
[Andrews et al. 2008] J. Andrews, S. Shakkottai, R. Heath, N. Jindal, M. Haenggi, R. Berry, D. Guo, M. Neely, S. Weber, S. Jafar, A. Yener. Rethinking information theory for mobile *ad hoc* networks. *IEEE Communications Magazine*. **46** (12), Dec-2008. pp. 94 – 101. Online (for example) at www.eecs.northwestern.edu/~dguo/publications/Andrews+08CM.pdf.

Authors' Abstract. *The subject of this article is the long standing open problem of developing a general capacity theory for wireless networks, particularly a theory capable of describing the fundamental performance limits of mobile* ad hoc *networks. A MANET is a peer-to-peer network with no preexisting infrastructure. MANETs are the most general wireless networks, with single-hop, relay, interference, mesh, and star networks comprising special cases. The lack of a MANET capacity theory has stunted the development and commercialization of many types of wireless networks, including emergency, military, sensor, and community mesh networks. Information theory, which has been vital for links and centralized networks, has not been successfully applied to decentralized wireless networks. Even if this was accomplished, for such a theory to truly characterize the limits of deployed MANETs it must overcome three key roadblocks. First, most current capacity results rely on the allowance of unbounded delay and reliability. Second, spatial and timescale decompositions have not yet been developed for optimally modeling the spatial and temporal dynamics of wireless networks. Third, a useful network capacity theory must integrate rather than ignore the important role of overhead messaging and feedback. This article describes some of the shifts in thinking that may be needed to overcome these roadblocks and develop a more general theory.*

Title: **Fourteen Points About Computational Connectivity**
Subtitle: With Apologies to Woodrow Wilson
Sub-subtitle: Essay Review of Rethinking Information Theory for Mobile *Ad Hoc* Networks

Laurence E. LaForge, PhD, President
The Right Stuff of Tahoe, Incorporated, www.The-Right-Stuff.com

Just over ten months prior to the armistice that ended World War I, Woodrow Wilson unveiled his *Fourteen Points*, a plan for lasting peace. Any resemblance between this essay review and Wilson's plan, including the number of points, is purely coincidental.

Introduction

Noting similarities between the position paper of [Andrews et al. 2008] and my own work on what I call computational connectivity, Scintera Networks CEO Art Reidel asked me to compare and contrast the two. As usual, Art's insights give recursive rise to ever-deepening critical analysis. I found [Andrews et al. 2008] to be immensely worthwhile, and what started out as a routine, annotated bibliographic entry developed into this fourteen-point essay review. The eclectic subject matter herein stems from my counter-cultured role as a Forrest Gump of sorts, privileged witness and occasional contributor at many intersections of science and engineering.

To get the ball rolling, jump to the last paragraph on page 94, which continues through to the top of page 95. Here, the authors take stock of the combinatorially explosive – whence computationally intractable – nature of a single-minded OSI-layer-2 approach:

> *Classical link-based information theory does not appear well suited to the role of describing network performance limits, any more than understanding the functionality of a single neuron gives insight on how the brain at large functions or a single transistor's behavior characterizes the behavior and capabilities of a modern CPU.* (1)

They are right, of course. Indeed, the above is similar in flavor to my 2006 prediction:

> *As concentrated abundances of photolithographically fabricated transistors have enabled microelectronic integration, concentrated abundances of wireless channels enable grid computing.* [LaForge and Turner 2006], p. 15 (2)

A key difference between [Andrews et al. 2008] and my decade or so of work leading up to [LaForge and Turner 2006]: *perspective*. Via Internet, my cursory inspection of the publications of the authors of [Andrews et al. 2008] reveals, as might be expected, that they have focused on *signal processing*. My own research, by contrast, applies *combinatorics* (*e.g.*, the number of eight-neighbor Jordan curves with finite, prescribed perimeter [LaForge, Huang, Agarwal 1994]); *computational geometry* (*e.g.*, delineation of Minkowski quotients of Kharitonov uncertainties [Fadali and LaForge 2001]); and *graph theory* (*e.g.*, solution to one-half of the Problem of Zarankiewicz [LaForge 2000], radix-invariant factorizations induced by Hamming ideals [LaForge 2008]). By and large, my mathematical machinations are by-products of solving problems about computer architecture in general, VLSI in particular (*e.g.*, [LaForge 1999]).

With respect to the analogy posed in passages (1) and (2), the *effective, systematic* use of transistors is largely due to researchers and practitioners who – over a remarkably brief timespan – figured out how to modularly partition the VLSI design space [Katz 1985], all-the-while respecting the underlying physics [Mead and Conway 1980]. In this essay review I highlight how my work extends that of my predecessors and colleagues. At the risk of omitting notable names, I relate their influence to the central theme of [Andrews et al. 2008]: a general capacity theory for wireless networks. I make no apology for fleshing out ideas and assertions with specific examples, historical tidbits, and personal anecdotes. These things really do matter.

