Network Robustness and Graph Topology

Anthony H. Dekker[†] **Bernard D. Colbert**[‡]

Defence Science and Technology Organisation (DSTO)
Canberra, Australia
† Email: `dekker@ACM.org`

Abstract

Two important recent trends in military and civilian communications have been the increasing tendency to base operations around an internal network, and the increasing threats to communications infrastructure. This combination of factors makes it important to study the robustness of network topologies. We use graph-theoretic concepts of connectivity to do this, and argue that node connectivity is the most useful such measure. We examine the relationship between node connectivity and network symmetry, and describe two conditions which robust networks should satisfy. To assist with the process of designing robust networks, we have developed a powerful network design and analysis tool called CAVALIER, which we briefly describe.

Keywords: Network Centric Warfare, Network robustness, Graph connectivity.

1 Introduction

There have been two important recent trends in both military and civilian communications. The first is network-centric operation, which bases organizational activity strongly around an internal network. In the civilian sphere, this is called e-commerce (and, in more recent developments, m-commerce). In the military sphere, this is called Network Centric Warfare (NCW). To quote Alberts *et al.* (1999):

"*We define NCW as an information superiority-enabled concept of operations that generates increased combat power by networking sensors, decision makers, and shooters to achieve shared awareness, increased speed of command, higher tempo of operations, greater lethality, increased survivability, and a degree of self-synchronization. In essence, NCW translates information superiority into combat power by effectively linking knowledgeable entities in the battlespace.*"

The second trend is the increasing threat to communications infrastructure. In the civilian sphere, the threat is from terrorist attacks, while in the military sphere this comes from the increasing tendency to view communications networks as high-value targets.

The first trend makes networks more important, while the second makes them more vulnerable. This dilemma makes it critically important to address network robustness, i.e. the continued ability of the network to perform its function in the face of attack.

Designers of communications networks must therefore assume that networks will be attacked, and that some of these attacks will result in damage. Robust networks will continue functioning in spite of such damage and outages (up to some level of damage).

In this paper we specifically focus on the robustness of the network *topology.* We use *graph theory* to investigate which network topologies are the most robust. Graph theory provides two different measures of *connectivity* which are possible ways of measuring robustness, and we argue that *node connectivity* is the most useful of these. We examine the relationship between node connectivity and the degree of symmetry of the network, and we suggest that it is important for robust networks to satisfy two conditions which we call *node-similarity* and *optimal connectivity.* We investigate the relationship between these conditions, and describe a number of ways to design robust networks which satisfy them. We explore Cayley graphs, random graphs, scale-free networks, and several areas of mathematics that shed light on robust network design.

To assist with the process of designing robust networks, we have developed a powerful network design and analysis tool called CAVALIER.

2 Graph Connectivity

A natural way to model the topology of a communications network is as an (undirected) *graph* consisting of *nodes* and *links.* For the purposes of analysing topology, we ignore any variation in the type of links. Robustness of the topology will come from the presence of *alternate paths*, which ensure that communication remains possible in spite of damage to the network.

If a graph has n nodes, then we say that the graph has *size* n. If a node has d outgoing links, we say that the node has *degree* d. The minimum degree $d_{\min}$ of the graph is the smallest of the node degrees, and the maximum degree $d_{\max}$ of the graph is the largest of the node degrees.

Definition 2.1 There are two concepts of connectivity for a graph which can be used to model network robustness:

(i) the *node connectivity* κ is the smallest number of nodes whose removal results in a disconnected or single-node graph.

(ii) the *link connectivity* λ is the smallest number of links whose removal results in a disconnected graph.

For example, K_n, the completely connected graph of size n, with each node connected to the $n-1$ others, has $\kappa = \lambda = n - 1$.

When modelling network robustness in the face of equipment failures (particularly for civilian networks) we would expect link connectivity to be the most useful. Random equipment failures, by affecting cables,

‡ Now at Telstra Research Laboratories, Melbourne, Australia.

interfaces, circuit boards, etc. would primarily put links out of action. On the other hand, when modelling network robustness of military networks in the face of combat (and indeed also of civilian networks in the face of terrorist activity), the major threat is the destruction of entire nodes (usually by explosive means). In this case, we would expect node connectivity to be the most useful in modelling robustness. In related work (Dekker 2004) we describe a simple combat simulation experiment which shows that this is, in fact, the case.

The following well-known theorem, due to Menger, provides an alternative formulation of node and link connectivity:

Theorem 2.2 *For any graph:*

(i) the node connectivity κ is the smallest number of node-distinct paths between any two nodes.

(ii) the link connectivity λ is the smallest number of link-distinct paths between any two nodes.

Proof. Corollaries 4.2 and 4.3 of Gibbons (1985) □

These connectivity measures can be calculated using the maximum-flow algorithm (Gibbons 1985), and we have developed a network design and analysis tool called CAVALIER which incorporates these calculations.

The CAVALIER tool also performs statistical and graph-theoretical network analyses, 2-dimensional and 3-dimensional visualisation (Dekker 2001), and has a simulation capability to assess network performance (Dekker 2003b). All the figures in this paper were produced using the CAVALIER tool.

There are well-known bounds on κ and λ:

Theorem 2.3 *For any graph, $\kappa \leq \lambda \leq d_{\min}$.*

Proof. Due to Whitney: see Theorem 5.1 of Harary (1969) or Theorem 2.9 of Gibbons (1985). □

If $\kappa = \lambda = d_{\min}$ for some graph, we say that the graph is *optimally connected*, since the node and link connectivities are as high as possible, i.e. the network is as robust as it could be, given the value of $d_{\min}$. In Section 3, we consider several strategies for designing optimally connected graphs.

3 Optimal Connectivity

Having discussed the importance of connectivity in modelling robustness of a network topology, we now review some results from the graph-theoretic literature relating connectivity of a graph to the degree of symmetry, and we describe several cases of optimally connected graphs, i.e. graphs with $\kappa = \lambda = d_{\min}$.

Definition 3.1 The following concepts of symmetry will be used:

(i) We say that a graph is *regular* if every node has the same degree; in this case we also speak of the degree d of the graph.

(ii) An *automorphism* of a graph is a permutation π of the nodes which preserves links, i.e. $a - b$ is a link if and only if $\pi a - \pi b$ is a link.

(iii) A graph is *node-similar* (more usually, *vertex-transitive*) if for any two nodes a and b there is an automorphism π such that $\pi a = b$.

(iv) A graph is *symmetric* if for any two links $a - b$ and $x - y$ there is an automorphism π such that $\pi a = x$ and $\pi b = y$.

