Capacity of Multi-Channel Wireless Networks with Random (c, f) Assignment*

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

ABSTRACT

With the availability of multiple unlicensed spectral bands, and potential cost-based limitations on the capabilities of individual nodes, it is increasingly relevant to study the performance of multi-channel wireless networks with channel switching constraints. To this effect, some constraint models have been recently proposed, and connectivity and capacity results have been formulated for networks of randomly deployed single-interface nodes subject to these constraints. One of these constraint models is termed random (c, f) assignment, wherein each node is pre-assigned a random subset of f channels out of c (each having bandwidth $\frac{W}{c}$), and may only switch on these. Previous results for this model established bounds on network capacity, and proved that when $c = O(\log n)$, the *per-flow* capacity is $O(W\sqrt{\frac{p_{rnd}}{n\log n}})$ and $\Omega(W\sqrt{\frac{f}{cn\log n}})$ (where $p_{rnd} = 1-(1-\frac{f}{c})(1-\frac{f}{c-1})...(1-\frac{f}{c-f+1}) \geq 1-e^{-\frac{f^2}{c}}$). In this paper we present a lower bound construction that matches the previous upper bound. This establishes the capacity as $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$. The surprising implication of this result is that when $f = \Omega(\sqrt{c})$, random (c, f) assignment yields capacity of the same order as attainable via unconstrained switching. The routing/scheduling procedure used by us to achieve capacity requires synchronized route-construction for all flows in the network, leading to the open question of whether it is possible to achieve capacity using asynchronous procedures.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Performance, Theory

*This research is supported in part by NSF grant CNS 06-27074, US Army Research Office grant W911NF-05-1-0246, and a Vodafone Graduate Fellowship.

Keywords

Wireless Networks, Capacity, Multiple Channels, Switching Constraints, Random (c, f) Assignment

1. INTRODUCTION

There has been much recent interest in exploiting the availability of multiple channels in wireless networks. The transport capacity of such networks has also been studied under various assumptions on availability/capability of radio-interfaces.

It was shown in [7] that for a single-channel single-interface scenario, in a randomly deployed network, per-flow capacity scales as $\Theta(\frac{W}{\sqrt{n\log n}})$ bits/s under a Protocol Model of interference, and that if the available bandwidth W is split into c channels, with each node having a dedicated interface per channel, the results remain the same.

While many existing standards, e.g., IEEE 802.11a, 802.11b, 802.15.4 allow for multiple channels, nodes are typically hardware-constrained and have much fewer interfaces. This issue was studied in [9], under a model where nodes were capable of switching their interface(s) to any channel. It was shown that given c available channels of bandwidth $\frac{W}{c}$ each, and $1 \leq m \leq c$ interfaces per node, capacity depends solely on the ratio $\frac{c}{m}$. For a random network, and the Protocol Model, three capacity regions were established. Most relevant to our work is the region $\frac{c}{m} = O(\log n)$, where they showed that capacity is the same as for $m = c$, i.e., $\Theta(\frac{W}{\sqrt{n\log n}})$ bits/s per-flow.

In [3], a case was made for the need to study the performance of multi-channel networks in situations where there are constraints on channel switching. This study was motivated on the basis of future low-cost transceiver designs involving limited tunability, as well as cognitive radio networks. As more spectrum becomes freely available for unlicensed use, cost concerns are very likely to lead to situations where individual nodes can operate only over a much smaller spectral range, and may possess heterogeneous capabilities. Thus it is quite relevant to study the impact of switching constraints, and attempt to quantify it.

Some constraint models were proposed in [3] to capture some expected constraints, and two such models were analyzed, viz., adjacent (c, f) assignment and random (c, f) assignment. The impact of restricted switching was quantified by the parameter f (where f is the number of channels an individual node may switch to). Results were presented for the regime $c = O(\log n)$. It was established that per-flow capacity is $\Theta(W\sqrt{\frac{f}{cn\log n}})$ for adjacent (c, f) assignment. For random (c, f) assignment, an upper bound of $O(W\sqrt{\frac{p_{rnd}}{n\log n}})$ and a lower bound of $\Omega(W\sqrt{\frac{f}{cn\log n}})$ were established.

In this paper, we establish that the per-flow capacity with random (c, f) assignment (under the Protocol Model of interference) for the regime $c = O(\log n)$ $(2 \leq f \leq c)$ is $\Theta(W\sqrt{\frac{p_{rnd}}{n \log n}})$ by presenting a capacity-achieving lower bound construction. It can be shown that $p_{rnd} \geq 1 - e^{-\frac{f^2}{c}}$. Thus, the somewhat surprising implication of this result is that when $f = \Omega(\sqrt{c})$, random (c, f) assignment yields capacity of the same order as attainable via unconstrained switching. Hence $\sqrt{c}$-switchability is sufficient to make order-optimal use of all c channels.

Interestingly, our capacity achieving routing/scheduling procedure requires that all routes be computed in lock-step. This leaves open the question of whether capacity can be achieved via asynchronous routing/scheduling procedures.

2. NOTATION AND TERMINOLOGY

Throughout this paper, we use the following standard asymptotic notation [4]:

- $f(n) = O(g(n))$ 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 [7], we say that the per flow network throughput is $\lambda(n)$ if each flow in the network can be guaranteed a throughput of at least $\lambda(n)$ with probability 1, as $n \to \infty$.

Whenever we use log without explicitly specifying the base, we imply the *natural* logarithm.

3. NETWORK MODEL

We consider a network of n *single-interface* nodes deployed uniformly at random over a unit torus. Each node is the source of exactly one flow. As in [7], each source S selects a destination by first fixing on a point D' uniformly at random, and then picking the node D (other than itself), that is closest to D'. The total bandwidth (data-rate) available is W, and it is divided into c channels of equal bandwidth $\frac{W}{c}$, where $c = O(\log n)$. We assume that $c \geq 2$, as $c = 1$ implies that $f = 1$ is the only possibility, which yields the degenerate single-channel case. We also assume $2 \leq f \leq c$. A justification for not allowing $f = 1$ for $c \geq 2$ is given in [3], [1], where it was shown that for the random (c, f) model (and also the adjacent (c, f) model described in [3]), $f = 1$ and $c \geq 2$ leads to zero capacity, as some flow will get no throughput w.h.p.

