

Robust Multi-Channel Wireless Networks

Arnold L. Rosenberg
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
Email: rsnbrg@cs.umass.edu

Abstract—One wants to deploy an n -node multi-channel wireless network in an environment that is inaccessible for repairs and/or that contains malicious adversaries. One wants to design the network to be *robust* in the following strong sense. Even if any set of $m < n$ nodes is disabled, one still wants all of the surviving $n - m$ nodes to be able to communicate with one another. We present a mathematical model for multi-channel wireless network that facilitates the problem of designing such networks with desired properties. We then present a *scalable, deterministic* design strategy that produces robust networks that are: (a) within a factor of 2 of optimal in size and complexity; (b) power-efficient, in that nodes attempt to communicate with nearer neighbors before trying more remote ones.

I. THE ROBUST-NETWORK PROBLEM

Advances in technology allow us to deploy increasingly complex unstructured (“*ad hoc*”) wireless networks in increasingly remote and vulnerable locations. These locations may render nodes inaccessible for repairs and/or may contain malicious adversaries who strive to disable the network. Such unfriendly environments raise the need for networks to be *robust*, in that they are able to function correctly even when nodes are disabled by happenstance (e.g., a hostile climate) or by malice. This paper develops a *scalable, deterministic* strategy for designing and organizing nodes in a network in a way that endows the network with the following very strong level of robustness. We provide the following guarantee. Say that the network has n nodes when deployed. If any $m < n$ of the network’s nodes are disabled, then all of the surviving $n - m$ nodes will still be able to communicate with one another. Our designs thus preserve the strongest possible level of internode connectivity in the face of a node-disabling attack, achieving for node-based attacks the kind of resilience sought by, e.g., [7] (which focuses on attacks that “cut” links). Moreover, we achieve this robustness with a network organization that is within a factor of 2 of optimal in complexity, using a measure of complexity that is appropriate for our model.

Our designs take account also of the fact that wireless networks usually lack *a priori* structure, which raises the need for *topology control*: the imposition of a logical structure on the network, that assists in routing internode communications. We facilitate topology control by allowing one to impose any of a large family of logical topologies onto a network. Of course, the richness of the available family depends on the complexity of the design.

In Section II, we describe both the design strategy that is our main contribution and the mathematical construct—*hypergraphs* [5]—that allows us to be precise in both definitions and analyses. Section III develops the analyses that substantiate our claims concerning the quality of our designs.

Related work. A comprehensive survey of algorithmic studies of ad hoc networks appears in [24]. Topics there that touch on the present study include topology control and routing—albeit using different models than we propose here. Studies such as [16], [17], [18], [21], [25] aim to increase the throughput of wireless networks by exploiting the multiple channels available the IEEE 802.11 standards. Not surprisingly, the issue of robustness in the face of threats to disable both the nodes of networks and their communication links has received much study. One finds in [7] a novel approach to measuring robustness in the face of threats to network links, as well as an extensive review of the relevant literature. Most studies of robustness in the face of threats to network nodes focus on demonstrating how *non*-robust various Internet-related graphs/networks are; exemplary studies that expose the low level of robustness in such networks appear in [8], [12], [14], [19], [22], [32]. In [2], one finds a distinctive—and important—notation of robustness, measured via network *diameter* (maximum internode distance) rather than degree of connectivity.

Several studies share our goal of achieving network robustness by guaranteeing embeddability of all small paths. Two notable studies, the nonconstructive design in [3] and its realization in [1], propose designs with less total channel capacity than ours, but at the expense of weaker notions of robustness: they embed sizable paths into surviving nodes only with high probability and only when at most a prespecified fraction of the nodes are disabled. (For our designs, the probability of success and the “prespecified fraction” are both 1.) Moreover, their constructive designs involve some randomness in the construction of the network and/or the association of nodes with communication channels. Kindred to [1], [3], but aimed specifically at trees are the designs in [4], [13]. In [29], we survey the costs and benefits of various approaches to robustness via embeddability of all small paths. The series of papers [9], [27], [28], [29] introduced strong universality and designed efficient strongly universal hypergraphs for a broad range of graph families, defined either by topology (e.g., paths and binary trees) or by ease of separation.