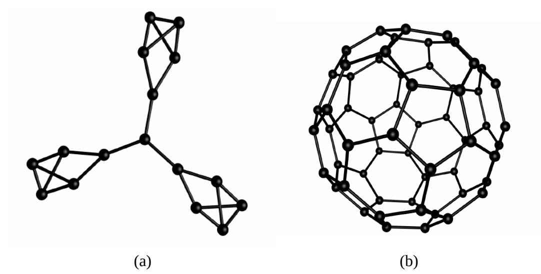

Figure 1: Two Regular Graphs

Our CAVALIER tool includes a module to enumerate all the automorphisms of a graph, and to check the conditions of node-similarity and symmetry (for small or sparse graphs, where this is feasible).

Essentially the condition of node-similarity says that all nodes "look the same," while symmetry says that all links "look the same." Clearly a symmetric graph must be node-similar, and a node-similar graph must be regular, but the converses of these implications do not hold. Figure 1(a) shows a graph which is regular but not node-similar. The "soccer-ball" graph in Figure 1(b) is node-similar (every node is the intersection of a pentagon and two hexagons) but not symmetric (some links join a pentagon and a hexagon, while others join two hexagons). The "soccer-ball" is one of the 13 semi-regular polyhedra described by Archimedes and Kepler, all of which are node-similar (essentially by definition).

Examples of symmetric graphs include the regular polyhedra: tetrahedron ($d = 3$), cube ($d = 3$), octahedron ($d = 4$), dodecahedron ($d = 3$), and icosahedron ($d = 5$). The q-dimensional hypercube (with $n = 2^q$ and $d = q$) is also symmetric, and has been used in a highly parallel computing (van de Goor 1989). The "Connection Machine" (Hillis 1985) was a supercomputer constructed as a 12-dimensional hypercube.

The *torus* (a rectangular $p \times q$ grid with $p, q \geq 3$, where the top edge is connected to the bottom, and the left to the right) is node-similar, and symmetric if $p = q$. The torus has also been used in highly parallel computing, and has the advantage of requiring fewer links than the hypercube.

The property of being node-similar is of value in its own right for parallel computing, since it tends to facilitate routing and load-balancing (van de Goor 1989). For communications network design, node-similarity also facilitates routing, and ensures that the impact of losing a node does not depend on which node is lost. The property of symmetry is also of value, because when links "look the same," traffic on each link is likely to be approximately equal.

Node-similar networks are particularly appropriate for *decentralised* Network Centric Warfare (Dekker 2003a), since all nodes are of equal importance. As military forces move from "platform-centric" to "network-centric" organisational structures, decentralised architectures become more important, since they provide no high-value targets to the enemy. On the other hand, decentralised forces can still focus their attention on a given point, through "swarming" (Arquilla & Ronfeldt 2000) or self-synchronisation (Alberts *et al.* 1999) behaviour.

Within the civilian telecommunications infrastructure, making the major communications backbone node-similar would also have the advantage of providing no high-value targets to terrorist attackers.

If a network is both node-similar and optimally connected, then it provides maximum resistance to node destruction. The literature of graph theory contains some interesting results relating symmetry to connectivity, proved by Mader and Watkins in 1970/71:

Theorem 3.2 *For any connected node-similar graph of degree d:*

(i) $\lambda = d$.

(ii) $\kappa \geq \frac{2}{3}(d+1)$.

(iii) if $d \leq 4$, *then* $\kappa = d$.

(iv) if the graph is symmetric, then $\kappa = d$.

Proof.

(i) Lemma 3.3.3 of Godsil & Royle (2001) or Theorem 3.7 of Babai (1996).

(ii) Theorem 3.4.2 of Godsil & Royle (2001) or Theorem 3.7 of Babai (1996).

(iii) From Theorem 2.3 and (ii).

(iv) Theorem 3.7 of Babai (1996). □

Note that nothing in general can be said about the connectivity of merely regular graphs. For example, the graph in Figure 1(a) has degree 3, but $\kappa = \lambda = 1$. Note also that the bound in (ii) is tight. For example, we can find graphs with $d = 5$ and $\kappa = 4$ (see Figure 3.4 of Godsil & Royle (2001)).

As a consequence of (iii), the "soccer-ball" graph in Figure 1(b) is optimally connected with $\kappa = \lambda = d = 3$ (in fact, all 13 of the Archimedean polyhedra are optimally connected, even the two with $d = 5$). Also by (iii), the torus is optimally connected with $\kappa = \lambda = d = 4$.

We can generalise the torus to the q-dimensional case, i.e. an $m_1 \times m_2 \times \cdots \times m_q$ rectangular grid with each $m_i \geq 3$ and opposite ends connected:

Theorem 3.3 *The q-dimensional hypertorus is optimally connected with* $\kappa = \lambda = d = 2q$.

Proof. Any q-dimensional hypertorus can be made symmetric by merging hyperlayers of adjacent nodes, and $\kappa = \lambda = d = 2q$ for the reduced hypertorus, by Theorem 3.2. For each node-distinct path in the reduced hypertorus, there is a corresponding path in the original, so the result follows by Theorem 2.2. □

An alternative way of designing optimally connected graphs involves *group theory*, an area of abstract algebra with a long tradition:

Definition 3.4 A *group* is a set containing a constant e, and equipped with a binary operator $\oplus$ and a unary operator $\ominus$ such that:

(i) $(x \oplus y) \oplus z = x \oplus (y \oplus z)$.

(ii) $x \oplus e = x$.

(iii) $e \oplus x = x$.

(iv) $x \oplus (\ominus x) = e$.

(v) $(\ominus x) \oplus x = e$.

Definition 3.5 The important group-theoretic concepts for our purpose are *generation* and *Cayley graphs*:

(i) A group G is *generated by* a set S if the elements of G can all be built up using $\oplus$, $\ominus$, e, and the elements of S.

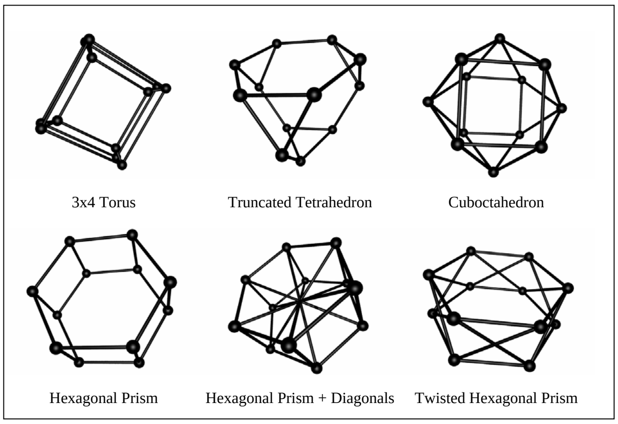

Figure 2: Some Minimal Cayley Graphs with 12 Nodes

(ii) G is *minimally generated by* S if G is generated by S but not by any proper subset of S (in particular, this means that $e \notin S$).