4. RELATED WORK

It was shown by Gupta and Kumar [7] that, for a single-channel single-interface scenario, the per flow capacity in a random network scales as $\Theta(\frac{W}{\sqrt{n \log n}})$ bits/s. The throughput-delay trade-off was studied in [6], and it was shown that the optimal trade-off is given by $D(n) = \Theta(nT(n))$ where $D(n)$ is delay, and $T(n)$ is throughput. In the multi-channel context, an interesting scenario arises when the number of interfaces m at each node may be smaller than the number of available channels c. This issue was analyzed in [9] and it was shown that the capacity results are a function of the channel-to-interface ratio $\frac{c}{m}$. It was also shown that in the random network case, there are three distinct capacity regions: when $\frac{c}{m} = O(\log n)$, the per-flow capacity is $\frac{W}{\sqrt{n \log n}}$, when $\frac{c}{m} = \Omega(\log n)$ and also $O\left(n\left(\frac{\log\log n}{\log n}\right)^2\right)$, the per flow capacity is $\Theta(W\sqrt{\frac{m}{nc}})$, and when $\frac{c}{m} = \Omega\left(n\left(\frac{\log\log n}{\log n}\right)^2\right)$, the per-flow capacity is $\Theta(\frac{Wm\log\log n}{c \log n})$. Connectivity and capacity of multi-channel wireless networks with channel switching constraints were considered in [3]. Results were presented for two specific constraint models, viz., adjacent (c, f) assignment and random (c, f) assignment. It was shown that when $c = O(\log n)$, capacity with adjacent (c, f) assignment scales as $\Theta(W\sqrt{\frac{f}{cn \log n}})$. For random (c, f) assignment, it was shown that capacity is $O(W\sqrt{\frac{p_{rnd}}{n \log n}})$ and $\Omega(W\sqrt{\frac{f}{cn \log n}})$. In this paper, we show that the capacity for this model is actually $\Theta(W\sqrt{\frac{p_{rnd}}{n \log n}})$.

5. RANDOM (c, f) ASSIGNMENT

In this section we briefly describe the random (c, f) assignment model first introduced in [3, 1], and summarize some already proven results that will be useful in proving the lower bound on capacity. 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 $p_{rnd} = 1 - (1 - \frac{f}{c})(1 - \frac{f}{c-1})...(1 - \frac{f}{c-f+1})$. The proofs of the following are available in [1], and also [2]:

LEMMA 1. *For $c \geq 2$ and $2 \leq f \leq c$, the following holds:*

$$\frac{cp_{rnd}}{f} \leq \min\{\frac{c}{f}, 2f\} \quad (1)$$

LEMMA 2. $\min\{\frac{c}{f}, 2f\} \leq \sqrt{2c}$

5.1 Sufficient Condition for Connectivity

It was stated and proved in [3] that, for random (c, f) assignment, if $\pi r^2(n) = \frac{800\pi \log n}{p_{rnd} n}$, then the network is connected w.h.p. We summarize the proof idea here, to provide important context for the results in this paper.

The unit torus is divided into square cells of area $a(n) = \frac{100 \log n}{p_{rnd} n}$. It can be shown that there are at least $\frac{50 \log n}{p_{rnd}}$ nodes in each cell w.h.p. $r(n)$ is set to $\sqrt{8a(n)}$. Within each cell, $\frac{2 \log n}{p_{rnd}}$ nodes are chosen uniformly at random, and set apart as *transition facilitators*. At least $\frac{48 \log n}{p_{rnd}}$ nodes remain in each cell, and they 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 some adjacent cell, there is no backbone candidate node with which it can communicate is less than $(1 - p_{rnd})^{\frac{48 \log n}{p_{rnd}}} \leq \frac{1}{e^{48 \log n}} = \frac{1}{n^{48}}$. Applying union bounds over all 8 adjacent cells of a node, and 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 $\frac{8}{n^{47}}$.

Associated with each node x, there is a set of nodes $\mathcal{B}(x)$ called the backbone for x. $\mathcal{B}(x)$ is constituted as follows: Cells already covered by the 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, till all cells are filled. Based on previous arguments, $\mathcal{B}(x)$ eventually covers all cells with probability at least $1-\frac{8}{n^{47}}$. For any node-pair x and y, if $\mathcal{B}(x)\cap\mathcal{B}(y)\neq\phi$ the two are obviously connected. Suppose the two backbones are disjoint. Then x and y are still connected if there exists a cell where the member of $\mathcal{B}(x)$ (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 z. q_x and q_y can communicate directly with probability 1 if they share a common channel. Thus the case to handle is that where no cell has q_x and q_y sharing a channel.

If they do not share a common channel, consider the event that there exists a third node amongst the *transition facilitators* in the cell through whom they can communicate. Thus, 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. This is elaborated further.

Consider a third node z amongst the transition facilitators in the same cell as q_x and q_y. Consider a situation where z enumerates its f channels in some uniformly random order, and then inspects the first two channels, checking whether the first one is common with q_x, and checking whether the second one is common with 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 do so is $p_{xy}\geq 1-(1-p_z)^{\frac{2\log n}{p_{rnd}}}$.

Thereafter application of the union bound over all cells, and all node pairs suffices to prove the result.

6. SUMMARY OF OUR RESULTS

In the rest of this paper, we describe a construction that achieves a per-flow throughput of $\Omega(W\sqrt{\frac{p_{rnd}}{n\log n}})$ for $c=O(\log n)$. In light of the upper bound of $O(W\sqrt{\frac{p_{rnd}}{n\log n}})$ proved in [3], this establishes the capacity for random (c,f) assignment as $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$ in the regime $c=O(\log n)$. It is easy to see the following:

$$\begin{aligned} p_{rnd} &= 1-(1-\frac{f}{c})(1-\frac{f}{c-1})\ldots(1-\frac{f}{c-f+1}) \\ &\geq 1-(1-\frac{f}{c})^f \geq 1-e^{-\frac{f^2}{c}} \end{aligned} \quad (2)$$

Hence: $f=\Omega(\sqrt{c}) \implies p_{rnd}=\Omega(1)$. To illustrate, if we set $f=\sqrt{c}$, $p_{rnd}\geq 1-\frac{1}{e}>\frac{1}{2}$. In light of Eqn. (2), our result implies that $f=\Omega(\sqrt{c})$ suffices for achieving capacity of the same order as the unconstrained switching case [9]. For $f=\sqrt{c}$, the previously established lower bound of $\Omega(W\sqrt{\frac{f}{cn\log n}})$, would have yielded a capacity degradation by a factor of the order of $c^{\frac{1}{4}}$ compared to the unconstrained switching case. In general, one may see that the capacity may diverge from the previous lower bound when $\frac{f}{c}\to 0$, but $f\to\infty$.