Many studies employ hypergraph-based models with an

eye toward topology control, i.e., to achieve networks that “contain” topologically rich families of graphs. One finds in [31] a broadly applicable strategy for designing optimally small hypergraphs into which all other small hypergraphs can be embedded efficiently. A main result in [23] determines the smallest I-hypergraph that “contains” the n -node clique \mathcal{K}_n . In [6], an n -node hypergraph is designed into which one can embed any n -node binary tree. Mesh-structured processor arrays with buses (hyperedges) as well as point-to-point communication links (edges) are studied in [30]; a major conclusion is that physical considerations favor buses that are “contiguous,” i.e., are *interval* hypergraphs. Tangentially related is the study in [11], which models a wireless network as a (directed) hypergraph, but uses hyperedges to model single-step multicasts rather than channels for establishing point-to-point connections.

A final note: We use geographic proximity information gleaned via GPS (as suggested indirectly in [24]) to assign network nodes to communication channels. Alternative clustering criteria leave open the question of how nodes can group themselves without external assistance; this question is addressed in [20] via distributed agreement protocols.

II. OUR DESIGN STRATEGY

A. Overview

Our design strategy assumes that the nodes of our wireless network have access to a large collection of channels, which is consistent with the IEEE 802.11 standards. We advocate using these channels to establish point-to-point communication links. In common with sources such as those discussed in [24], that, using some mechanism such as GPS, the network’s nodes know their geographical locations. We advocate exploiting that knowledge in two notable ways:

- *to establish a node’s “affinity groups”* based on internode distances; each such group will control one communication channel; This use of location information can enhance power conservation by enabling a node to communicate with distant nodes only if no nearby node is responsive.
- *to endow nodes with “names” that enable a natural ordering among the nodes.* The bit representations of location coordinates “tell” nodes u and v which one precedes the other in the ordering. This use of location information can obviate algorithmically expensive tie-breaking procedures that distributed systems often require; cf. [20].

For each power of 2, $n = 2^\ell$, we design an n -node network \mathcal{N}_n . We endow \mathcal{N}_n with n communication channels; each node has access to¹ $\log n$ channels; one channel is available to all n nodes, two channels to disjoint sets of $n/2$ nodes each, and so on, ending with $n/2$ channels available to mutually disjoint sets of 2 nodes each. On average, each channel is available to $\log n$ nodes.

¹All logarithms have base 2.




Fig. 1. The 8-node wireless network \mathcal{N}_8 : filled squares are network nodes; “clouds” are channels along which nodes communicate; digits are channel names.

Channel assignments proceed via a type of “inorder” tree traversal [10]. Using physical proximity information gleaned via, say, GPS, each node of \mathcal{N}_n affiliates with other nodes, resulting in a distributed self-assignment of nodes to channels.

- 1) The “leader” node, v_0 —which is the node with the smallest name—chooses its closest neighbor, v_1 , as its co-occupant in channel c_1 , which becomes a 2-node channel.
- 2) v_0 then adds its two next closest nodes, v_2 and v_3 , to share channel c_2 —which becomes a 4-node channel—with it and v_1 . Having been so chosen, v_2 and v_3 appropriate channel c_3 as their 2-node channel.
- 3) v_0 then adds its four next closest nodes, v_4 , v_5 , v_6 , and v_7 , to share channel c_4 —which becomes an 8-node channel—with it, v_1 , v_2 , and v_3 . Having been so chosen, nodes v_4 , v_5 , v_6 , and v_7 determine which of them has the smallest name—say that it is v_4 . v_4 chooses its closest neighbor, say v_5 , as its co-occupant in channel c_4 , which becomes a 2-node channel, and it tells v_6 and v_7 to co-locate in channel c_5 , which becomes a 2-node channel.
- 4) Channel assignments continue in the indicated way: node v_0 establishes channels of increasing sizes and empowers the new residents of each new channel to employ this procedure recursively to establish their smaller channels.

Fig. 1 illustrates the preceding process with the ($n = 8$)-node instantiation of network \mathcal{N}_n .

Our prescribed communication protocol has a node access its “size”- k channel as a communication medium only after it has failed to communicate successfully using all of its smaller channels. Thus, if nodes share channels based on physical proximity, the protocol expends battery power only to the extent necessary to preserve connectivity.

We show in subsequent sections that the design exemplified by Fig. 1 achieves our goals.

