Hence all channels have $O(\log^6 n)$ non-detour-routed transitioning flows in the cell w.h.p.

A detour-routed flow is always in *ready-for- transition* mode. By Lemma 18, there are $O(\log^6 n)$ such flows traversing any cell. Each such flow can only traverse a cell twice along the SD' stretch. This yields $O(\log^6 n)$ detour-routed flows (including repeat traversals).

Also, the cell may be re-traversed by some flows on their additional last hop. By an argument similar to Lemma 11, there are $O(na(n))$ pseudo-destinations in the adjacent cells, and thus, from Eqn. 7, $O(na(n)) = O(\frac{\log n}{p_{rnd}}) \implies O(\log^2 n)$ such last hop flow traversals.

Thus the number of flows transitioning in any cell is $O(\log^6 n)$ w.h.p.

■

The backbone construction procedure is different from the one in the proof of Theorem 6 in that we take load-balancing into account. Thus we can describe the procedure for constructing the backbone $\mathcal{B}_p(x)$ of x as follows:

Given a cell $\mathcal{D}$, the 8 cells adjacent to cell $\mathcal{D}$ are denoted as $\mathcal{D}(k), 1 \leq k \leq 8$.

Let $m(X,i)$ denote the number of new flows that enter cell X in step i. Step 0 of the procedure occurs in the very beginning, when each node is assigned the flow for which it is the source. Since there are no flows entering a cell at this stage, we define $m(X,0)$ as the number of flows originating in cell X.

Fig. 3.   Cell D and neighboring cells during backbone construction

Each flow has a unique source. Then $m(\mathcal{D}, i) \leq \sum_{k=1}^{8} m(\mathcal{D}(k), i-1)$ since the flows entering cell D in step i must have entered one of the 8 adjacent cells in the previous hop (or originated from one of them, in case $i = 1$). The situation is illustrated in Fig. 3.

$\mathcal{B}_p(x)$ is constituted for all nodes as follows. Let $\mathcal{S} \cup \mathcal{D}_b$ be the subset of cells that must be covered by $\mathcal{B}_p(x)$ where $\mathcal{S}$ comprises cells traversed by the flow for which x is the source, and $\mathcal{D}_b$ comprises the cells traversed by flows for which it may be the destination. x is by default a member of $\mathcal{B}_p(x)$.

We consider backbone construction for the route each source to its pseudo-destination below. Some routes will require an additional last hop. However, note that the only last hop routes that may enter a cell will correspond to pseudo-destinations in the 8 adjacent cells. Then from Lemma 3, they are only $O(na(n))$ such pseudo-destinations, and thus only $O(na(n))$ such last-hop flows. Hence we can account for them separately.

   a) Expanding backbones to $\mathcal{S}$: We first cover cells in $\mathcal{S}$. Recall that we are only constructing the SD' part and not considering the possible additional last hop at this stage.

This has two sub-stages. In the first stage, we construct backbones for source nodes whose flow does not require a detour. In the second sub-stage we construct backbones for source nodes whose flow requires a detour.

20

Straight-line backbones:

This step proceeds in a hop-by-hop manner for all non-detour-routed flows in parallel (each of which has a unique source x).

Any cell of S in which there is already a node assigned to $\mathcal{B}_p(x)$ is called a filled cell. Thus initially x's cell is filled. We next consider the cell in S that is traversed next by the flow. We consider all nodes in that cell sharing one or more common channel with x. This provides a number of alternative channels on which to switch a flow into that cell.

Let h_{max} be the maximum hop-length of any non-detour-routed SD' route. Then, the procedure has $h_{max} = O(\frac{1}{\sqrt{a(n)}})$ steps. In step i, for each source node x whose flow has more than i hops, $\mathcal{B}_p(x)$ expands into the cell entered by x's flow on the i-th hop. Each cell D performs the following procedure:

Let the number of *proper* channels in $\mathcal{D}$ be c'. From Lemma 10, we know that $c' \geq c - \frac{f}{4} \geq \frac{3c}{4}$. Each flow that are enters cell $\mathcal{D}$ in step i has a previous hop-node in one of the 8 adjacent cells. Also note that, from Lemma 10, each previous hop node has at least $\lceil \frac{3f}{4} \rceil$ of cell $\mathcal{D}$'s *proper* channels available to it as choices (since it has f channels of which at most $\lfloor \frac{f}{4} \rfloor$ may be non-proper in cell $\mathcal{D}$). The backbones are extended by constructing a bipartite graph that aids load-balance.