Fig. 1 is a numerical plot (obtained by setting c to 10^4, and varying f from 2 to c) depicting how the probability p_{rnd} compares with the probability $p_{adj}^{max}=\min\{\frac{2f-1}{c-f+1},1\}$. Recall that p_{rnd} is the probability that two nodes share at least one channel in random (c,f) assignment, and p_{adj}^{max} is the upper bound on the probability that two nodes share at least one channel in adjacent (c,f) assignment [3]. It is quite remarkable that though both models allow nodes to switch between a subset of f channels, the additional degrees of freedom obtained via the random assignment model lead to a much quicker convergence of p_{rnd} toward 1. The results in [3] established that connectivity was the dominant constraint determining capacity for adjacent (c,f) assignment in the $c=O(\log n)$ regime. The lower bound in this paper for random (c,f) assignment matches the upper bound imposed by the connectivity constraint (see [3, 1]). Thus, the quick convergence of p_{rnd} to 1 leads to a quicker convergence of capacity towards that attainable via unconstrained switching.

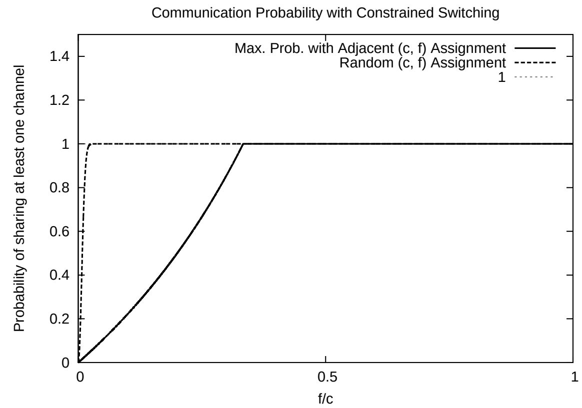

Figure 1: Comparison of probability of sharing a channel

It is to be noted that the lower bound of [3] was obtained using a much simpler construction than the one described in this paper. Thus the two constructions represent an interesting trade-off in capacity versus scheduling/routing complexity.

7. SOME USEFUL RESULTS

THEOREM 1. *(Chernoff Upper Tail Bound [10]) Let $X_1,\ldots,X_n$ be independent Poisson trials, where $Pr[X_i=1]=p_i$. Let $X=\sum_{i=1}^{n}X_i$. Then, for $0<\beta\leq 1$:*

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

LEMMA 3. *For all $0\leq x\leq 1:(1-x)\leq e^{-x}$.*

LEMMA 4. *Suppose we are given a unit torus with n points (or nodes) located uniformly at random, and the region is subdivided 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 $(100\alpha(n)+50)\log n$ points (or nodes), with high probability.*

LEMMA 5. *Suppose we are given a unit torus with n points (or nodes) located uniformly at random. Let us consider the set of all circles of radius R and area $A(n)=\pi R^2$ on the unit torus. 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 $(100\alpha(n)+50)\log n$ points (or nodes), with high probability.*

LEMMA 6. *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 $L_i = P_iQ_i$ satisfies the condition that no cell has more than $n\sqrt{a(n)}$ lines passing through it.*

THEOREM 2. *(Hall's Marriage Theorem [8], [11]) Given a set $\mathcal{S}$, let $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, \ldots \mathcal{T}_n\}$ be a finite system of subsets of $\mathcal{S}$. Then $\mathcal{T}$ possesses a system of distinct representatives if and only if for each k in $1,2,..,n$, any selection of k of the sets $\mathcal{T}_i$ will contain between them at least k elements of $\mathcal{S}$. Alternatively stated: for all $\mathcal{A} \subseteq \mathcal{T}$, the following is true: $|\cup\mathcal{A}| \geq |\mathcal{A}|$*

LEMMA 7. *The number of subsets of size k chosen from a set of m elements is given by $\binom{m}{k} \leq \left(\frac{me}{k}\right)^k$.*

THEOREM 3. *(Integrality Theorem [4]) If the capacity function of a network flow graph takes on only integral values, then the maximum flow x produced by the Ford-Fulkerson method has the property that $|x|$ is integer-valued. Moreover, for all vertices u and v, the value of $x(u,v)$ is an integer.*

8. LOWER BOUND ON CAPACITY

A lower bound of $\Omega(W\sqrt{\frac{f}{cn\log n}})$ for capacity with random (c, f) assignment was proved in [3, 1]. From Lemma 1, it follows that $\frac{\sqrt{\frac{f}{cn\log n}}}{\sqrt{\frac{p_{rnd}}{n\log n}}} = \Omega(\frac{1}{\sqrt{f}})$. Thus for $f < 100$, $\frac{\sqrt{\frac{f}{cn\log n}}}{\sqrt{\frac{p_{rnd}}{n\log n}}} = \Omega(1)$, and the construction presented in [3] (details in [1]) is asymptotically optimal. Thus, we propose to use this construction for $f < 100$ to achieve capacity.

We now present a construction that achieves $\Omega(W\sqrt{\frac{p_{rnd}}{n\log n}})$ when $f \geq 100$ (thus necessarily $c \geq 100$).

Traffic-model related results. We first state some results for the traffic model of [7] (which is also used in this paper). For proofs, please see [2].

LEMMA 8. *The number of flows for which any node is the destination is $O(\log n)$ w.h.p.*

LEMMA 9. *For large n, at least one node is the 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.*

Subdivision of network region into cells. We use a square cell construction (similar to that used in [6], and subsequently in [9], [3]). 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* [7], 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 in other cells within a distance $(2+\Delta)r(n)$ from some point in that cell. Since Δ is independent of n, the number of cells that interfere with 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 10. *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. The proof has been omitted due to space constraints. Please see [2]. □

Many of the intermediate results in the rest of this paper assume that the high-probability event of Lemma 10 holds.