1. Computer Architects and Signal Processing Specialists See MANETs Differently

For emphasis, let's contrast approaches to MANETs as epitomized by two camps: *computer architects*, on one hand, and *signal-processing specialists*, on the other. Of course, this distinction is somewhat artificial: no one work, person, or group in academia or industry can be said to purely and completely manifest the caricature evoked by these labels. For example, computer architects are not single-mindedly concerned with wrangling transistors in order to realize an instruction set. Nor, I submit, would many information theorists, such as researchers in error-correcting codes (*e.g.*, [MacWilliams and Sloane 1998]), and whom I categorize here among signal processing specialists, bill themselves as card-carrying members of that camp.

The advent of DSP chips represents perhaps the most striking *overlap* between computer architects and signal processing specialists [LaForge and Lin 1981]. In 1997, for example, Hamid Bolouri and I chaired the program for the IEEE International Conference on Innovative Systems in Silicon (previously, the IEEE Conference on Wafer Scale Integration). As our lead-off speaker for the communications session, Ray Simar, of Texas Instruments, described the then-novel very-long-instruction-word (VLIW) architecture of the TMS320C6201 [LaForge et al. 1997]. Fast forwarding to the present, the power inverter on which The Right Stuff of Tahoe works for Mariah Power is controlled by a Texas Instruments DSP successor to the TI '6201.

Moreover, many researchers and practitioners – from rank-and-file (like me) to famous – have contributed to both computer architecture and signal processing. Consider, for example, Claude Shannon [Lucky 2001 Giants], whose channel capacity formula is cited so abundantly in [Andrews et al. 2008]. Yet, along with John von Neumann, Shannon laid the foundations of computer fault tolerance, an area in which I did my doctoral research. With Edward Moore – after whom is named a classical bound on topological latency ([LaForge 2008], p. 20) – Shannon introduced the notion of a *quorum* [Moore and Shannon 1956], later adapted by Digital Equipment Corporation to its VAXCluster parallel-processing architecture.

With these caveats, it is nonetheless illuminating to distinguish the problem-solving mindsets of computer architects, on one hand, and signal processing specialists, on the other.

2. [Andrews et al. 2008] Properly Recognize the Curse of Dimensionality

If distributed MANET control is to have any hope of approximating optimality then the realtime algorithms that we embed in nodes *must* compute based on just a very few figures of merit. [Andrews et al. 2008] are to be commended for explicitly identifying the need for optimizing MANETs across a space of very few dimensions (p. 99). Among signal processing specialists, alas, this laudable position manifests as the exception rather than the rule.

Consider, for example, the Air Force Research Laboratory's ill-fated Transmission Hypercube Project (TH). As recounted in [LaForge and Turner 2006], signal processing specialists from Engenium Technologies really and truly believed that an ultra-high-fidelity link-level optimization (~300,000 variables on ~36 nodes) could (and should) be linearized (questionable); that a centralized master node could solve the attendant linear program quickly enough; and that commands for channel assignment could be (re)-issued from the master (somehow), in time to keep up with changes in both the mission and the MANET environment.

Not surprisingly, the Engenium approach turned out to be intractable. In the ensuing benchmark competition on 100 nodes, my team delivered throughput that bested that of Engenium, as well as three other companies, by 231% [LaForge and Turner 2006]. Adopting a cross-disciplinary approach that would be considered stock-in-trade for computer architects, our winning solution combined Charles Leiserson's FAT-trees [Leiserson 1985] with parallel schedule decomposition, all the while respecting the physics of directed antennas [LaForge 2004 MANETs].

3. [Andrews et al. 2008] Enumerate a Credible Shortlist of Figures of Merit

As a figure-of-merit shortlist, p. 99 of [Andrews et al. 2008] proffers: 1) throughput; 2) delay; and 3) reliability. These criteria are both credible and consistent with the four-dimensional space that I have espoused on a number of occasions, and in a number of places.

In the context of spaceflight grid computing, for example, James Turner and I cast cross-layer optimization in terms of: 1) tolerance to faults (*e.g.*, nodes disabled by radiation); 2) throughput (as measured by, say, aggregate channel capacity); 3) latency (measured by, say, topology diameter); and 4) cost (*e.g.*, power, mass, or dollars) [LaForge and Turner 2006].

One difference between computer architects and signal processing specialists is revealed by contrasting their approaches to fault tolerance and reliability. The difference is a consequence of the respective lenses through which the two camps view throughput.