(iii) If G is generated by S, the *Cayley graph* $\Gamma(G, S)$ is the graph whose nodes are the elements of G, and whose links are $x \,—\, s \oplus x$ for every $x \in G$ and $s \in S$.

(iv) If G is minimally generated by S, for some G and S, then $\Gamma(G, S)$ is called a *minimal Cayley graph*.

A simple example of a group is the set $\{0, 1, 2, 3, 4\}$ where:

(i) $e = 0$.

(ii) $x \oplus y = x + y \pmod 5$.

(iii) $\ominus x = 5 - x \pmod 5$.

It is easy to verify that the rules of Definition 3.4 are satisfied, and that the group is minimally generated by the set $\{1\}$. The corresponding minimal Cayley graph is a pentagon with links $0 — 1 — 2 — 3 — 4 — 0$.

There are usually several Cayley graphs (even several minimal Cayley graphs) for a group G, depending on the choice of S. Conversely, different groups may have the same (minimal) Cayley graphs. For example, there are 5 groups containing 12 elements (generally written $\mathbb{Z}_{12}$, $\mathbb{Z}_2 \times \mathbb{Z}_6$, D_6, A_4, and Q_{12}), but at least 7 minimal Cayley graphs with 12 nodes. Six of these are illustrated in Figure 2. For a list of all Cayley graphs with 12 nodes, see Giudici & Abreu (2000).

The importance of minimal Cayley graphs for network design lies in the fact that they are optimally connected (in the sense defined above):

Theorem 3.6 *Let G be a group minimally generated by S, and* $\Gamma(G, S)$ *the corresponding minimal Cayley graph. Then:*

(i) $\Gamma(G, S)$ *is regular and node-similar.*

(ii) $\Gamma(G, S)$ *has degree d equal to the size of the set* $\{x_1, \ldots, x_m, \ominus x_1, \ldots, \ominus x_m\}$, *where* $S = \{x_1, \ldots, x_m\}$.

(iii) $\Gamma(G, S)$ *has* $\kappa = \lambda = d$.

Proof.

(i) Proposition 16.2 of Biggs (1993).

(ii) Considering possible links $y — x_i \oplus y$ and $(\ominus x_j \oplus y) — y$.

Graph	Groups				
	$\mathbb{Z}_{12}$	$\mathbb{Z}_2 \times \mathbb{Z}_6$	D_6	A_4	Q_{12}
3×4 Torus	$\{3, 4\}$ $3^4{=}4^3{=}e$ $3{\oplus}4{=}4{\oplus}3$	$\{a, b, c\}$ $a^3{=}b^2{=}c^2{=}e$ $x{\oplus}y{=}y{\oplus}x$	$\{a^2, a^3, b\}$ $a^6{=}b^2{=}e$ $b{\oplus}a{=}a^5{\oplus}b$		$\{a, b\}$ $a^3{=}b^4{=}e$ $b{\oplus}a{=}a^2{\oplus}b$
Truncated Tetrahedron				$\{a, b\}$ $a^3{=}b^2{=}(a{\oplus}b)^3{=}e$	
Cuboctahedron				$\{a, a{\oplus}b\}$ $a^3{=}b^2{=}(a{\oplus}b)^3{=}e$	
Hexagonal Prism		$\{a, b\}$ $a^6{=}b^2{=}e$ $a{\oplus}b{=}b{\oplus}a$	$\{a, b\}$ $a^6{=}b^2{=}e$ $b{\oplus}a{=}a^5{\oplus}b$		
Hex Prism + Diagonals	$\{2, 3\}$ $3^4{=}e,\ 2^3{=}3^2$ $2{\oplus}3{=}3{\oplus}2$				$\{b, a^2{\oplus}b^2\}$ $a^3{=}b^4{=}e$ $b{\oplus}a{=}a^2{\oplus}b$
Twisted Hex Prism					$\{b, b{\oplus}a\}$ $a^3{=}b^4{=}e$ $b{\oplus}a{=}a^2{\oplus}b$

Table 1: Generators for Some Minimal Cayley Graphs with 12 Nodes

(iii) Theorem 3.7 of Babai (1996) and note following. □

Groups can therefore provide a source of optimally connected networks, by means of minimal generator sets S for particular groups G, producing minimal Cayley graphs $\Gamma(G, S)$. Candidate groups of the required size may be found in textbooks on group theory (Fraleigh 1976, Humphreys 1996), or in Internet resources on group theory (Wilson *et al.* 2003).

The "soccer-ball" graph in Figure 1(b) is, in fact, a minimal Cayley graph for the so-called alternating group on five letters (A_5), generated by the set $\{a, b\}$, where:

(i) $a \oplus a \oplus a \oplus a \oplus a = e$.

(ii) $b \oplus b = e$.

(iii) $a \oplus b \oplus a \oplus b \oplus a \oplus b = e$.

This has degree 3, since $b = \ominus b$. If the generating set $\{b, a \oplus b\}$ is used instead, the resulting minimal Cayley graph corresponds to a different Archimedean polyhedron: the truncated dodecahedron, i.e. a dodecahedron with vertices cut off to form triangular and decagonal (ten-sided) faces.

Our CAVALIER tool includes a module that generates minimal Cayley graphs from group descriptions like these, using *Knuth-Bendix completion* (Baader & Nipkow 1998). Table 1 shows the group descriptions used to generate the graphs in Figure 2. In this table, a^2 refers to $a \oplus a$, a^3 to $a \oplus a \oplus a$, etc. An alternative approach to generating Cayley graphs, using *semi-direct products*, is given in Dineen (1991).

It is also possible to generate optimally connected graphs from other graphs, replacing nodes by completely connected subgraphs. This allows a large optimally connected graph to be constructed from a simpler one:

Theorem 3.7 *Consider an optimally connected regular graph of degree $d \geq 2$, with some or all nodes replaced by completely connected subgraphs of d nodes. Then this graph is also optimally connected and regular with $\kappa = \lambda = d$.*

Proof. It suffices to consider the replacement of a single node by a subgraph. There are d node-distinct paths between nodes in the subgraph (including one path via the rest of the graph), and also d node-distinct paths between nodes in the subgraph and the other nodes (based on the paths in the original graph). □

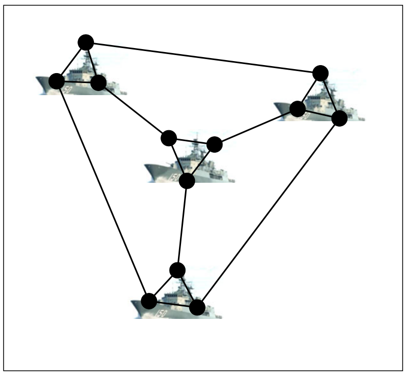

Figure 3: Truncated Tetrahedron Network for CEC

For polyhedral graphs with $d = 3$, this is a process of truncating vertices to form triangular faces (e.g. the truncated dodecahedron or truncated cube). For $d \geq 4$ we refer to the process as *mutilation*.