We also state the following facts:

$$\frac{f}{c} \leq p_{rnd} \leq 1 \tag{4}$$

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

$$f(n) = O(na(n)) \implies f(n) = O(\frac{n\sqrt{a(n)}}{c}) \tag{5}$$

$$f(n) = O(\frac{1}{\sqrt{a(n)}}) \implies f(n) = O(\frac{n\sqrt{a(n)}}{c}) \tag{6}$$

Some properties of $SD'D$ routing. Recall that we use the traffic model of [7], where each source S first chooses a pseudo-destination D', and then selects the node D nearest to it as the actual destination. In [7], 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. In our case, it may not always suffice to use $SD'D$ routing (we elaborate on this later). However, this is still an important component of our routing procedure, and so we state and prove the following lemmas (similar results were stated in [1]) for $SD'D$ routing:

LEMMA 11. *Given only straight-line SD' routing (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. Let us consider the straight-line part SD' of an $SD'D$ route. 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 [6] (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 $Pr[X_i^k = 1] = a(n) = \frac{250\max\{\log n, c\}}{p_{rnd}n}$.

Let $X_i = \sum_{k=1}^{n} X_i^k$. Then $E[X_i] = na(n)$. Also, for a given i, the X_i^k's are independent [6]. Then by application of the Chernoff bound from Thereom 1 (with $\beta = \frac{1}{4}$):

$$\begin{aligned} Pr[X_i \geq \frac{5E[X_i]}{4}] &\leq \exp(-\frac{E[X_i]}{48}) \\ \therefore Pr[X_i \geq \frac{1250\max\{\log n, c\}}{4p_{rnd}}] & \\ \leq \exp(-\frac{250\max\{\log n, c\}}{48p_{rnd}}) &< \frac{1}{n^5} \end{aligned} \tag{7}$$

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. □

LEMMA 12. *The number of flows for which any single node is the destination is $O(na(n))$ w.h.p.*

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

LEMMA 13. *If a node is the destination of some flow, then that flow's pseudo-destination must lie within either the same cell, or an adjacent cell w.h.p.*

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

LEMMA 14. *The number of $SD'D$ routes that traverse any cell is $O(n\sqrt{a(n)})$ w.h.p.*

PROOF. The proof for this lemma is largely based on a proof in [6]. It has been omitted due to space constraints. Please see [2]. □

Having stated and proved these lemmas, we now establish some properties of the spatial distribution of channels, and thereafter describe our scheduling/routing procedure further:

DEFINITION 1. *We define a term M_u where $M_u = \lceil \frac{9fna(n)}{25c} \rceil = \lceil \frac{90f\max\{\log n,c\}}{cp_{rnd}} \rceil$.*

Then the following holds:

LEMMA 15. *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 as* candidates *to examine, then, in each cell, amongst those $\frac{180\max\{\log n,c\}}{p_{rnd}}$* candidate *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. The proof has been omitted due to space constraints. Please see [2]. □

Similar to the construction for connectivity from [3] that we briefly summarized in Section 5.1, we will construct a backbone for each node. *However, since our concern is not merely connectivity but also capacity, these backbones need to be constructed carefully, to ensure that no bottlenecks are formed*.

Conditioning on Lemma 10, there are at least $\frac{200\max\{\log n,c\}}{p_{rnd}}$ nodes in each cell w.h.p. 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*.

DEFINITION 2. (Proper Channel) *A channel i is deemed* proper *in cell $\mathcal{D}$ if it occurs in at least M_u backbone candidate nodes in $\mathcal{D}$.*

LEMMA 16. *For each cell of the network, the following is true w.h.p.: if the number of* proper *channels in the cell is c', then $c' \geq c - \lfloor \frac{f}{4} \rfloor \geq c - \lfloor \frac{c}{4} \rfloor \geq \lceil \frac{3c}{4} \rceil \geq \frac{3c}{4}$.*

PROOF. The proof follows from Lemma 10 and Lemma 15. □

Besides, we can also show the following:

LEMMA 17. [1] *Consider any cell $\mathcal{D}$. Let $\mathcal{W}_i$ be the set of all nodes in the 8 adjacent cells $\mathcal{D}(k), 1 \leq k \leq 8$, that are capable of switching on channel i.*

For a set of nodes $\mathcal{B}$, define $\mathcal{C}(\mathcal{B}) = \{j | j$ proper in $\mathcal{D}$ and $\exists u \in \mathcal{B}$ capable of switching on $j\}$. If $f \geq 100$, the following holds w.h.p.:

$$\forall \text{ channels } i, \forall \mathcal{B} \subseteq \mathcal{W}_i \text{ such that } |\mathcal{B}| = \lceil \frac{fna(n)}{4c} \rceil : \ |\mathcal{C}(\mathcal{B})| \geq \lceil \frac{3c}{8} \rceil$$

This is true for all cells $\mathcal{D}$.

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

[1]This can be viewed as a special variant of the Coupon Collector's problem [10], where there are c different types of coupons, and each box has a random subset of f different coupons. Some other somewhat different variants having multiple coupons per box have been considered in work on coding, e.g., [5].

8.1 Routing and channel assignment

Partial Backbones. 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 in [3]. 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 a flow can be routed along these backbones from its source to its destination.

LEMMA 18. *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 cells (hops) (recall that $2 \leq f \leq c$ and $c = O(\log n)$) , it will have found an opportunity to make the transition w.h.p. If the routes of each of the n flows get to traverse at least $\frac{c^2}{f^2}$ distinct cells (note that each individual route needs to traverse at least so many distinct cells; two different flows may have common cells on their respective routes), 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 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 5.1, we summarized a proof from [3] showing 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 (Lemma 10), 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 $p_{xy} \geq 1-(1-p_z)^{\frac{20\log n}{p_{rnd}}}$.