As computer architects, we prefer to think of channels as discrete Tinker Toy sticks, with numerical weights that reflect, for example, the (fractional Shannon) capacity, or perhaps the transmit or receive power. Indeed, this graph-theoretic model is exactly that incorporated into The Right Stuff of Tahoe's Connection Foundry software, and used to devise a winning solution to the competition mentioned in Section 2 above. Cognizant that *network* delay is *not* the reciprocal of link bit rate, we optimize objective functions that characterize *how to use* these Tinker Toy sticks, subject to constraints spanning just a very few dimensions. We adopt this approach in part because it is consistent with, and easily admits, the capacitated-flow model of [Ford and Fulkerson 1956]. Bolstering this synergy, the nodal (or link) fault tolerance equals the Menger vertex (or edge) connectivity of the underlying, unweighted graph (for definitions and further explanations of these terms, see [LaForge et al. 2006] or [LaForge 2008]).

Signal processing specialists, by contrast, tend to be concerned with issues for maintaining and maximizing message-bit throughput across individual node-to-node channels: *e.g.*, multi-path reflection, path-loss exponents, bit error rate, or how modulation or error correction must adapt to, say, Rician or Rayleigh fading. To a computer architect, important problems such as these belong in the link-level control subsystem, and should be properly housed there.

It is a *very* formidable problem for OSI Layers 2 (and 1) to sense and actuate channels in a way that delivers within a target envelope of message-bit capacity and cost, or that interrupts (say) an embedding of Connection Foundry algorithms whenever a channel's envelope has been breeched. It is commensurately formidable for embedded Connection Foundry algorithms to translate OSI Layer 5^{+} mission requirements into a multi-hop topology; to prescribe which nodes get connected by channels having a target envelope of properties; to reconfigure the topology in realtime as the mission or MANET environment changes; *and* to do so in a distributed fashion.

Profiting from the lessons of VLSI design, therefore, I suggest that the general capacity theory sought by [Andrews et al. 2008] is best begot, not by squeezing evermore computation into the link (or physical) layers, but rather by modularly partitioning link (and physical) level behavior from the portion of the stack that computes *whether and how to forge*, *close, and use* links. That is, by separating OSI Layers 2 (and 1) from the portion of the stack which *computes connectivity*, in the graph-theoretic sense. Originating at the University of Pisa, the MobileMan framework is consistent with such partitioning – indeed with most of the ideas I espouse in this essay review – albeit with single-hop 802.11 configurations of less than two dozen nodes [Conti et al. 2004].

4. Feasible Regions Are Well-Suited to a Worthy Shortlist of MANET Figures of Merit

P. 99 of [Andrews et al. 2008] mentions – almost in passing – a "capacity region". This potent concept is similar to, and at face value consistent with, feasible regions of design and operation as advocated and surveyed in [LaForge and Turner 2006]. In the domain of VLSI architecture, I credit IEEE Fellow Jim Meindl as the progenitor of such feasible regions (*Ibid.*, Figure 5).

Consistent with the throughput-delay-reliability triplet of [Andrews et al. 2008], my pending patent prescribes Meindl-VLSI-style power × delay as bracketing the link-channel state envelope requested by algorithms ported from computational connectivity, and embedded into nodes:

> *... At any node, the quantified status of links and potential links (on channels), whether active or not, whether measured or estimated from a combination of analytics and indirect measurements, and whether explicitly or implicitly represented, constitutes the nodal link-channel state. The aggregate of nodal*

> *link-channel states is the network-wide link-channel state. My invention realizes a major innovation: synthesis of network-wide optimal topologies, using observations of link-channel states, and actions that control link-channel states, that are as localized as the problem will permit. ... subsystem(s) which actually sample(s) link-channel parameters, or which actually control(s) localized link-channels ... need not be an integral part of my invention.*
>
> *... power × delay is a fundamental optimization criterion for my invention. For given power, end-to-end ... channel capacity ... is the reciprocal of the minimum delay, averaged over all paths between two end nodes. Alternatively, there is a minimum power to sustain a given feasible end-to-end throughput – or, equivalently, to sustain a given feasible delay. ... the per-bit value of the power × delay can be expressed in watt-seconds or, alternatively, in joules. With power × delay as our baseline, my invention can compatibly employ alternative measures of power, such as the number of links ... or other measures of delay, such as the reciprocal of a fraction of the Shannon channel capacity.*
>
> *... power × delay gives rise to envelopes of link optimality, where the tradeoff is allowed to float. ...*
>
> *... Power × delay is fundamental to a flexible link-channel model. It provides the basis for an interoperable interface between my invention and other embedded subsystems, such as those which implement physical, data link, routing, and transport layer functions. ...* [LaForge 2008 Pat Pend] (3)