Networks of this kind are particularly useful for implementing *Cooperative Engagement Capability* (CEC) in the Naval environment (Perry *et al.* 2002). CEC requires a highly robust communications network which can link together sensors and anti-missile systems on different ships. This permits serious threats to be engaged, even when several ships have been seriously damaged. CEC can be implemented using truncated/mutilated networks, with small completely connected fibre-optic subnetworks on each ship, and fast low-latency encrypted point-to-point links between ships. Figure 3 shows such a network based on a truncated tetrahedron. In this case, $\kappa = \lambda = 3$, i.e. the network can survive the loss of any two nodes, while remaining completely connected. Similarly, for eight ships, a truncated cube could be used.

The truncation/mutilation process can be performed repeatedly and/or only on some nodes. The resulting asymmetric networks are also optimally connected, and may be more useful in real-world environments.

To summarise our discussion of optimal connectivity, the graphs in Table 2 are all node-similar and optimally connected with $\kappa = \lambda = d$.

Optimally-connected networks can therefore be designed by finding graphs which are symmetric (Theorem 3.2), or node-similar with degree $d \leq 4$ (Theorem 3.2), or minimal Cayley graphs (Theorem 3.6), or which can be reduced to these cases while preserving connectivity (e.g. Theorems 3.3 and 3.7). Theorem 3.7 also allows us to make local adjustments to networks derived from the other cases.

4 Random Graphs

There is a final way to design (if that is the right word) optimally-connected networks. Surprisingly, it involves making connections at random, i.e. with a fixed probability p of an edge between any pair of nodes (the Erdős-Rényi model). This strategy is successful when the graph is large enough, and is useful militarily for e.g. "swarms" of many low-cost unmanned aerial vehicles or UAVs (Parunak *et al.* 2002). The connections must, however, be made genuinely at random, rather than depending on the physical distances between nodes.

The following result shows that with probability approaching 100%, such random graphs will be optimally connected:

Theorem 4.1 *For any randomly generated graph of size n, the probability that $\kappa = \lambda = d_{\min}$ approaches 1 as $n \to \infty$.*

Proof. Theorem 7.6 of Bollobás (2001) and note following. □

Convergence here is surprisingly rapid. In a simple test of 200,000 random graphs with size n ranging from 7 to 30 and average degree $\sqrt{n}$, we found that the percentage of optimally connected graphs increased from 94.8% for $n = 7$ to 99.98% for $n = 30$. The percentages fitted very closely (with a correlation of 0.97) to the curve $100 - 20e^{-0.2n}$. This result is indicative only, but if extrapolation of this was valid, the percentage of optimally connected graphs for $n = 100$ would be approximately 99.9999999%.

Random graphs are in general not node-similar (indeed, in most cases, as $n \to \infty$, the probability approaches 100% that the only automorphism is the identity: see Theorem 9.3 of Bollobás (2001)). However, for random graphs, nodes are equally important in a statistical sense: since links are placed randomly, no node is privileged by design.

5 Graph Diameter and Link Load

An important graph-theoretic concept is *diameter.* The diameter D of a graph is the longest of all the shortest paths between pairs of nodes. In general, we wish the diameter of networks to be low, since long paths between nodes contribute to both longer message transmission times and greater load on links. Dineen (1991) and Hafner (1995) describe some techniques for constructing graphs of small diameter.

The worst case for the diameter of node-similar graphs is a ring of size n, where the diameter is $\lfloor \frac{n}{2} \rfloor$. The following result sets a bound on the best case for graph diameter:

Theorem 5.1 *For any regular graph of size n, with degree $d \geq 3$ and diameter D:*

(i) $n \leq \frac{d(d-1)^D - 2}{d-2}$ *(the Moore bound)*

(ii) $D \geq \frac{\log(n-1)}{\log d}$

Proof.

(i) Theorem 10.1 of Bollobás (2001).

(ii) From (i), $n - 1 \leq d^D$, and the result follows. □

Note that the bound in (ii) is tight. For example, the completely connected graph K_n has $d = n-1$ and $D = 1$, while the *Petersen Graph* has $n = 10$, $d = 3$, and $D = 2$ (see Figure 6.14 of Gibbons (1985)). On the other hand, there are examples where no graph exists with diameter $\lceil \log(n-1)/\log d \rceil$, e.g. for $n = 16$ and $d = 4$, $\log(n-1)/\log d = 1.95$, but all graphs of degree $d = 4$ and diameter $D = 2$ have size $n \leq 15$ (see note following Theorem 10.1 of Bollobás (2001)).

For *random* regular graphs, the diameter will be approximately double the bound in (ii): see Theorem 10.15 of Bollobás (2001).

Table 3 shows several families of node-similar graphs which are optimally connected with $\kappa = \lambda = d$, together with their diameters. For the torus/hypertorus, the best case (a 3^q hypertorus) and worst case (a $3 \times m$ torus) are shown explicitly. The prism is a double ring, with links between the rings forming a ring of squares. The antiprism is a double ring with twice as many connections between the rings forming a ring of triangles (facing alternately up and down).

In some ways more important than the diameter is the *average distance* D_{ave} between pairs of nodes. This is bounded by the diameter:

Theorem 5.2 *For any node-similar graph with size n and diameter D:*

$$\frac{nD}{2(n-1)} \leq D_{\text{ave}} \leq D$$

Proof. Consider an arbitrary node a. There is at least one node b at a distance D from a, and by node-similarity the tree of nodes from a "looks like" the tree of nodes from b. The result follows by the triangle inequality. □

The lower bound in this theorem is tight: equality holds for e.g. the octahedron ($n = 6$, $D = 2$), the icosahedron ($n = 12$, $D = 3$), the dodecahedron ($n = 20$, $D = 5$), the hypercube ($n = 2^q$, $D = q$), or the ring of even size ($n = 2m$, $D = m$). For the tightness of the upper bound, consider the following family of graphs:

Definition 5.3 The *Hamming graph* H_m^q is the graph whose nodes are all the vectors $(a_1, \ldots, a_q)$ with $a_i \in \{0, \ldots, m-1\}$, and with links between vectors that differ in exactly one position, i.e. with a Hamming distance of 1 (Lidl & Pilz 1984).