Hence, the probability that this happens in none of the $\frac{c^2}{f^2}$ distinct cells is at most $(1-p_{xy})^{\frac{c^2}{f^2}} \leq (1-p_z)^{\frac{20c^2\log n}{f^2 p_{rnd}}} < (1-\frac{f^2}{c^2})^{\frac{20c^2\log n}{f^2 p_{rnd}}} \leq e^{-\frac{20\log n}{p_{rnd}}} \leq \frac{1}{n^{20}}$ (from Lemma 3). Applying the union bound over all n flows, the probability that all flows are able to transition is at least $1-\frac{1}{n^{19}}$. □

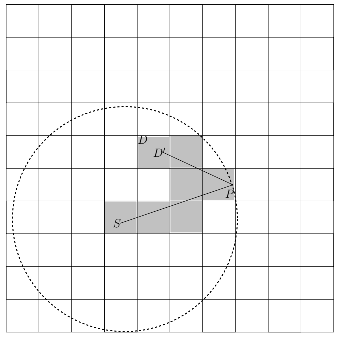

Figure 2: Illustration of detour routing

Hence, we require each route to have at least $\frac{c^2}{f^2}$ distinct hops (note that this is not a tight bound on the minimum number of required hops). Resultantly, we cannot stipulate that *all* flows be routed along the (almost) straight-line path $SD'D$. If $SD'D$ is short, a detour may be required to ensure the minimum route-length, akin to detour-routing in the constructions of [3]. Such flows are said to be *detour-routed*.

Flow Transition Strategy. As per our strategy, a non-detour-routed flow is initially in a *progress-on-source-backbone* mode, and keeps to the source backbone till there are only $\frac{c^2}{f^2}$ distinct intermediate cells 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 comprises $h \geq \frac{c^2}{f^2}$ distinct intermediate cells (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 in the constructions described in [3, 1] (consider a circle of this radius centered around S, choose a point P on the circle, and follow the route $SPD'D$).

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

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

LEMMA 20. *The number of flows traversing any cell is $O(n\sqrt{a(n)})$ w.h.p. even with detour routing.*

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

LEMMA 21. *The number of flows traversing any cell in* ready-for-transition *mode is $O(\log^6 n)$ w.h.p.*

PROOF. The proof has been omitted due to space constraints. Please see [2]. □

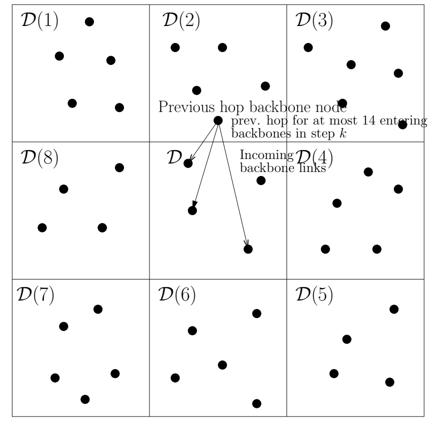

Figure 3: Cell $\mathcal{D}$ and neighboring cells during backbone construction

Backbone Construction. The backbone construction procedure is required to 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 $\mathcal{D}(j), 1 \leq j \leq 8$ (Fig. 3). $\mathcal{B}_p(x)$ is constituted 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 from each source to its pseudo-destination below. Some routes require an additional last hop to reach the actual destination node. However, from Lemma 13, the only such last hop routes that may enter a cell correspond to pseudo-destinations in the 8 adjacent cells. Then applying Lemma 4 to the set of pseudo-destinations, they are only $O(na(n))$ such pseudo-destinations, and thus only $O(na(n))$ such last-hop flows entering the cell. Hence we can account for them separately.

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.

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 $\mathcal{S}$ 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 then consider the cell in $\mathcal{S}$ 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 the flow can enter that cell.

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

The backbones are extended by constructing bipartite graphs that aid load-balance.

LEMMA 22. *If $f \geq 100$, then it is possible to devise a backbone construction procedure, such that, after step h_{max} of the backbone construction procedure for $\mathcal{S}$ (for non-detour-routed flows), each cell has $O(\frac{n\sqrt{a(n)}}{c})$ incoming backbone links on a single channel, and each node appears on $O(\frac{n\sqrt{a(n)}}{c})$ (source) backbones, w.h.p.*

PROOF. This proof assumes that the high probability events in Lemma 10, Lemma 11, Lemma 16, and Lemma 17 occur.

We present an inductive argument. Recall that we are expanding backbones to cover cells in $\mathcal{S}$. *At each step of the (inductive) construction, we first have a channel-allocation phase, followed by a node-allocation phase.* We prove that after step k 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 k. Thus after step k, it appears in a total of $O(14k) \implies O(k)$ backbones.
- *Invariant 2:* No more than $\lfloor \frac{5na(n)}{c} \rfloor$ new backbone links enter the cell on a single channel during step k. Thus, in total $O(\frac{kna(n)}{c})$ incoming backbones (entering the cell) are assigned (incoming links) on a single channel after step k.

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. (6)).

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

If Invariant 1 holds at the end of step $k-1$, then Invariant 2 continues to hold after the channel-allocation phase of step k. If Invariant 2 holds after the channel-allocation phase of step k, then Invariant 1 will continue to hold after the node-allocation phase of step k, and thus both Invariants 1 and 2 will hold at the end of step k.

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 an 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 $k-1$. Consider a particular cell $\mathcal{D}$ during step k.

Let the number of *proper* channels in $\mathcal{D}$ be c'. From Lemma 16, we know that $c' \geq c - \lfloor \frac{f}{4} \rfloor \geq \frac{3c}{4}$ for each cell. Each flow that enters cell $\mathcal{D}$ in step k has a previous hop-node in one of the 8 adjacent cells. Also note that, from Lemma 16, 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 can switch on f channels of which at most $\lfloor \frac{f}{4} \rfloor$ may be non-proper in cell $\mathcal{D}$).

Channel-Allocation. 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 (source) backbones that enter the cell $\mathcal{D}$ in step k. From Lemma 11, it proceeds that $|\mathcal{L}| \leq \lfloor \frac{5na(n)}{4} \rfloor$. 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 $\mathcal{D}$, i.e., $|\mathcal{P}| = c' \lfloor \frac{5na(n)}{c} \rfloor$.