Lemma 21: After step h_{max} of the backbone construction procedure for S (for non-detour-routed flows), no cell has more than $O(\frac{n\sqrt{a(n)}}{c})$ incoming backbone links on a single channel, and no node appears on more than $O(\frac{n\sqrt{a(n)}}{c})$ (source) backbones.

Proof: Recall that we are expanding backbones to cover cells in S. The proof proceeds by induction. We prove that after step i of the backbone construction procedure, the following two invariants hold for *all* cells of the network:

- *Invariant 1:* Each node is assigned at most 14 new incoming backbone links during step i. Thus after step i, it appears in a total of $O(14i)$ backbones.
- *Invariant 2:* No more than $\lfloor \frac{5na(n)}{c} \rfloor$ new backbone links enter the cell on a single channel during step i. Thus, in total $O(\frac{ina(n)}{c})$ incoming backbones (entering the cell) are assigned (incoming links) on a single channel after step i.
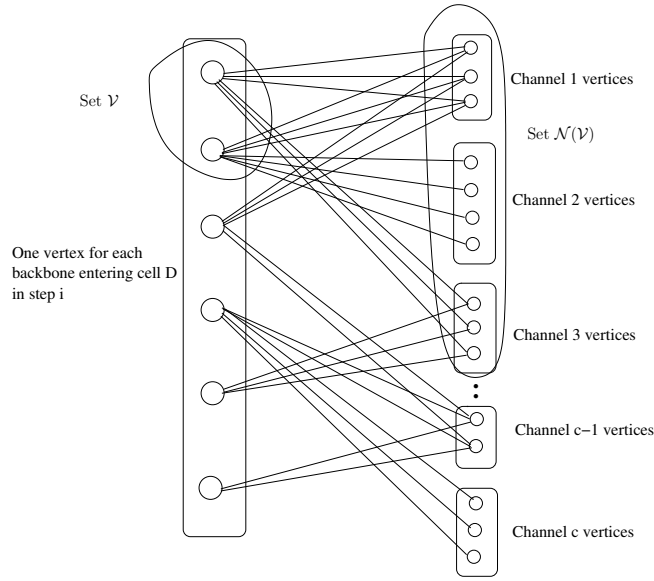
21

Fig. 4. Bipartite Graph for Cell $\mathcal{D}$ in step i

If the above two Invariants hold, then it is easy to see that after h_{max} steps, cell $\mathcal{D}$ will have no more than $\frac{5h_{max}na(n)}{c} = O(\frac{n\sqrt{a(n)}}{c})$ backbone links assigned to any single channel, and no node occurs on more than $14h_{max} \implies O(\frac{1}{\sqrt{a(n)}}) \implies O(\frac{n\sqrt{a(n)}}{c})$ backbones (from Eqn. 9).

We prove that the Invariants hold, by induction, as follows:

If Invariant 1 holds after step $i-1$, then Invariant 2 holds after step i. If Invariant 2 holds after step i, then Invariant 1 will also continue to hold after step i.

Base Case:

Before the procedure begins, at step 0, each node is assigned to its own backbone, for which it is effectively the origin (and this can be viewed as a single backbone link incoming to this node from a imaginary super-source). Thus after Step 0, Invariant 1 holds trivially, and Invariant 2 is irrelevant, and thus trivially true.

Inductive Step:

Suppose Invariants 1 and 2 held at the end of step $i-1$.

Consider the particular cell $\mathcal{D}$ during step i.

22

Construct a bipartite graph with two sets of vertices (Fig. 4); one set (call it $\mathcal{L}$) has a vertex corresponding to each of the $m(\mathcal{D}, i)$ (source) backbones that enter the cell $\mathcal{D}$ in step i, and the other set (call it $\mathcal{P}$) has $\lfloor \frac{5na(n)}{c} \rfloor \leq \frac{5na(n)}{c}$ vertices for each proper channel i in cell $\mathcal{D}$.