B. A Model for Multi-Channel Wireless Networks

Hypergraphs and Embeddings. We model a (wireless) network \mathcal{N} formally as a *hypergraph* [5]. A network \mathcal{N} thus has a set N of *nodes* and a *multiset*² C of *channels*, which are sets of nodes; in graph-theoretic terminology, our channels are *hyperedges*. A network node could be a sophisticated computer or a sensor with only modest capabilities; no matter which

²In contrast to a set, a *multiset* can contain multiple instances of each of its elements.




Fig. 2. The network \mathcal{N}_8 of Fig. 1 depicted as an I-hypergraph.

point in this range the node sits at, it has a communication module that allows it to communicate with other nodes—specifically, those it shares a channel with. As Fig. 1 suggests, a channel can contain any number of nodes, and a node can belong to any number of channels.

Two special genres of hypergraph play a role in our study.

- An *Interval hypergraph* (I-hypergraph, for short) [29] is a hypergraph \mathcal{N} :
 - whose node-set has the form $\{0, 1, \dots, k\}$ for some integer $k \geq 0$;
 - each of whose hyperedges has the form $\{h, h + 1, \dots, h + \ell\}$ for some integers $h, \ell \geq 0$ (so, is an *interval*).

We remarked earlier that we advocate having the nodes of a wireless network impose an ordering on themselves (using their unique node-names). Thereby, any channel-based wireless network can organize itself into an I-hypergraph. Fig. 2 depicts the network \mathcal{N}_8 of Fig. 1 visualized as an I-hypergraph. The reader can easily extrapolate from Figs. 1 and 2 to networks/I-hypergraphs \mathcal{N}_n with node-sets of arbitrary power-of-2 cardinalities.

- A *graph* is a hypergraph in which: each hyperedge—now called an *edge*—has precisely two nodes; distinct edges are distinct sets (they can share one node but not two).

Complexity measures. Let \mathcal{N} be a network with node-set $N_{\mathcal{N}}$ and hyperedge-set $C_{\mathcal{N}}$. For the purpose of analyses, we characterize \mathcal{N} in terms of the following parameters:³

- The *channel complexity* of \mathcal{N} , $|C_{\mathcal{N}}|$: the number of channels that \mathcal{N} possesses
- The *size* of channel c , $|c|$: the number of nodes that c contains
- The (total) *SIZE* of \mathcal{N} , $SIZE(\mathcal{N})$: the sum of the sizes of \mathcal{N} 's channels
- The *degree* of $v \in N_{\mathcal{N}}$: the number of channels that v belongs to

Communication and Embeddings. We view a network \mathcal{N} as employing *point-to-point* communications, in the following sense. Focus on an arbitrary node u of \mathcal{N} , which belongs to channels c_1, c_2, \dots, c_k . At any time, node u can send/receive a message to/from a single node from each channel c_i (a total of k messages). This model endows \mathcal{N} 's nodes with an unusual degree of communication power, probably more than can be realized in practice. However, accommodating this much power means that our connectivity guarantees will certainly

extend to any network having more modest communication power. Formally, this communication model means that a communication among \mathcal{N} 's nodes is a special form of graph “contained in” \mathcal{N} 's underlying hypergraph. We make this assertion precise.

An *embedding* of the graph $\mathcal{G} = (N_{\mathcal{G}}, C_{\mathcal{G}})$ into the hypergraph $\mathcal{H} = (N_{\mathcal{H}}, C_{\mathcal{H}})$ comprises a pair of one-to-one mappings, $\langle \mu_n, \mu_c \rangle$:

- The *node mapping* $\mu_n : N_{\mathcal{G}} \rightarrow N_{\mathcal{H}}$ associates a unique (host) node of \mathcal{H} with each (guest) node of \mathcal{G} .
- The *channel mapping* $\mu_c : C_{\mathcal{G}} \rightarrow C_{\mathcal{H}}$ associates an *appropriate* unique channel/hyperedge of \mathcal{H} with each edge of \mathcal{G} .

A channel mapping μ_c is “*appropriate*” if, for each edge $(u, v) \in C_{\mathcal{G}}$, the nodes $\mu_n(u)$ and $\mu_n(v)$ both belong to the hyperedge $\mu_c(u, v)$; symbolically, $\{\mu_n(u), \mu_n(v)\} \subseteq \mu_c(u, v)$.

We say that a network/hypergraph *contains* any graph that is embeddable into it.

C. Strong Universality: Robustness Formalized

Our notion of network robustness resides in the following notion of *strong universality*.