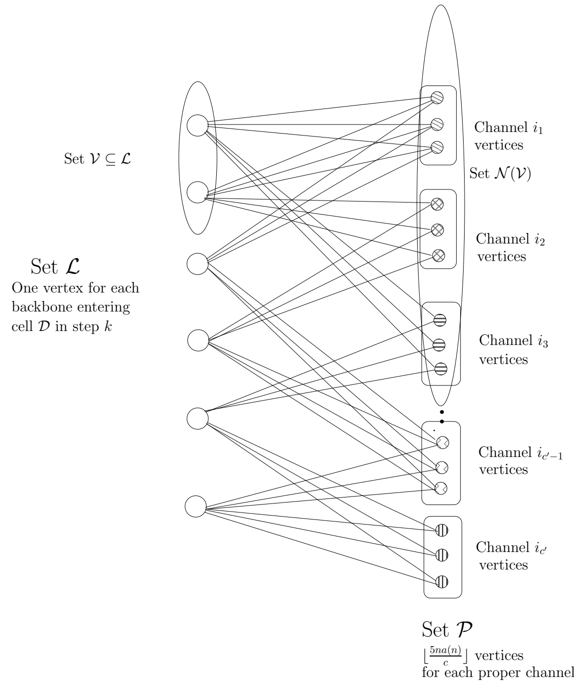

Figure 4: Bipartite Graph for Cell $\mathcal{D}$ in step k

A backbone vertex is connected to all the vertices for the channels proper in $\mathcal{D}$ 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 2). 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 $\mathcal{V}$.

We first note the following:

$$\begin{aligned}\lceil \frac{3f}{4} \rceil \lfloor \frac{5na(n)}{c} \rfloor &\geq \frac{3f}{4}\left(\frac{5na(n)}{c} - 1\right) = \frac{15fna(n)}{4c} - \frac{3f}{4} \\ &\geq \frac{15fna(n)}{4c} - \frac{3fna(n)}{1000c} \geq \frac{29fna(n)}{8c} (\because na(n) \geq 250c)\end{aligned} \tag{8}$$

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 $|\mathcal{N}(\mathcal{V})| \geq \frac{3f}{4}\left(\frac{5na(n)}{c} - 1\right) \geq \frac{29fna(n)}{8c}$ (from Eqn. 8). 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 $\frac{29fna(n)}{8c}$. Note that since Invariant 1 held till end of step $k-1$, no more than 14 backbone links were assigned to any single node in $\bigcup_{k=1}^{8} \mathcal{D}(k)$ in step $k-1$.

Intuitively, in order to show that $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$ for all such $\mathcal{V}$, we first state and prove the observation that if a channel overload condition occurs, resulting in $|\mathcal{N}(\mathcal{V})| < |\mathcal{V}|$ for some $\mathcal{V}$, then the overload must also manifest itself in some *channel-aligned* subset (i.e. a subset where all flows have some *common* proper channel i available to them). Thus, to show that no overload condition occurs, it suffices to show that no overload condition occurs in any of these *critical* channel-aligned subsets, which can be shown using Lemma 17. The argument is formalized as follows:

Let $\mathcal{V}_i$ be the set comprising all sets $\mathcal{U}_i \subseteq \mathcal{L}$, such that all backbone vertices in $\mathcal{U}_i$ have channel i associated with them (i.e., all backbone vertices in $\mathcal{U}_i$ have i available to them as a valid proper channel choice for entering $\mathcal{D}$).

Claim (a). $\forall \mathcal{U} \in \bigcup_{i \text{ proper in } \mathcal{D}} \mathcal{V}_i$:

$$|\mathcal{U}| \geq \lceil \frac{29fna(n)}{8c} \rceil \implies |\mathcal{N}(\mathcal{U})| \geq |\mathcal{L}|$$

Proof of Claim (a): We know that $\mathcal{U} \in \mathcal{V}_i$ for some i that is proper in $\mathcal{D}$. Also, since no node can be the previous hop in step k of more flows than those assigned to it in step $k-1$, and Invariant 1 held after step $k-1$, it proceeds that no previous hop node is common to more than 14 entering backbone links. Let $\mathcal{A}$ be the set of distinct previous hop nodes associated with $\mathcal{U}$. Then $|\mathcal{A}| \geq \frac{1}{14}|\mathcal{U}| \geq \frac{1}{14}(\frac{29fna(n)}{8c}) \geq \frac{fna(n)}{4c} + \frac{fna(n)}{112c} > \frac{fna(n)}{4c} + 1 \geq \lceil \frac{fna(n)}{4c} \rceil$ (note that $\frac{fna(n)}{c} \geq 250f \geq 500 > 112$). Observe that $\mathcal{A}$ thus contains at least one subset $\mathcal{B}$ satisfying $|\mathcal{B}| = \lceil \frac{fna(n)}{4c} \rceil$. Recognizing that all members of $\mathcal{A}$, and hence all members of $\mathcal{B}$, are capable of switching on channel i, we can invoke Lemma 17 on $\mathcal{B}$, to obtain that when $f \geq 100$, $|\mathcal{C}(\mathcal{B})| \geq \lceil \frac{3c}{8} \rceil$. This yields: $\mathcal{N}(\mathcal{U}) \geq |\mathcal{C}(\mathcal{B})| \lfloor \frac{5na(n)}{c} \rfloor \geq |\mathcal{C}(\mathcal{B})| \left(\frac{5na(n)}{c} - 1\right) \geq \lceil \frac{3c}{8} \rceil \left(\frac{5na(n)}{c} - 1\right) \geq \frac{15na(n)}{8} - \lceil \frac{3c}{8} \rceil \geq \frac{15na(n)}{8} - \frac{3}{8}\left(\frac{na(n)}{250}\right) - 1 \geq \frac{5na(n)}{4} \geq |\mathcal{L}|$.

Claim (b). Consider a set $\mathcal{V} \subseteq \mathcal{L}$. Then:

$$|\mathcal{N}(\mathcal{V})| < |\mathcal{V}| \implies \exists i \text{ proper in } \mathcal{D}, \mathcal{S}_i \subseteq \mathcal{V} \text{ s.t. } : \quad \mathcal{S}_i \in \mathcal{V}_i \text{ and } |\mathcal{S}_i| \geq \lceil \frac{29fna(n)}{8c} \rceil \tag{9}$$