Note that H_m^1 is the completely connected graph K_m; H_2^q is the hypercube; and H_3^q is the 3^q hypertorus. The graph H_m^q is symmetric; and has degree $d = q(m-1)$, diameter $D = q$, and average distance given by the following theorem:

Theorem 5.4 *The Hamming graph H_m^q has average distance:*

$$D_{\text{ave}} = \frac{q(m-1)m^{q-1}}{m^q - 1}$$

Proof. By induction on q. □

For $q = 1$, the Hamming graph H_m^q has $D_{\text{ave}} = D = 1$, and for $q > 1$, $D_{\text{ave}} \to q = D$ as $m \to \infty$, i.e.

Graph	Size n	Number of Links	$\kappa = \lambda = d$	Symmetric
Tetrahedron	4	6	3	yes
Cube	8	12	3	yes
Octahedron	6	12	4	yes
Dodecahedron	20	30	3	yes
Icosahedron	12	30	5	yes
Soccer Ball	60	90	3	—
Truncated Tetrahedron	12	18	3	—
Truncated Cube	24	36	3	—
Truncated Dodecahedron	60	90	3	—
Mutilated Icosahedron	60	150	5	—
Fully Connected	n	$n(n-1)/2$	$n-1$	yes
Hypercube	2^q	$q2^{q-1}$	q	yes
Torus/Hypertorus	$m_1 m_2 \cdots m_q$	$q m_1 m_2 \cdots m_q$	$2q$	if m_i equal
Minimal Cayley Graph	$\|G\|$	$\kappa\|G\|/2$	$\|\{x, \ominus x \mid x \in S\}\|$	sometimes

Table 2: Some Optimally Connected Node-Similar Graphs

Graph	Size n	$\kappa = \lambda = d$	Diameter D	Symmetric
Lower Bound	n	$d \geq 3$	$D \geq \frac{\log(n-1)}{\log d}$	—
Fully Connected	n	$n-1$	1	yes
Hypercube	2^q	q	q	yes
(Hyper)torus	$m_1 \cdots m_q$	$2q$	$\lfloor \frac{m_1}{2} \rfloor + \cdots + \lfloor \frac{m_q}{2} \rfloor$	if m_i equal
3^q Hypertorus	3^q	$2q$	q	yes
$3 \times m$ Torus	$3m$	4	$1 + \lfloor \frac{m}{2} \rfloor = 1 + \lfloor \frac{n}{6} \rfloor$	if $m = 3$
Ring	n	2	$\lfloor \frac{n}{2} \rfloor$	yes
Prism	$2m$	3	$1 + \lfloor \frac{m}{2} \rfloor = 1 + \lfloor \frac{n}{4} \rfloor$	if $m = 4$
Antiprism	$2m$	4	$\lfloor \frac{m+1}{2} \rfloor = \lfloor \frac{n+2}{4} \rfloor$	if $m = 3$
Twisted Prism	$2m$	4	$\max\left(2, \lfloor \frac{n}{4} \rfloor\right)$	yes

Table 3: Diameters of Some Optimally Connected Node-Similar Graphs

the upper bound of Theorem 5.2 is tight for all values of D.

The average distance between nodes is important in determining the amount of communication traffic on links. For the purpose of analysis, we assume that:

(i) A unit amount of traffic is exchanged between every pair of nodes.

(ii) Traffic between two nodes is sent along the shortest path.

(iii) If they are several shortest paths, traffic is divided equally between them.

For symmetric networks, every link "looks the same," and hence every link carries the same traffic load L. For non-symmetric networks, some links can carry more traffic than others, and so we consider the maximum traffic load $L_{\max}$ (i.e. on the most heavily loaded link). The following result places bounds on L and $L_{\max}$:

Corollary 5.5 *For any node-similar graph with size n, degree d and diameter D:*

(i) For symmetric graphs, $L = \frac{(n-1)D_{\mathrm{ave}}}{d} \geq \frac{nD}{2d}$

(ii) In other cases, $L_{\max} \geq \frac{(n-1)D_{\mathrm{ave}}}{d} \geq \frac{nD}{2d}$

Proof. Since a total traffic of $\frac{n(n-1)D_{\mathrm{ave}}}{2}$ is divided among $\frac{nd}{2}$ links. □

Table 4 shows the values of L or $L_{\max}$ for some families of optimally connected node-similar graphs, and Table 5 shows the values for a selection of 12-node optimally connected node-similar graphs. The K_6 prism is a prism with the top and bottom faces completely connected. The six graphs marked with asterisks are the minimal Cayley graphs in Figure 2.

6 Robustness and Link Load

Previously we have considered the impact of node destruction in terms of potentially disconnecting a network. However, node destruction can also have a dramatic impact in terms of link load (L or $L_{\max}$). Several serious failures of the US electrical power distribution network have resulted from this phenomenon.

Optimal connectivity is important in reducing the impact of node destruction on link load. In general, there are at least λ link-disjoint paths between two nodes, and traffic can be distributed over the shortest of these paths, to avoid congestion. If $\kappa < \lambda$, the number of link-disjoint paths, after the loss of a critical node, may drop dramatically from λ to $\kappa - 1$. However, if $\kappa = \lambda$, node destruction can destroy at most one of the λ link-disjoint paths between nodes.

For a ring, the impact of node loss on link load is very serious: for large rings, the load on the most seriously affected link will approximately double. To be precise, the percentage increase in traffic on the most seriously affected link is $100\lfloor \frac{n-3}{2} \rfloor / \lfloor \frac{n+1}{2} \rfloor\%$, which approaches 100% as n becomes large.

For any topology based on a double ring, such as a prism, antiprism, or twisted prism, the increase in link load also approaches 100%. For a topology based on a triple ring, such as the $3 \times m$ torus, the increase

Graph	$\kappa = \lambda = d$	Link Load L or $L_{\max}$	Increase on Node Loss
Lower Bound	d	$L, L_{\max} \geq \frac{nD}{2d} \geq \frac{n\log(n-1)}{2d\log d}$	$0 \ldots 100\%$
Fully Connected	$n-1$	$L = 1$	0%
Hamming Graph H_m^q	$q(m-1)$	$L = m^{q-1}$	$\longrightarrow 0\%$
Ring	2	$L = \frac{1}{2}\lfloor\frac{n}{2}\rfloor\lfloor\frac{n+1}{2}\rfloor$	$\longrightarrow 100\%$
Prism	3	$L_{\max} = \max\left(\frac{n}{2}, \lfloor\frac{n}{4}\rfloor\lfloor\frac{n+2}{4}\rfloor\right)$	$\longrightarrow 100\%$
Antiprism	4	$L_{\max} = \lfloor\frac{n}{4}\rfloor\left(\lfloor\frac{n+2}{4}\rfloor - \frac{1}{2}\right)$	$\longrightarrow 100\%$
Twisted Prism	4	$L = \frac{1}{2} + \frac{n}{4}\lfloor\frac{n}{4}\rfloor - \frac{1}{2}\lfloor\frac{n}{4}\rfloor^2$	$\longrightarrow 100\%$
$3 \times m$ Torus	4	$L_{\max} = \frac{3}{2}\lfloor\frac{n}{6}\rfloor\lfloor\frac{n+3}{6}\rfloor$	$\longrightarrow 50\%$

Table 4: Some Optimally Connected Node-Similar Graphs with n Nodes

in link load approaches 50%. On the other hand, for the completely connected network K_n, there is no increase in link load when a node is destroyed. For the other Hamming graphs H_m^q (e.g. the hypercube or 3^q hypertorus) the increase in link load may be nonzero due to asymmetries created by node loss, but approaches 0% as $m \to \infty$ or $q \to \infty$. Table 4 summarises these results, and Table 5 shows the increase in link load for a selection of 12-node graphs.