A backbone vertex is connected to all the vertices for the channels on which its previous hop node can switch (and which are therefore valid channel choices for entering the cell $\mathcal{D}$). We show that there exists a matching that pairs each backbone vertex to a unique channel vertex, through an argument based on Hall's marriage theorem (Theorem 5). Thus, we seek to show that for all $\mathcal{V} \subseteq \mathcal{L}$, $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$, where $\mathcal{N}(\mathcal{V}) \subseteq \mathcal{P}$ is the union of the neighbor-sets of all vertices in $\mathcal{V}$.

Consider the following two cases:

Case 1: $|\mathcal{V}| < \frac{29fna(n)}{8c}$: Consider any set $\mathcal{V}$ of backbone vertices such that $|\mathcal{V}| < \frac{29fna(n)}{8c}$. Then, since there are at most $\lfloor \frac{f}{4} \rfloor$ non-proper channels in a cell, every previous hop node has at least $\lceil \frac{3f}{4} \rceil \geq \frac{3f}{4}$ *proper* channel choices. For each proper channel there are $\lfloor \frac{5na(n)}{c} \rfloor \geq \frac{5na(n)}{c} - 1$ associated channel vertices. Thus we obtain that $|\mathcal{N}(\mathcal{V})| \geq \frac{3f}{4} \left(\frac{5na(n)}{c} - 1 \right) \geq \frac{15fna(n)}{4c} - \frac{3f}{4} \geq \frac{15fna(n)}{4c} - \frac{3fna(n)}{1000c} \geq \frac{29fna(n)}{8c} (\because na(n) \geq 250c)$. Thus $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$.

Case 2: $|\mathcal{V}| \geq \frac{29fna(n)}{8c}$: Now consider sets $\mathcal{V}$ of size at least $\frac{29fna(n)}{8c}$. Since Invariant 1 held till end of step $i-1$, no more than 14 backbone links were assigned to any single node in $\overset{8}{\underset{k=1}{\cup}} \mathcal{D}(k)$ in step $i = 1$. Since no node can be previous hop in step i of more flows than those assigned to it in the previous step, no previous hop node is common to more than 14 entering backbone links. Thus, the number of distinct previous hop nodes associated with these entering links is at least $\frac{1}{14}\left(\frac{29fna(n)}{8c} \right) \geq \frac{fna(n)}{4c} + \frac{fna(n)}{112c} > \frac{fna(n)}{4c} + 1 \geq \lceil \frac{fna(n)}{4c} \rceil$ (note that $\frac{fna(n)}{c} \geq 250f \geq 500 > 112$). Then from Lemma 15, and the fact that $\mathcal{V}$ is associated with at least one subset of $\lceil \frac{fna(n)}{4c} \rceil$ previous hop nodes, $\mathcal{N}(\mathcal{V})$ has vertices corresponding to at least $\lceil \frac{3c}{8} \rceil$ *proper* channels, and thus $|\mathcal{N}(\mathcal{V})| \geq \lceil \frac{3c}{8} \rceil \lfloor \frac{5na(n)}{c} \rfloor \geq \frac{3c}{8}\left(\frac{5na(n)}{c} - 1 \right) \geq \frac{15na(n)}{8} - \frac{3c}{8} > \frac{5na(n)}{4}$ (from the observation that $na(n) = \frac{250 \max\{\log n, c\}}{p_{rnd}} \geq 250c$). Recall that from Lemma 14, that no more than $\lfloor \frac{5na(n)}{4} \rfloor \leq \frac{5na(n)}{4}$ flows enter cell $\mathcal{D}$ on any hop. Thus for all possible sets $\mathcal{V}$, $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$.

Hence, by application of Hall's marriage theorem (Theorem 5), each backbone vertex can be matched with a unique channel vertex, and the corresponding backbone will be assigned to the channel with which this vertex is associated. Thus all backbones get assigned a channel, and (since there are at most $\lfloor \frac{5na(n)}{c} \rfloor$ channel vertices for one channel) no more than $\lfloor \frac{5na(n)}{c} \rfloor$ incoming backbone links are assigned to any

single channel.

While Hall's marriage theorem proves that such a matching exists, the matching itself can be computed using the Ford-Fulkerson method [9] on the bipartite graph [2].