Proof of Claim (b): Suppose $|\mathcal{N}(\mathcal{V})| < |\mathcal{V}|$. Let us denote by $\mathcal{S}_i \subseteq \mathcal{V}$ the set of all backbone vertices in $\mathcal{V}$ that are associated with channel i (i.e., have channel i available as a valid proper channel choice for entering cell $\mathcal{D}$). Consider the bipartite sub-graph $G_{\mathcal{V}}$ induced by $\mathcal{V} \cup \mathcal{N}(\mathcal{V})$, and assign all edges unit capacity. Construct the graph $G_{\mathcal{V}} \cup \{s,t\}$ where s is a source node having a unit capacity edge to all vertices $v \in \mathcal{V}$, and t is a sink node, connected to each vertex $u \in \mathcal{N}(\mathcal{V})$ via a unit capacity edge. We try to obtain a (s,t) flow g such that all edges (s,v) are saturated. Each vertex $v \in \mathcal{V}$ sub-divides the unit of flow received from s equally amongst all edges (v,u) outgoing from it. Since each vertex has edges to vertices of at least $\frac{3f}{4}$ channels, this yields at least $\frac{3f}{4}\left(\frac{5na(n)}{c} - 1\right) \geq \frac{29fna(n)}{8c}$ edges (see Eqn. 8). Thus each $v \in \mathcal{V}$ contributes at most $\frac{8c}{29fna(n)}$ units of flow to a vertex $u \in \mathcal{N}(\mathcal{V})$, i.e., $g(v,u) \leq \frac{8c}{29fna(n)}$. Hence no vertex $u \in \mathcal{N}(\mathcal{V})$ gets more than $h(u) = \sum_{v \in \mathcal{S}_i} g(v,u) = \frac{8c|\mathcal{S}_i|}{29fna(n)}$ units of flow, where i is the channel corresponding to vertex u. Resultantly, if $|\mathcal{S}_i| \leq \lfloor \frac{29fna(n)}{8c} \rfloor$ for all channels i that are proper in cell $\mathcal{D}$, this implies that $h(u) \leq 1$, and setting $g(u,t) = h(u)$ yields the desired (s,t) flow. Hence g is a valid flow that allows a unit of flow to pass through each vertex $v \in \mathcal{V}$. From the Integrality Theorem (Theorem 3), we can obtain an integer-capacity flow that yields a matching of size $|\mathcal{V}|$. Thus, from Hall's marriage theorem (Theorem 2), $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$ (else a matching of size $|\mathcal{V}|$ could not have existed). This yields a contradiction. Thus there must exist a proper channel i, and $\mathcal{S}_i \subseteq \mathcal{V}$ such that $\mathcal{S}_i \in \mathcal{V}_i$ and $|\mathcal{S}_i| > \lfloor \frac{29fna(n)}{8c} \rfloor$. Since set-cardinality must necessarily be an integer, it proceeds that $|\mathcal{S}_i| \geq \lceil \frac{29fna(n)}{8c} \rceil$, and Eqn. (9) holds.

Claim (c). $\forall \mathcal{V} \subseteq \mathcal{L}$ such that $|\mathcal{V}| \geq \frac{29fna(n)}{8c} : |\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$

Proof of Claim (c): Suppose $|\mathcal{N}(\mathcal{V})| < |\mathcal{V}|$ for some such $\mathcal{V}$. Then, from Claim (b), there exists a set $\mathcal{S}_i \subseteq \mathcal{V}$ such that $\mathcal{S}_i \in \mathcal{V}_i$, and $|\mathcal{S}_i| \geq \lceil \frac{29fna(n)}{8c} \rceil$. Thus $\mathcal{S}_i$ qualifies as a set to which Claim (a) applies. Invoking Claim (a) on this set $\mathcal{S}_i$, it follows that $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{N}(\mathcal{S}_i)| \geq |\mathcal{L}| \geq |\mathcal{V}|$. This yields a contradiction. Hence: $|\mathcal{N}(\mathcal{V})| \geq |\mathcal{V}|$.

Therefore, by application of Hall's marriage theorem (Theorem 2), 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 each proper 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 [4] on a flow network obtained from the bipartite graph by adding a source with an edge to each vertex in $\mathcal{L}$, a sink to which each vertex in $\mathcal{P}$ has an edge, and assigning unit capacity to all edges.

Thus Invariant 2 continues to hold after the channel-allocation phase of step k.

Node-Allocation. Having determined the channel each backbone should use to enter cell $\mathcal{D}$, we need to assign a node in cell $\mathcal{D}$ 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 $\mathcal{D}$ in step k. The second set (call it $\mathcal{R}$) comprises 14 vertices for each *backbone candidate* node in cell $\mathcal{D}$. 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 in the preceding channel-allocation phase.

Each vertex $x \in \mathcal{F}$ has degree at least $14M_u$, since it is assigned to a *proper* channel, which by definition has at least M_u representatives in cell $\mathcal{D}$, each of which has 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 any set $\mathcal{V} \in \mathcal{F}$.

Since no channel is assigned more than $\lfloor \frac{5na(n)}{c} \rfloor$ entering backbone links in this step, the vertices in $\mathcal{V}$ are cumulatively associated with at least $m \geq \frac{|\mathcal{V}|}{\lfloor \frac{5na(n)}{c} \rfloor}$ 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 $m \geq \frac{|\mathcal{V}|}{\lfloor \frac{5na(n)}{c} \rfloor}$ proper channels is at least $\frac{|\mathcal{V}|M_u}{f\lfloor \frac{5na(n)}{c} \rfloor} \geq \frac{|\mathcal{V}|\lceil \frac{9fna(n)}{25c} \rceil}{\frac{5fna(n)}{c}} \geq \frac{9|\mathcal{V}|}{125}$, and as each node has 14 vertices, it follows that $|\mathcal{N}(\mathcal{V})| \geq 14\left(\frac{9|\mathcal{V}|}{125}\right) \geq \frac{126|\mathcal{V}|}{125} > |\mathcal{V}|$.

Then 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 the backbone corresponding to vertex x in cell $\mathcal{D}$ (the matching can again be computed via the Ford-Fulkerson method). Since there are at most 14 vertices associated with a node, no node is assigned more than 14 incoming backbone links in step k, and Invariant 1 continues to hold after the node-allocation phase of step k.

Thus we have shown that both Invariants 1 and 2 continue to hold after step k.

Hence after step h_{max} (where $h_{max} \leq \frac{2}{\sqrt{a(n)}}$), each cell $\mathcal{D}$ has $O(\frac{h_{max} n a(n)}{c}) \implies O(\frac{n\sqrt{a(n)}}{c})$ entering backbone links per channel, and each node appears on $O(h_{max}) = O(\frac{1}{\sqrt{a(n)}}) \implies O(\frac{n\sqrt{a(n)}}{c})$ (from Eqn. (6)) source backbones. □

Detour backbones: From Lemma 19 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 since, as a consequence of Lemma 16, each node can switch on at least $\lceil \frac{3f}{4} \rceil$ channels that are proper in the next cell).