Any network design which includes large subrings potentially suffers from the problems of ring networks at a local scale. For example, the "soccer-ball" and the truncated dodecahedron both have 60 nodes and 90 links, but the "soccer-ball" (which is made up of hexagons and pentagons) has $L_{\max} = 96.67$ (increasing 35% on node loss), while the truncated dodecahedron (which is made up of triangles and decagons) has $L_{\max} = 151$ (increasing 39% on node loss).

7 Two-Level Networks

We can combine the advantages of high node connectivity and small diameter in networks by relaxing the condition of node-similarity slightly, and permitting the use of *communications hubs*. A *two-level network* has a *base graph* and a much smaller *hub graph*, both of which are node-similar on their own:

Definition 7.1 A *two-level network* $T(G, H, h)$ consists of an optimally connected node-similar *base graph* G, an optimally connected node-similar *hub graph* H with nodes $H_1, \ldots, H_m$, and a surjective mapping h from base nodes to hub nodes, such that each node x in G is connected to $h(x)$ in H. We write $\widehat{H_i}$ for the set of nodes in G connected to H_i, and require that for each H_i and H_j there is an automorphism of $T(G, H, h)$ that maps $\widehat{H_i}$ to $\widehat{H_j}$.

The following theorem shows how two-level networks combine high node connectivity with small diameter:

Theorem 7.2 *Let $T(G, H, h)$ be a two-level network where the base graph G has node connectivity $\kappa \geq 2$ and diameter D, and the hub graph H has node connectivity $\kappa' \geq \kappa$ and diameter $D' \leq D - 2$. Then $T(G, H, h)$ is optimally connected with node connectivity $\kappa + 1$ and diameter $2 + D'$.*

Proof. The diameter result is trivial. For connectivity, consider the impact of deleting κ nodes. If these are all deleted from G, connectivity is maintained via H, and vice versa. If some nodes are deleted from G and some from H, both G and H remain connected, and the number of nodes remaining in H will be at least one more than the number of nodes deleted from G, so that at least one link will remain to connect G and H. □

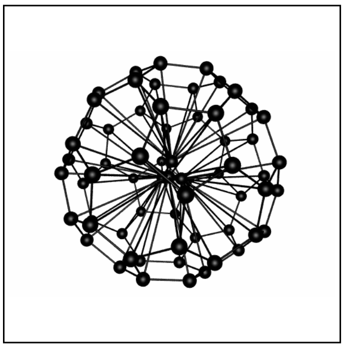

Figure 4: Two-Level "Soccer-Ball" Network

Figure 4 shows an example two-level network where the base graph is a "soccer-ball," and the hub graph is a tetrahedron. The resulting network has node connectivity $\kappa = 4$ and diameter $D = 3$, while the "soccer-ball" alone had $\kappa = 3$ and $D = 9$.

The following theorem places bounds on the link load for an important class of two-level networks:

Theorem 7.3 *Let $T(G, H, h)$ be a two-level network where the base graph G has degree $d = \kappa \geq 2$ and size n, and the hub graph H is the completely connected graph K_m. Then:*

(i) The load on links in the base graph G is at most $3d^2 + \frac{n}{m}$.

(ii) The load on links between G and H is at most $\frac{3}{2}(n - d) + m - 1$.

(iii) The load on links in the hub graph H is at most $(1 + \frac{n}{m})^2$.

Proof. By considering cases. Traffic between nodes up to a distance 1 apart will go via the base graph, traffic between nodes a distance 4 or more apart will go via the hub graph, and traffic between nodes a distance 2 or 3 apart will be divided between G and H. □

For the network in Figure 4, these bounds are 42, 88.5, and 256, while the actual maximum link loads of each type are 19.83, 54.83, and 198.42.

In the event of losing a single hub node, the maximum link loads in the base graph will increase substantially, up to values comparable to those in the base graph alone. For example, for the example in Figure 4, base link loads increase 390% to 96.45 (compared to 96.67 for the "soccer-ball" alone). Loads

Graph	κ=λ=d	Diameter D	Link Load L or $L_{\max}$	Increase on Node Loss
Ring (*symmetric*)	2	6	18	67%
Truncated Tetrahedron*	3	3	10	25%
Ring + 6-chords	3	3	9	22%
Hexagonal Prism*	3	4	9	17%
Hexagonal Antiprism	4	3	7.5	53%
Ring + 4-chords	4	3	6	9%
3×4 Torus*	4	3	6	0%
Hex Prism + Diagonals*	4	2	5	3%
Twisted Prism* (*symmetric*)	4	3	5	23%
Cuboctahedron* (*symmetric*)	4	3	4.75	25%
Ring + 4- & 6-chords	5	2	3.667	14%
Icosahedron (*symmetric*)	5	3	3.6	32%
K_6 Prism	6	2	6	0%
K_{12} (*symmetric*)	11	1	1	0%

Table 5: Some Optimally Connected Node-Similar Graphs with 12 Nodes

on links to hubs can also more than double. However, the most heavily loaded links will usually still be those between hubs, where the increase is up to $50(m-1)/(m-2)\%$. For the example in Figure 4, this bound is 75%, while the actual increase is 50% (from 198.42 to 297.6).

Both the maximum link load and the increase on hub loss are improved if hub regions $\widehat{H_i}$ are interleaved. As an example, consider a base graph which is a ring of 30 nodes, and a hub graph which is a triangle of 3 nodes. If $\widehat{H_i}$ is a group of 10 adjacent nodes, the maximum link load $L_{\max}$ is 108.2, increasing 89% (to 204) on hub loss. On the other hand, if $\widehat{H_i}$ consists of every third node in the ring, the maximum link load is 37.67, increasing only 5% (to 39.5) on hub loss.

Two-level networks can therefore achieve good node connectivity and small diameter, but robustness requires that links in the base graph have sufficient excess capacity to pick up the massive load increase that can occur if hub nodes are lost. This may not always be feasible, and it may be preferable to instead use node-similar graphs of small diameter, where these exist (Dineen 1991, Hafner 1995).