Let Γ be a finite family of graphs. The hypergraph $\mathcal{H} = (N_{\mathcal{H}}, C_{\mathcal{H}})$ is *strongly universal* for Γ if, given any set $W \subseteq N_{\mathcal{H}}$: for every graph $\mathcal{G} = (N_{\mathcal{G}}, C_{\mathcal{G}})$ in Γ for which $|N_{\mathcal{G}}| \leq |W|$, there is an embedding of \mathcal{G} into \mathcal{H} with $\mu_n(N_{\mathcal{G}}) \subseteq W$.

The intended interpretation. The network \mathcal{N} represents a wireless network whose nodes intercommunicate along \mathcal{N} 's channels. The graph family Γ represents a collection of desired communication patterns. Each $\mathcal{G} \in \Gamma$ represents a specific point-to-point communication pattern, which can be realized in \mathcal{N} just when one can assign a unique channel of \mathcal{N} to each of \mathcal{G} 's edges: this assignment is embodied in our notion of embedding.

Under the preceding representation, the mapping μ_n assigns logical nodes to functioning physical ones, while μ_c assigns communication links to the channels that will realize them. The appropriateness condition assures that any pair of nodes that are supposed to use a channel c both have access to that channel; the one-to-one condition assures that a channel is used to realize at most one edge, modeling our assumption that each channel is dedicated to a single communication.

Motivating our notion of “strong universality:” the set $W \subseteq N_{\mathcal{H}}$ contains the operational network nodes, while nodes in $W - N_{\mathcal{H}}$ are disabled; one wants to realize the communication pattern \mathcal{G} on the good nodes of \mathcal{H} .

Our main results, informally. Our particular focus is on the family $\Gamma = \Pi_n$ whose elements are all *path(-graphs)* having n or fewer nodes: the m -node path is the graph \mathcal{P}_m whose node-set is $N_{\mathcal{P}_m} = \{1, 2, \dots, m\}$ and whose edge-set is $C_{\mathcal{P}_m} = \{(i, i + 1) \mid i \in \{1, 2, \dots, m - 1\}\}$.

The main results of this paper are:

- (a) Each I-hypergraph \mathcal{N}_n is strongly universal for Π_n .
- (b) \mathcal{N}_n is within a factor of 2 of optimal in SIZE among I-hypergraphs that are strongly universal for Π_n .

³We denote by $|S|$ the cardinality of the set S .

III. \mathcal{N}_n 's OPTIMAL STRONG UNIVERSALITY FOR Π_n

We specify the I-hypergraphs \mathcal{N}_n of Fig. 2 formally. For each power of 2, $n = 2^\ell$:

- \mathcal{N}_n 's nodes form the set $N_{\mathcal{N}_n} = \{0, 1, \dots, n-1\}$.
- \mathcal{N}_n 's channels/hyperedges are as follows.
For each $k \in \{0, 1, \dots, \ell\}$, $C_{\mathcal{N}_n}$ contains 2^k disjoint channels of size $n/2^k$:

$$\{\{an/2^k, an/2^k + 1, \dots, (a+1)n/2^k - 1\} \mid a = 0, 1, \dots, 2^k - 1\}$$

One easily generalizes this construction to parameters n that are not powers of 2.

Theorem 1: (a) For all n , any I-hypergraph that contains the SIZE- $(n \cdot \log n)$ n -node I-hypergraph \mathcal{N}_n as a sub-I-hypergraph is strongly universal for the family Π_n .

(b) Any n -node network that is strongly universal for the family Π_n must have $\text{SIZE} > \frac{1}{2}n \cdot \log n$.

Proof: (a) We show that \mathcal{N}_n is strongly universal for Π_n , whence, trivially, any super-I-hypergraph of \mathcal{N}_n is. To embed any path-graph \mathcal{P}_m with $m < n$ into \mathcal{N}_n : (1) use the identity node-mapping,

$$\mu_n(v) = v \quad \text{for all } v \in \{1, 2, \dots, n\};$$

(2) use the channel-mapping μ_c that maps each edge $(u, v) \in C_{\mathcal{P}_m}$ to the smallest channel of \mathcal{N}_n that contains both u and v . By construction, this channel exists and is unique.