Thus Invariant 2 holds. Having determined the channel each backbone should use to enter cell D, we now need to assign a node in cell D to each backbone. For this, we again construct a bipartite graph. In this graph, the first set of vertices (call it $\mathcal{F}$) comprise a vertex for each backbone entering cell D in step *i*. The second set (call it $\mathcal{R}$) comprises 14 vertices for each *backbone candidate* node in cell D. A vertex *x* in $\mathcal{F}$ has an edge with a vertex *y* in $\mathcal{R}$ iff the actual *backbone candidate* node associated with *y* is capable of switching on the channel assigned to the backbone associated with vertex *x*.

From Lemma 10, it proceeds that each vertex $x \in \mathcal{F}$ has degree at least $14M_u$, since it is assigned to a proper channel, which has at least M_u representatives in cell $\mathcal{D}$, each of which have 14 associated vertices in $\mathcal{R}$. Also recall that $M_u = \lceil \frac{9fna(n)}{25c} \rceil$. Once again we seek to show that for all $\mathcal{V} \subseteq \mathcal{F}, |\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$. Consider the following two cases:

Case 1: $|\mathcal{V}| < 14M_u$: Consider any set $\mathcal{V} \subseteq \mathcal{F}$ such that $|\mathcal{V}| < 14M_u$. Then, by our observation that each vertex in $\mathcal{F}$ has degree at least $14M_u$, it immediately proceeds that $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$.

---

[2]It is interesting to consider whether load-balance would continue to hold even if we follow simpler procedures. We have shown in [1] that for random (c, f) assignment, a per-flow throughput of $\Theta(W\sqrt{\frac{f}{cn\log n}})$ is achievable with a much simpler construction. That construction is of interest despite not achieving optimal capacity since it provides a trade-off between throughput and routing/scheduling complexity. In fact when f is a small constant, the asymptotic capacity for both constructions is within a small constant factor of each other. However, it is also useful to consider whether simpler procedures can allow one to achieve the optimal capacity. As an illustration, consider a procedure where a backbone link is assigned to the least-loaded of all channels available to it. If this procedure can be proved to yield optimal load-balance, it would have useful practical implications toward potentially indicating that even simple protocols can suffice for good performance. This problem is a special variant of the problem of throwing balls into bins with the power of d choices. The problem of throwing a balls into b bins with d choices was studied in [10]. In [11], a balls-and-bins technique is used to obtain fractional matchings in graphs. However these results yield probability bounds polynomial in number of bins. In our case, the bins (channels) are $O(\log n)$ (where n is number of nodes), and we need much stronger bounds to ensure that global overload probability goes to 0, and thus a simple adaptation of existing balls-into-bins proofs does not suffice. Our case also has additional constraints, e.g., the number of choices available to each ball is $\Theta(f)$, and the number of balls (traversing source backbones) decreases with increase in f.

Also of interest is the possibility of having optimal-capacity achieving procedures where backbones are constructed sequentially, or even better, completely asynchronously (recall that the simpler construction possesses these properties, but yields sub-optimal capacity). If such a procedure can be shown to achieve good load balance, it has useful protocol implications in that when a new flow is admitted, routes for existing flows do not need to be re-organized to ensure load-balance.

Case 2: $|\mathcal{V}| \geq 14M_u$: Consider sets $\mathcal{V}$ of size $\alpha M_u, \alpha \geq 14$. Since no channel is assigned more than $\frac{5na(n)}{c}$ entering backbone links in this step, the vertices in $\mathcal{V}$ are cumulatively associated with at least $\frac{\alpha M_u}{\frac{5na(n)}{c}} \geq \frac{\frac{18\alpha f na(n)}{25c}}{\frac{5na(n)}{c}} \geq \frac{18\alpha f}{125}$ distinct proper channels. Since each of these channels have at least M_u *backbone candidate* nodes capable of switching on them, and any one node can only switch on up to f proper channels, this implies that the number of nodes in cell $\mathcal{D}$ cumulatively associated with these $\frac{9\alpha f}{125}$ channels is at least $\frac{9\alpha M_u}{125}$, and as each node has 14 vertices, it follows that $|\mathcal{N}(\mathcal{V})| \geq 14\left(\frac{9\alpha M_u}{125}\right) \geq \frac{126\alpha M_u}{125} > \alpha M_u > |\mathcal{V}|$.