8 Scale-Free Networks

Scale-free networks were introduced by Barabási & Albert (1999) and have attracted a great deal of interest (Barabási 2002). Scale-free networks grow by a process of *preferential attachment.* In particular, an r-linked scale-free graph grows by incrementally adding nodes, and connecting each new node by r links to existing nodes. The nodes that these links go to are chosen randomly with probability proportional to their degree (it is possible for some or all of these links to go to the same node). The properties of scale-free graphs have been derived by Bollobás and others (Bollobás 2001):

Theorem 8.1 *For an r-linked scale-free graph of size n, with $r \geq 2$:*

(i) With probability approaching 1 as $n \to \infty$, the graph is connected.

(ii) As $n \to \infty$, the diameter D approaches $\frac{\log n}{\log \log n}$.

(iii) The number of nodes of degree d is approximately

$$\frac{2r(r+1)n}{(d+r+1)(d+r+2)(d+r+3)}$$

Proof. Theorem 10.29 of Bollobás (2001) and note following. □

Note that the expected diameter does not depend on r, provided that $r \geq 2$. Note also that the average distance D_{ave} will be close to the diameter D, as was the case with other large graphs of small diameter, such as the Hamming graphs H_m^q (Theorem 5.4). As a consequence of (iii), a log-log plot of the number of nodes against degree will be expected to fit a straight line with slope -3.

The World-Wide Web is very close to a scale-free graph, if links are considered to be undirected. Broder *et al.* (2000) studied a large sample of 204 million web pages, and found that:

(i) There was an average of 7 links per page, so that we can estimate $r = 7$, although in fact the out-degree of pages varied from 0 to very large values (following a power-law distribution). One result of this is that the Web is not completely connected.

(ii) 92% of the sample was connected, and 28% was strongly connected, in the sense of pages being mutually reachable by following directed links.

(iii) The average distance D_{ave} between connected pages was 6.83, considering links to be undirected (i.e. links can be followed both forwards and backwards). This value is very similar to the theoretical value for the diameter D ($\log n / \log \log n = 6.48$), so that we can estimate $D \approx 7$ (it is interesting to note that n, D, and D_{ave} for the Web are all very similar to the Hamming graph H_{16}^7, although that graph has 7 times as many links).

(iv) A log-log plot of the number of pages against in-degree has slope -2.1. This differs from the predicted value of -3, either because of non-linearities in the relationship between probability and degree during preferential attachment (Barabási *et al.* 2000), or because of the distribution of out-degrees (Mossa *et al.* 2002).

Scale-free graphs have small diameter, but are not particularly robust against deliberate attack (Albert & Barabási 2002, Bollobás & Riordan 2003), and have low connectivity ($\kappa = 1$). However, we can combine the advantages of node-similar networks (or the Erdős-Rényi random networks described in Section 4) with the low diameter of scale-free networks, in a way which generalises the two-level networks of Section 7:

Corollary 8.2 *Consider an optimally connected node-similar base graph G, with size n and node connectivity κ, and construct a 2-linked scale-free graph on the nodes of G, adding the new links to the existing ones. Then with probability approaching 1 as $n \to \infty$:*

(i) The resulting hybrid graph is optimally connected, with node connectivity $\kappa + 1$.

(ii) The resulting hybrid graph has diameter $D \approx \log n / \log \log n$ (or the diameter of G, if smaller).

Proof. Since the new links add one path between each pair of nodes. □

For example, if we begin with a node-similar optimally connected graph with 1000 nodes and $\kappa = 3$, and add 2 links per node, we will obtain an optimally connected network with node connectivity 4, and diameter $D \approx 4$.

There is some evidence that the World-Wide Web is already a hybrid network of this kind (with an Erdős-Rényi random base graph), since when all pages of in-degree 10 or more are removed, more than 50% of the Web remains connected (Broder *et al.* 2000).

As was the case with two-level networks, most of the traffic will travel via hubs (i.e. nodes of high degree), but in the event of hub loss, links in the base network may experience massive increases in load. Robustness therefore requires that links in the base graph have sufficient excess capacity to pick up this load increase, and, as before, it may be preferable to instead use node-similar graphs of small diameter, where these exist (Dineen 1991, Hafner 1995).

9 Conclusions

We have discussed the graph-theoretic concepts of *node connectivity* and *link connectivity* as measures of network robustness, and argued that node connectivity is most appropriate for modelling the robustness of network topologies in the face of possible node destruction. This is important both for military networks and for civilian networks facing possible terrorist activity.

The most robust networks are *optimally connected*, which means that the node connectivity is as high as possible, given the node degrees.

An important class of networks are the *node-similar* networks, where every node "looks the same." This means that no node is a particularly high-value target, since that the impact of losing a node does not depend on which node is lost. Node-similar networks are particularly appropriate for *decentralised* forms of Network Centric Warfare. We have described five ways of designing node-similar networks which are optimally connected:

(i) We can apply trial and error to network design, calculating the node connectivity for each option. The CAVALIER tool which we have developed has the capacity to do this at the touch of a button.

(ii) All node-similar networks of degree $d \leq 4$ are optimally connected. The "soccer-ball" graph and the torus are two examples.

(iii) Minimal Cayley graphs, derived from the mathematical objects called *groups*, using minimal sets of *generators*, are node-similar and optimally connected. We have provided a number of examples of minimal Cayley graphs, generated using our CAVALIER tool.

(iv) Symmetric graphs, in which all links "look the same," are optimally connected. Symmetric graphs also have advantages in terms of balancing communications load across links. The hypercube is an example of a symmetric graph.

(v) A network derived from some optimally connected network by means of a process which preserves node-distinct paths will also be optimally connected. One such process is truncation/mutilation, which replaces nodes of degree d by completely connected subnetworks of d nodes. The resulting networks are candidates for use particularly in the naval environment.

We have also considered the impact of node destruction on the communications load of links (if node loss causes overload of some links, the result may be almost as serious as if those links were destroyed). This has provided a number of additional network design principles:

(i) Optimal connectivity has the additional advantage of reducing the impact of node loss on link load.

(ii) The *ring* is a very poor network design, since the impact of node loss on traffic load approaches 100%.

(iii) Double rings such as the prism or antiprism are also poor network designs, since the impact of node loss on traffic load also approaches 100%.

(iv) Network designs which include large *subrings* potentially suffer from the same problem at a local scale.

(v) Link load decreases as the degree of a network, and hence the number of links, increases.

(vi) In general, a lower diameter D can also result in a lower link load: the lower bound on link load is proportional to D. However, Table 5 shows that increases in diameter can be associated with decreases in link load.

(vii) Symmetric networks have lower maximum link load, since traffic is equally balanced across links.