We verify by induction that the pair of mappings $\langle \mu_n, \mu_c \rangle$ does, indeed, embed \mathcal{P}_m into \mathcal{N}_n . Say, for induction, that every I-hypergraph \mathcal{N}_p with $p < n$ is strongly universal for Π_p . Say further that the m nodes $\{v_0, v_1, \dots, v_{m-1}\}$ of \mathcal{N}_n are operational. Let $\widehat{\mathcal{N}}$ be the smallest sub-I-hypergraph of \mathcal{N}_n whose node-set contains all of the v_i . For simplicity, rename the nodes of $\widehat{\mathcal{N}}$ as $0, 1, \dots, 2^{q+1} - 1$, using the monotonic renaming map ρ that gives each node u of $\widehat{\mathcal{N}}$ the new name $\rho(u) \in \{0, 1, \dots, 2^q - 1\}$, while retaining $\widehat{\mathcal{N}}$'s nodes' natural order from \mathcal{N}_n ; this automatically gives each operational v_i a new name, $w_i = \rho(v_i)$. Let \mathcal{N}' and \mathcal{N}'' be, respectively, the “lefthand” and “righthand” sub-I-hypergraphs of $\widehat{\mathcal{N}}$ obtained by removing its “universal” channel $\{0, 1, \dots, 2^{q+1} - 1\}$.⁴ Finally, let index k be such that $w_0 < \dots < w_{k-1} \leq 2^q - 1 < w_k < \dots < w_{m-1}$. By induction:

- one can embed a copy of \mathcal{P}_k into \mathcal{N}' , using nodes w_0, \dots, w_{k-1} ;
- one can embed a copy of \mathcal{P}_{m-k} into \mathcal{N}'' , using nodes w_k, \dots, w_{m-1} .

In order to embed \mathcal{P}_m into $\widehat{\mathcal{N}}$, and, thereby, into \mathcal{N}_n , we need only use the preceding two embeddings, augmented with the use of channel $\{0, 1, \dots, 2^{q+1} - 1\}$ of $\widehat{\mathcal{N}}$ to implement the edge between w_{k-1} and w_k in \mathcal{P}_m . On the one hand, this hyperedge is available to implement this edge because we began our embedding process by removing this channel for the “lefthand” and “righthand” sub-embeddings. On the other hand, this channel is, by definition, the smallest channel

⁴The meaning of “lefthand” and “righthand” here should be clear from Fig. 2.

of \mathcal{N}_n that contains both w_{k-1} and w_k . We thus have the desired embedding.

(b) Say that we are given an arbitrary n -node network \mathcal{N} that is strongly universal for the family Π_n . For each $k \in \{0, 1, \dots, \log n\}$, one must be able to embed the $m_k = 2^{k+1}$ -node path-graph \mathcal{P}_{m_k} into \mathcal{N} , using nodes⁵

$$\begin{aligned} &0, n/2^k - 1, n/2^k, n/2^{k-1} - 1, n/2^{k-1}, \dots \\ &n - n/2^{k-1} - 1, n - n/2^{k-1}, \\ &n - n/2^k - 1, n - n/2^k, n - 1 \end{aligned}$$

of \mathcal{N} as homes for \mathcal{P}_{m_k} 's nodes. Since each edge of \mathcal{P}_{m_k} is embedded in a unique channel, the embedding of \mathcal{P}_{m_k} accounts for an aggregate SIZE of $\geq n + 2^{k+1} - 2$ in \mathcal{N} : 2^k edges of \mathcal{P}_{m_k} are channel-mapped to distinct channels of sizes no smaller than $n/2^k$, and $2^k - 1$ edges are channel-mapped to distinct channels of sizes no smaller than 2.

Since each value of k thus contributes $\Omega(n)$ to $\text{SIZE}(\mathcal{N})$, it appears that we have the desired bound of $\Omega(n \log n)$ for $\text{SIZE}(\mathcal{N})$. However, we must be more careful in our reckoning, because long channels that we used to embed the path $\mathcal{P}_{m_{k_1}}$ into \mathcal{N} for some small k_1 can be reused when we embed the path $\mathcal{P}_{m_{k_2}}$ into \mathcal{N} for some $k_2 > k_1$. We show now, by induction on the path-parameter k , that such reuse can save only a factor of 2 in our bound on $\text{SIZE}(\mathcal{N})$. Call a channel of \mathcal{N} *big* if it contains more than two nodes. We bound from below the sizes of \mathcal{N} 's big channels.

$k = 0$. In order to embed \mathcal{P}_{m_0} into \mathcal{N} , \mathcal{N} must contain one channel of size n . This contributes n to $\text{SIZE}(\mathcal{N})$.