Then by invoking Hall's Marriage Theorem again, each vertex $x \in \mathcal{F}$ can be matched with a unique vertex $y \in \mathcal{R}$, and the actual network node associated with y is deemed the backbone representative for backbone x in cell $\mathcal{D}$. Since there are at most 14 vertices associated with a node, no node is assigned more than 14 incoming backbone links in step i, and Invariant 2 continues to hold.

Thus after step $h_{max} \leq \frac{2}{\sqrt{a(n)}}$, each cell has $O(\frac{h_{max}na(n)}{c}) \implies O(\frac{n\sqrt{a(n)}}{c})$ backbone links per channel, and each node appears on $O(\frac{1}{\sqrt{a(n)}}) \implies O(\frac{n\sqrt{a(n)}}{c})$ (from Eqn. 9) source backbones. ∎

Detour backbones:

From Lemma 18 the number of additional flows traversing a cell due to detour routing is only $O(\log^6 n)$, and each such flow will at most traverse the cell twice. Thus detour flows do not pose any significant load-balancing issue at any cell, and we can grow the backbones in $\mathcal{S}$ for these flows in any manner possible, i.e. by assigning links to any eligible node/channel (at least one eligible node is guaranteed to exist).

Additional last hop: We now account for the possible additional last hop that some flows may have, yielding an additional cell in $\mathcal{S}$ (in addition to those traversed by the straight-line from source to pseudo-destination). We already argued that at most $O(na(n)) \implies O(\frac{n\sqrt{a(n)}}{c})$ flows (from Eqn. 8) enter any cell on their additional last hop. Thus, even if their backbone links are assigned to the same channel/node, we would still have $O(\frac{n\sqrt{a(n)}}{c})$ flows per node and channel in any cell for the $\mathcal{S}$ stage.

b) Expanding backbone to $\mathcal{D}_b - \mathcal{S}$: In this stage $\mathcal{B}_p(x)$ expands into the cells traversed by flows for which x is the destination. Note that by our routing strategy a flow will only attempt to switch to the destination backbone when it enters *ready-for-transition* mode. From Lemma 20, the number of flows traversing a cell in *ready-for-transition* mode is $O(\log^6 n)$, which is negligible compared to the total number of traversing flows. Thus flows on their destination backbone do not pose any major load-balance issues,

and the backbones can be expanded into cells of $\mathcal{D} - \mathcal{S}$ by assigning links to any eligible node/channel.

B. Balancing Load within a Cell

Per-Channel Load:

Lemma 22: The number of flows that enter any cell on a given channel is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.

Proof: A flow on route $\mathcal{D}_1, \mathcal{D}_2, ..., \mathcal{D}_{j-1}, \mathcal{D}_j....$ may enter a cell $\mathcal{D}_j$ on a channel i if (1) the flow is in *progress-on-source-backbone* mode, or it is in *ready-for-transition* mode, but is yet to find a transition into the destination backbone , and i is the shared channel between the source backbone nodes in $\mathcal{D}_{j-1}$ and $\mathcal{D}_j$ (2) the flow has already made a transition, and i is the shared channel between the destination backbone nodes in $\mathcal{D}_{j-1}$ and $\mathcal{D}_j$

We first consider the flows that enter a cell in *progress-on-source-backbone* mode, i.e., are proceeding on their source backbones. Recall that these are all non-detour-routed flows, since detour-routed flows are always in *ready-for-transition* mode. Then the number of such flows that traverse any cell on a single channel is $O(\frac{n\sqrt{a(n)}}{c})$ from Lemma 21.

Thus, for flows in *progress-on-source-backbone* mode, we can say that no single channel has more than $O(\frac{n\sqrt{a(n)}}{c})$ flows entering on it in any cell.

We now need to account for the fact that some of these flows may be in the *ready-for-transition* mode. From Lemma 20 there are $O(\log^6 n)$ flows traversing any cell in *ready-for-transition* mode w.h.p. (recall that these include the detour-routed flows, and the possible additional last D'D hop). Thus regardless of whether they are still on their source backbone, or have already made the transition to their destination backbone, no channel can have more than $O(\log^6 n)$ such flows entering the cell.

Hence the number of flows entering on a single channel is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p. for each cell of the network.
∎

Lemma 23: The number of flows that leave any given cell on a given channel is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.