(viii) Network diameter can be reduced by adding communications hubs, including the use of *scale-free networks*. In this case, robustness requires a high-capacity base network which can absorb the increase in load resulting from loss of a hub, but this may not always be feasible.

We therefore suggest that military networks, or civilian communications backbones, be node-similar and optimally connected, with degree as high as feasible, diameter as low as feasible, symmetric if possible, and containing no large subrings. Group theory and (minimal) Cayley graphs provide a useful way of generating such networks.

10 Acknowledgements

The authors are indebted to Richard Taylor, John Dekker, and 3 anonymous referees for comments on earlier versions of this paper, to Ed Kazmierczak for discussions on Knuth-Bendix completion, to Reinaldo Giudici for discussions on Cayley graphs, and to John Poirier for discussions on "swarming." Figures 1, 2 and 4 were produced by interfacing CAVALIER to the Persistence of Vision™ ray-tracing package.

References

Albert, R. & Barabási, A.-L. (2002), 'Statistical mechanics of complex networks', *Reviews of Modern Physics* **74**, 47–97. Available electronically at www.nd.edu/~networks/Papers/review.pdf.

Alberts, D., Garstka, J. & Stein, F. (1999), *Network Centric Warfare: Developing and Leveraging Information Superiority*, C4ISR Cooperative Research Program Publications Series, Department of Defense, USA. Available electronically at www.dodccrp.org/Publications/pdf/ncw_2nd.pdf.

Arquilla, J. & Ronfeldt, D. (2000), *Swarming and the Future of Conflict*, RAND. Available electronically at www.rand.org/publications/DB/DB311/DB311.pdf.

Baader, F. & Nipkow, T. (1998), *Term Rewriting and All That*, Cambridge University Press.

Babai, L. (1996), Automorphism groups, isomorphism, reconstruction, *in* R. L. Graham, M. Groetschel & L. Lovasz, eds, 'Handbook of Combinatorics', Elsevier. Also Technical Report TR-94-10, University of Chicago, available electronically at www.cs.uchicago.edu/files/tr_authentic/TR-94-10.ps.

Barabási, A.-L. (2002), *Linked: The New Science of Networks*, Perseus Publishing.

Barabási, A.-L. & Albert, R. (1999), 'Emergence of scaling in random networks', *Science* **286**, 509–512. Available electronically at www.nd.edu/~networks/Papers/science.pdf.

Barabási, A.-L., Albert, R. & Jeong, H. (2000), 'Scale-free characteristics of random networks: the topology of the world wide web', *Physica A* **281**, 69–77. Available electronically at www.nd.edu/~networks/Papers/proceeding.pdf.

Biggs, N. (1993), *Algebraic Graph Theory*, 2nd edn, Cambridge University Press.

Bollobás, B. (2001), *Random Graphs*, 2nd edn, Cambridge University Press.

Bollobás, B. & Riordan, O. (2003), 'Robustness and vulnerability of scale-free random graphs', *Internet Mathematics* **1**(1), 1–35.

Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, A. & Wiener, J. (2000), Graph structure in the web, *in* 'Proceedings of the 9th International World Wide Web Conference', pp. 247–256. Available electronically at www.almaden.ibm.com/cs/k53/www9.final/.

Dekker, A. H. (2001), Visualisation of social networks using CAVALIER, *in* 'Proc. Australian Symposium on Information Visualisation', Vol. 9 of *Conferences in Research and Practice in Information Technology*, Sydney, Australia, pp. 49–55. Available electronically at crpit.com/confpapers/CRPITV9Dekker.pdf.

Dekker, A. H. (2003*a*), 'Centralisation and decentralisation in network centric warfare', *Journal of Battlefield Technology* **6**(2), 23–28.

Dekker, A. H. (2003*b*), Using agent-based modelling to study organisational performance and cultural differences, *in* 'MODSIM 2003 International Congress on Modelling and Simulation', Townsville, Queensland. Available electronically at mssanz.org.au/modsim03/Media/Articles/Vol%204%20Articles/1793-1798.pdf.

Dekker, A. H. (2004), 'Simulating network robustness'. To appear.

Dinneen, M. J. (1991), Algebraic methods for efficient network construction, Master's thesis, Department of Computer Science, University of Victoria, Victoria, British Columbia, Canada. Available electronically at www.cs.auckland.ac.nz/~mjd/networks/mjdMSc.pdf.

Fraleigh, J. B. (1976), *A First Course in Abstract Algebra*, 2nd edn, Addison Wesley.

Gibbons, A. (1985), *Algorithmic Graph Theory*, Cambridge University Press.

Giudici, R. & Abreu, M. (2000), On generating cayley's graphs, *in* 'Thirty-first Southeastern International Conference on Combinatorics, Graph Theory, and Computing', Florida Atlantic University, Boca Raton, Florida, USA. Available electronically at www.ma.usb.ve/~giudici/Marien.ps.

Godsil, C. & Royle, G. (2001), *Algebraic Graph Theory*, Springer Verlag.

Hafner, P. (1995), Large cayley graphs and digraphs with small degree and diameter, Technical Report CDMTCS-005, Department of Mathematics, University of Auckland. Available electronically at http://www.cs.auckland.ac.nz/CDMTCS/researchreports/004cayley.pdf.

Harary, F. (1969), *Graph Theory*, Addison-Wesley.

Hillis, W. D. (1985), *The Connection Machine*, MIT Press.

Humphreys, J. F. (1996), *A Course in Group Theory*, Oxford University Press.

Lidl, R. & Pilz, G. (1984), *Applied Abstract Algebra*, Springer Verlag.

Mossa, S., Barthélémy, M., Stanley, H. E. & Nunes Amaral, L. A. (2002), 'Truncation of power law behavior in "scale-free" network models due to information filtering', *Physical Review Letters* **88**(13). Available electronically at polymer.bu.edu/hes/articles/mbsa02.pdf.

Parunak, H. V. D., Purcell, M. & O'Connell, R. (2002), Digital pheromones for autonomous coordination of swarming UAV's, *in* 'American Institute of Aeronautics and Astronautics (AIAA) First Technical Conference and Workshop on Unmanned Aerospace Vehicles, Systems, and Operations'.

Perry, W., Button, R., Bracken, J., Sullivan, T. & Mitchell, J. (2002), *Measures of Effectiveness for the Information-Age Navy: The Effects of Network-Centric Operations on Combat Outcomes*, RAND. Available electronically at www.rand.org/publications/MR/MR1449/.

van de Goor, A. J. (1989), *Computer Architecture and Design*, Addison-Wesley.

Wilson, R., Walsh, P., Tripp, J., Suleiman, I., Rogers, S., Parker, R., Norton, S., Linton, S. & Bray, J. (2003), 'ATLAS of finite group representations'. Available electronically at www.mat.bham.ac.uk/atlas/v2.0/.