$k = 1$. In order to embed \mathcal{P}_{m_1} into \mathcal{N} , \mathcal{N} must contain two channels of sizes $\geq n/2$ each. One of these channels may be the size- n channel from the case $k = 0$; the other must be new. Thus, this value of k contributes at least an additional $n/2$ to $\text{SIZE}(\mathcal{N})$.

$k = 2$. In order to embed \mathcal{P}_{m_2} into \mathcal{N} , \mathcal{N} must contain four channels of sizes $\geq n/4$ each. One of these channels may be the size- n channel from the case $k = 0$; one may be the size- $(\geq n/2)$ channel that was “new” in the case $k = 1$; the other two must be new. Thus, this value of k contributes at least an additional $n/2$ to $\text{SIZE}(\mathcal{N})$.

Arbitrary $k < \log n$. In order to embed \mathcal{P}_{m_k} into \mathcal{N} , \mathcal{N} must contain 2^k channels of sizes $\geq n/2^k > 2$ each. The cases defined by smaller values of k have mandated \mathcal{N} 's having 2^{k-1} channels that each has size $\geq n/2^k$; hence, we can use these channels to satisfy the current demand for channel capacity. However, even with this reuse of preexisting big channels, this case still calls for 2^{k-1} new channels of sizes $\geq n/2^k$ each to satisfy the current demand. Thus, this value of k contributes at least an additional $n/2$ to $\text{SIZE}(\mathcal{N})$.

Totalling the channel sizes exposed by the preceding analysis, we conclude that $\text{SIZE}(\mathcal{N}) > \frac{1}{2}n \cdot \log n$, as was claimed.

IV. CONCLUSIONS

We have proposed a model for multi-channel wireless networks that allows one simple, scalable, deterministic designs

⁵To enhance legibility, we list nodes with repetitions; in fact, each node appears once in the set.




Fig. 3. An I-hypergraph that is strongly universal for binary trees having ≤ 8 nodes.

of networks that are robust in the face of any malicious attack that strives to disable the network by disabling some its nodes. Our designs guarantee that all surviving nodes maintain connectivity with one another.

The structural regularity of our designs suggests that our strategy can be extended to preserve even more structure in a network, and this is, indeed, the case. In [28], we derived SIZE-optimal I-hypergraphs that are strongly universal for binary trees; the n -node version of these I-hypergraphs, whose structure is illustrated for the case $n = 8$ in Fig. 3, have SIZE $\Theta(n \log^2 n)$.

In [29], we extrapolated from our topology-specific studies in [27], [28] to construct I-hypergraphs that are strongly universal for any family of graphs that is defined in terms of its separation characteristics; these designs are often SIZE-optimal and are never far from that. Clearly, the more structure one wants to guarantee to preserve in the face of attacks, the greater the necessary expenditure of resources. It remains to be seen if there are other approaches to this problem that represent other points along the cost-benefit axis.

ACKNOWLEDGMENTS

This research was supported in part by NSF Grant CCF-0342417. It is a pleasure to thank George Bissias for helpful comments.

REFERENCES

- [1] N. Alon and F.R.K. Chung, "Explicit constructions of linear-sized fault-tolerant networks." *Discr. Math.* 72 (1988) 15–19.
- [2] R. Albert, H. Jeong, A. Barabasi, "Error and attack tolerance of complex networks." *Nature* 406 (2000) 378–382.
- [3] J. Beck, "On size Ramsey number of paths, trees, and circuits, I." *J. Graph Th.* 7 (1983) 115–129.
- [4] J. Beck, "On size Ramsey number of paths, trees, and circuits, II." Manuscript (1983).
- [5] C. Berge, "*Graphs and Hypergraphs.*" North-Holland, 1973.
- [6] S.N. Bhatt and C.E. Leiserson, "How to assemble tree machines." In *Advances in Computing Research 2*, (F.P. Preparata, ed.) JAI Press, Greenwich, CT (1984) 95–114.
- [7] G.D. Bissias, B.N. Levine, A.L. Rosenberg, "Bounding damage due to link destruction with applications to the Internet" (Poster) *ACM SIGMETRICS'07* (2007).