Fig. 5.   Two additional transition links lying wholly within the cell

Proof: The flows that leave a cell fall into two categories (1) those that originate at some node in the cell (2) those that entered the cell but did not terminate there (i.e. were relayed through the cell). The former can be no more than the number of nodes in the cell, i.e. $O(na(n)) = O(\frac{\log n}{p_{rnd}}) = O(\log^2 n)$ (from Eqn. 7). For the latter, note that the flows that leave the cell, must then enter one of the 8 adjacent cells on that channel (as the corresponding backbone link for a flow leaves the current cell, and enters an adjacent cell). Thus, flows leaving the cell on a channel can be no more than 8 times the maximum number of flows entering a cell on any one channel, which has been established as $O(\frac{n\sqrt{a(n)}}{c}) = O(\sqrt{\frac{n\log n}{c}})$ in Lemma 22. Hence, the total number of flows leaving any given cell on a given channel is $O(\frac{n\sqrt{a(n)}}{c}) + O(\log^2 n) \implies O(\frac{n\sqrt{a(n)}}{c})$ w.h.p. ∎

Lemma 24: The number of additional transition links scheduled on a channel within any cell is $O(\log^6 n)$ w.h.p.

Proof: Recall that transition strategy outlined in the proof of Lemma 16, whereby the flow locates a cell along the route where the source backbone node q_x, and destination backbone node q_y are connected through a third node z. This yields two additional links $q_x \to z$, and $z \to q_y$ that lie entirely within the cell (Fig. 5). Note that the number of flows performing this transition in the cell can be no more than the number of flows traversing the cell in *ready-for-transition* mode. From Lemma 20 there are $O(\log^6 n)$ such flows traversing any cell w.h.p. In the worst case, we can count 2 additional links for each such flow as being all assigned to one channel. The result thus proceeds. ∎

a) Per-Node Load:

Lemma 25: The number of flows that are assigned to any one node in any cell is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.

27

Proof: A node is always assigned the single flow for which it is the source. A node is also assigned flows terminating in that cell and for which it is the destination, and from Lemma 11 there are at most $D(n) = O(na(n))$ such flows for any node w.h.p. Besides, a node may be assigned flows that are in the *ready-to-transition* mode, for which it facilitates a transition (if it is a *transition faciliator* node), or on whose destination backbone it figures. There are $O(\log^6 n)$ such transitioning flows in a cell w.h.p. from Lemma 20. Thus a node can only have $O(log^6 n)$ such flows assigned.

We now consider the flows in *progress-on-source-backbone* mode that do not originate in the cell. These nodes are on their source-backbone, and by Lemma 21, all nodes have at most $O(\frac{n\sqrt{a(n)}}{c})$ such flows assigned each. Thus, the resultant number of assigned flows per node is $1 + D(n) + O(\log^6 n) + O(\frac{n\sqrt{a(n)}}{c}) \implies O(\frac{n\sqrt{a(n)}}{c})$. ∎

C. Transmission Schedule

As mentioned earlier, from the Protocol Model assumption, each cell can face interference from at most a constant number β of nearby cells. Thus, if we consider the resultant cell-interference graph, it has a chromatic number at most $1 + \beta$. Hence, we can come up with a global schedule having $1 + \beta$ unit time slots in each round. In any slot, if a cell is active, then all interfering cells are inactive. The next issue is that of intra-cell scheduling. We need to schedule transmissions so as to ensure that at any time instant, there is at most one transmission on any given channel in the cell. Besides, we also need to ensure that no node is expected to transmit or receive more than one packet at any time instant.

We construct a conflict graph based on the nodes in the active cell, and its adjacent cells (note that the hop-sender of each flow shall lie in the active cell, and the hop-receiver shall lie in one of the adjacent cells), as follows: we create a separate vertex for each flow traversing the cell. Since the flow has an assigned channel on which it operates in that particular hop, each vertex in the graph has an implicit asociated channel. Besides, each vertex has an associated pair of nodes corresponding to the hop endpoints. Two vertices are connected by an edge if (1) they have the same associated channel, or (2) at least one of their associated nodes is the same. The scheduling problem thus reduces to obtaining a vertex-coloring of this graph. If we have a vertex coloring, then it ensures that (1) a node is never simultaneously sending/receiving for more than one flow (2) no two flows on the same channel are active simultaneously. Thus, the number of neighbors of a graph vertex is upper bounded by the number of flows entering/leaving the active cell on that channel, and the number of flows assigned to the flow's two