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 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. (5)) 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.

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 21, the total number of flows traversing a cell in *ready-for-transition* mode is $O(\log^6 n)$ (counting possible repeat traversals), which is much smaller than $O(\frac{n\sqrt{a(n)}}{c})$. Thus flows on their destination backbone do not pose any major load-balance issues, and the backbones can be expanded into cells of $\mathcal{D}_b - \mathcal{S}$ by assigning links to any eligible node/channel.

8.2 Proving load-balance within a cell

We now show that no channel or interface bottlenecks form in the network when our described construction is used.

Per-Channel Load

LEMMA 23. *The number of flows that enter any cell on any single channel is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.*

PROOF. A flow on route $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_{j-1}, \mathcal{D}_j \ldots$ 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 $\mathcal{D}_j$, or (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 enter any cell on a single channel is $O(\frac{n\sqrt{a(n)}}{c})$ from Lemma 22.

We now need to account for the fact that some of these flows may be in the *ready-for-transition* mode. From Lemma 21 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 with their repeat traversals counted separately, 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 $O(\frac{n\sqrt{a(n)}}{c}) + O(\log^6 n) \implies O(\frac{n\sqrt{a(n)}}{c})$ w.h.p. for each cell of the network. □

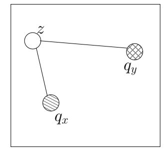

Figure 5: Two additional transition links for a flow lying wholly within the cell

LEMMA 24. *The number of flows that leave any cell on any single channel is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.*

PROOF. 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 $O(\frac{n\sqrt{a(n)}}{c})$ in Lemma 23. Hence, the total number of flows leaving any cell on a single channel is also $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p. □

LEMMA 25. *The number of additional transition links scheduled on any single channel within any cell is $O(\log^6 n)$ w.h.p.*

PROOF. Recall the transition strategy outlined in the proof of Lemma 18, 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 $q_x \rightarrow z$, and $z \rightarrow q_y$ 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 21 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 proceeds from this observation. □

Per-Node Load

LEMMA 26. *The number of flows that are assigned to any single node in any cell is $O(\frac{n\sqrt{a(n)}}{c})$ w.h.p.*

PROOF. A node is always assigned the single flow for which it is the source. A node is also assigned flows for which it is the destination, and from Lemma 8 there are at most $D(n) = O(\log n)$ 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 facilitator* 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 21. 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 from Lemma 22, each node has at most $O(\frac{n\sqrt{a(n)}}{c})$ such flows assigned. Thus, the resultant number of assigned flows per node is $1 + D(n) + O(\log^6 n) + O(\frac{n\sqrt{a(n)}}{c}) \implies O(\frac{n\sqrt{a(n)}}{c})$. $\square$

8.3 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 (a graph with a vertex for each cell, and an edge between two vertices if the corresponding cells can interfere with each other), 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, except for transition links, for which both lie in the active cell) as follows: we create a separate vertex for each flow for which a node in the cell needs to transmit data (we count repeat traversals or additional transition links as distinct flows for the purpose of scheduling; these have been accounted for while bounding the number of flows in a cell in previous lemmas). Since the flow has an assigned channel on which it operates in that particular hop, each vertex in the graph has an implicit associated 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 requiring a transmission in 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). It can be seen from Lemma 23, Lemma 24, Lemma 25 and Lemma 26 that the degree of the conflict graph is $O(\frac{n\sqrt{a(n)}}{c}) + O(\frac{n\sqrt{a(n)}}{c}) + O(\log^6 n) + O(\frac{n\sqrt{a(n)}}{c}) + O(\frac{n\sqrt{a(n)}}{c}) = O(\frac{n\sqrt{a(n)}}{c})$ (note that $O(\log^6 n) \implies O(\frac{n\sqrt{a(n)}}{c})$, since $\frac{n\sqrt{a(n)}}{c} = \Omega(\sqrt{\frac{n}{\log n}})$). Thus the graph can be colored in $O(\frac{n\sqrt{a(n)}}{c})$ colors.

Hence, 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. Moreover, 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 4. *When $c = O(\log n)$, and $2 \le f \le c$, the per-flow network capacity with random (c, f) assignment is $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$.*

9. CONCLUSION

We have established the capacity of a random network with random (c, f) assignment, for $c = O(\log n)$, $2 \le f \le c$. Our result indicates that capacity is $\Theta(W\sqrt{\frac{p_{rnd}}{n\log n}})$. Thus, when $f = \Omega(\sqrt{c})$, one can achieve capacity of the same asymptotic order as with unconstrained switching. There still remain some interesting open questions pertaining to the random (c, f) model, in terms of what is achievable via strictly asynchronous routing/scheduling. However, the results in this paper, along with prior results in [3], have been able to demonstrate that it may be possible to achieve good throughput characteristics even when devices are subject to switching constraints. Designing practically feasible and efficient protocols for networks of devices with constrained switching ability is still an open and interesting problem domain.

10. REFERENCES

[1] V. Bhandari and N. H. Vaidya. Connectivity and Capacity of Multi-Channel Wireless Networks with Channel Switching Constraints. Technical Report, CSL, UIUC, January 2007.

[2] V. Bhandari and N. H. Vaidya. Capacity of Multi-Channel Wireless Networks with Random (c, f) Assignment. Technical Report, CSL, UIUC, June 2007.

[3] V. Bhandari and N. H. Vaidya. Connectivity and Capacity of Multi-Channel Wireless Networks with Channel Switching Constraints. In *Proceedings of IEEE INFOCOM*, Anchorage, Alaska, May 2007.

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

[5] C. Fragouli, J. Widmer, and J.-Y. Le 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.

[6] A. E. Gamal, J. P. Mammen, B. Prabhakar, and D. Shah. Throughput-Delay Trade-off in Wireless Networks. In *Proceedings of IEEE INFOCOM*, 2004.

[7] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, IT-46(2):388–404, March 2000.

[8] P. Hall. On Representatives of Subsets. *J. London Math. Soc.*, s1-10(37):26–30, 1935.

[9] P. Kyasanur and N. H. Vaidya. Capacity of Multi-channel Wireless Networks: Impact of Number of Channels and Interfaces. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, pages 43–57. ACM Press, 2005.

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

[11] H. Perfect. The Mathematics of AGMs. *The Mathematical Gazette*, 53(383):13–19, Feb. 1969.