- [8] B. Bollobas and O. Riordan, "Robustness and vulnerability of scale-free random graphs." *Internet Mathematics* 1 (2003) 1–35.
- [9] F.R.K. Chung and A.L. Rosenberg, "Minced trees, with applications to fault-tolerant VLSI processor arrays." *Math. Syst. Th.* 19 (1986) 1–12.
- [10] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein (1999): *Introduction to Algorithms* (2nd Edition). MIT Press, Cambridge, Mass.
- [11] S. Deb, M. Effros, T. Ho, D.R. Karger, R. Koetter, D.S. Lun, M. Medard, N. Ratnakar, "Network coding for wireless applications: a brief tutorial." *Intl. Wkshp. on Wireless Ad-hoc Networks* (2005)
- [12] A. Flaxman, A. Frieze, J. Vera, "Adversarial deletion in a scale free random graph process." *16th Ann. ACM-SIAM Symp. on Discrete Algorithms* (2005).
- [13] J. Friedman and N. Pippenger, "Expanding graphs contain all small trees." *Combinatorica* 7 (1987) 71–76.
- [14] P. Holme, B.J. Kim, C.N. Yoon, S.K. Han, "Attack vulnerability of complex networks." *Physics Review* (2002).
- [15] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Medard, "Resilient network coding in the presence of byzantine adversaries." *IEEE INFOCOM* (2007).
- [16] N. Jain, S. Das, A. Nasipuri, "A multichannel CSMA MAC protocol with receiver-based channel selection for multihop wireless networks." *IEEE Intl. Conf. on Computer Communications and Networks* (2001).
- [17] P. Kyasanur and N.H. Vaidya, "Routing and interface assignment in multi-channel multi-interface wireless networks." *IEEE Wireless Communications and Networking Conf.* (2005).
- [18] C. Liu, Z. Liu, Y. Liu, H. Zhao, T. Zhao, W. Yan, "A clustering-based channel assignment algorithm and routing metric for multi-channel wireless mesh networks." *5th Intl. Symp. on Parallel and Distributed Processing and Applications (ISPA'07)*. In *Lecture Notes in Computer Science 4742*, Springer-Verlag, New York (2007) 832–843.
- [19] D. Magoni, "Tearing down the Internet." *IEEE J. Selected Areas in Communications* (2003).
- [20] H. Moser and B. Thallner, "Construction of a fault-tolerant wireless communication topology using distributed agreement." *Wkshp. on Dependability Issues in Wireless Ad Hoc Networks and Networks* (2006) 35–44.
- [21] A. Nasipuri, J. Zhuang, S. Das, "A multichannel CSMA MAC protocol for multihop wireless networks." *IEEE Wireless Communications and Networking Conf.* (1999).
- [22] S.-T. Park, A. Khrabrov, D.M. Pennock, S. Lawrence, C.L. Giles, L.H. Ungar, "Static and dynamic analysis of the Internet's susceptibility to faults and attacks." *IEEE Conf. on Computer Communications* (2003).
- [23] G.L. Peterson and Y.-H. Ting, "Trade-offs in VLSI for bus communication networks." Tech. Rpt. 111, Univ. Rochester (1982).
- [24] R. Rajaraman, "Topology control and routing in ad hoc networks: a survey." *SIGACT News* 33 (June, 2002) 60–73.
- [25] A. Raniwala, K. Gopalan, T. Chiueh, "Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks." *ACM SIGMOBILE Review* 8, (2004) 50–65.
- [26] A.L. Rosenberg, "Routing with permuters: Toward reconfigurable and fault-tolerant networks." Tech. Rpt. CS-1981-13, Duke Univ. (1981).
- [27] A.L. Rosenberg, "On designing fault-tolerant VLSI processor arrays." In *Advances in Computing Research 2*, (F.P. Preparata, ed.) JAI Press, Greenwich, CT, (1984) 181–204.
- [28] A.L. Rosenberg, "A hypergraph model for fault-tolerant VLSI processor arrays." *IEEE Trans. Comp., C-34* (1985) 578–584.
- [29] A.L. Rosenberg, "Interval hypergraphs." In *Graphs and Algorithms* (R.B. Richter, ed.) *Contemporary Mathematics* 89, Amer. Math. Soc., Providence (1989) 27–44.
- [30] Q. Stout, "Meshes with multiple buses." *27th IEEE Symp. Foundations of Computer Science* (1986) 264–273.
- [31] L.G. Valiant, "Universality considerations in VLSI circuits." *IEEE Trans. Comp., C-30* (1981) 135–140.
- [32] S. Zhou, and R.J. Mondragon, "Redundancy and robustness of AS-level Internet topology and its models." *Electronics Lets.* (2004).