hop endpoints (both hop-sender and hop-receiver). Thus, it can be seen from Lemmas 22, 23, 24 and 25 that the degree of the conflict graph is $O(\frac{n\sqrt{a(n)}}{c}) + O(\frac{n\sqrt{a(n)}}{c}) + O(\log^6 n) + O(\frac{n\sqrt{a(n)}}{c}) = O(\frac{n\sqrt{a(n)}}{c})$ (note that $O(\log^6 n) \implies O(\frac{n\sqrt{a(n)}}{c})$), since we showed in Eqn. 8 that $\frac{n\sqrt{a(n)}}{c} = \Omega(\sqrt{\frac{n}{\log n}})$). Thus the graph can be colored in $O(\frac{n\sqrt{a(n)}}{c})$ colors.

Thus the cell-slot (which can be assumed to be of unit time) is divided into $O(\frac{n\sqrt{a(n)}}{c}) = O(\frac{\sqrt{\frac{n\log n}{p_{rnd}}}}{c})$ equal length subslots, and all traversing flows get a slot for transmission. This implies that each flow gets a $\Omega(c\sqrt{\frac{p_{rnd}}{n\log n}})$ fraction of the time. Also recall that each cell gets at least one slot in $1 + \beta$ slots, where β is a constant, and each channel has bandwidth $\frac{W}{c}$. Thus each flow gets a throughput of at least $\left(\frac{1}{1+\beta}\right)\left(\frac{W}{c}\right)\Omega(c\sqrt{\frac{p_{rnd}}{n\log n}}) = \Omega(W\sqrt{\frac{p_{rnd}}{n\log n}})$.

We thus obtain the following theorem:

Theorem 7: With a (c, f)-random channel assignment, where $c = O(\log n)$: whenever c, f satisfy: $f > 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$, the network capacity is $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$ per flow.

V. CONCLUSION

We have presented a tight bound for capacity with random (c, f) assignment ($c = O(\log n)$), whenever c, f satisfy: $f > 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$. This still leaves a small gap in terms of what optimal capacity is when $f \leq 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$. Note that our earlier results in [1] have shown that the capacity is $\Omega(W\sqrt{\frac{f}{cn\log n}})$ and $O(W\sqrt{\frac{p_{rnd}}{n\log n}})$ for all $2 \leq f \leq c(c = O(\log n))$, and these bounds continue to apply. We conjecture that even when $f \leq 10(1 + \log\frac{192}{5} + \log\frac{c}{f})$ too, the tight capacity bound would yield $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$.

REFERENCES

[1] V. Bhandari and N. H. Vaidya, "Connectivity and capacity of multichannel wireless networks with channel switching constraints," Technical Report, CSL, UIUC (revised Sept. 2006 (original version dated August 2006)).

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

[3] M. Mitzenmacher and E. Upfal, *Probability and computing*. Cambridge University Press, 2005.

[4] A. E. Gamal, J. P. Mammen, B. Prabhakar, and D. Shah, "Throughput-delay trade-off in wireless networks." in *Proceedings of IEEE INFOCOM*, 2004.

[5] "Hall's marriage theorem," Online at http://planetmath.org/encyclopedia/HallsMarriageTheorem.html.

[6] D. Dubhashi and D. Ranjan, "Balls and bins: a study in negative dependence," *Random Struct. Algorithms*, vol. 13, no. 2, pp. 99–124, 1998.

[7] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, W. M. McEneany, G. Yin, and Q. Zhang, Eds. Boston: Birkhauser, 1998, pp. 547–566.

[8] C. Fragouli, J. Widmer, and J.-Y. L. Boudec, "On the benefits of network coding for wireless applications," in *4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2006.

[9] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. MIT Press, 1990.

[10] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal, "Balanced allocations," *SIAM J. Comput.*, vol. 29, no. 1, pp. 180–200, 2000.

[11] R. Motwani, R. Panigrahy, and Y. Xu, "Fractional matching via balls-and-bins," in *RANDOM*, 2006.