5. Modularize Routing As Separate from Topology

Alas, signal processing specialists have all-too-often (*e.g.*, [Cruz and Santhanam 2003], [Kawadia and Kumar 2003]) co-mingled packet *routing* (OSI Layer 3) with *topology* (no designated OSI layer), where the latter is taken to mean the aggregate of paths, each spanning Tinker-Toy node-to-node channels, over which packets travel. I submit that this confusion is as much a barrier to *acceptance* of a general theory of network capacity as any of the three roadblocks enumerated by [Andrews et al. 2008]. Or, as I put it in my pending patent [LaForge 2008 Pat Pend]:

> *As a further departure from the prior art, my invention need not rely on instantaneous, local, or OSI-Layer-3 traffic demand to drive topology. Routing is often – and regrettably – co-mingled with topology. Exemplified by Tropos Networks' PWRP, such co-mingling can result in the explosive spawning of links.* (4)

Modularizing the control software, firmware, or hardware that configures MANET topology, along with discrete weighting as described in Section 3 above, readily allows us to *tune* – not only Ford-Fulkerson capacity, Menger connectivity, and fault tolerance – but *latency* as well [LaForge et al. 2006]. MANET network paths tend to be too few and too long, and topology takes top prize as the most under-tapped facet of cross-layer optimization. Hui Zhang, Carnegie Mellon researcher and founder of Conviva (www.Conviva.com, formerly Rinera Networks) understands only too well the importance of this point. Here's how one of Zhang's students summed things up:

> *Get the topology right first, then look at routing.* [Akella 2004] (5)

6. Approximations Can Be Fine, So Go Ahead and Optimize Flow and Capacity Separately

Given that all of the authors of [Andrews et al. 2008] are on faculty at a U. S. university, their common-sense advocacy of "functional capacity" – allowing that it falls short of theoretically best possible (Figure 1; Capacity Approximation Techniques, p. 99) – is refreshing and timely.

By analogy, and hearkening to early champions of VLSI architecture, it is instructive to reflect on the evolution of the grid model for rectangular placement and routing. Consider this vignette of the early 1980's. At that time I worked at Digital Equipment Corporation on CHAS, one of the first industrial CAD suites for collaboration across the multi-layered VLSI design space. On the practical side of things, the simplicity of the grid model made it easier to implement layout compaction; *e.g.*, a recursive algorithm based on the Lipton-Tarjan separator theorem for planar netlists ([Ullman 1984], Sec. 3.5). The grid model also paved the way for powerful theoretical results, such as tight bounds on the area and maximum wirelength of H-trees. That said, VLSI devices and channels for interconnecting them need not, in principle, conform to right-angle geometries. However, and analogous to the functional capacity advanced by [Andrews et al. 2008], the grid model provides a computationally tractable, albeit suboptimal, approximation to packing limits. Similarly, and even though the netlist within a conducting layer is planar, multiple layers admit nonplanar netlists. Even in the 1980's, our routing software (for boards as well as integrated circuits) satisfactorily handled non-zero crossing numbers via quasi-planar heuristics that minimized inter-layer contacts. The grid model, on the other hand, hit a brief speedbump.

CHAS ("CHip ASsembler") was the brainchild of Carol Peters, project leader for whom I worked as technical developer [LaForge 1982]. Leading up to CHAS, the trend at DEC (and at other computer manufacturers) was to assign a large entourage of CAD engineers to service the team responsible for delivery of each CPU. As computer and hardware goals became increasingly ambitious, complexity management, scalability, and multi-disciplinary integration emerged as major hurdles. Emblematic of the magnitude these problems: the hodge-podge of capture and verification tools that CAD engineers cobbled together, time and again, in different combinations, and from one CPU project to another. Understanding that there had to be a better, less expensive way, Engineering Vice President Gordon Bell threw his support behind Carol Peters and CHAS.

One day, as we were frenetically programming the first release of CHAS, Dan Dobberpuhl, perhaps our most capable MOS architect (and later a vice president at Broadcom, www.Broadcom.com), called us in for a conference with Lance Glasser (www.lanceglasser.com), then a professor at MIT. Lance proceeded to show us a novel layout for a NAND gate: researchers at the University of California, Berkeley, had implemented a nonrectangular placement and routing CAD system that generated curvy transistors, curvy interlayer contacts, and curvy wires. The curvy layouts were indeed more compact than those achieved with the right-angle geometries of the venerable grid model, and in many cases nearly achieved theoretical limits on packing density. While an interesting idea, it took us less than a day to conclude that the modest improvement in packing density realized by the Berkeley system would be more than offset by the cost that curvy layout would impose in other areas (*e.g.*, modifications to our design rule checking software, in turn tailored to DEC's integrated circuit fabrication processes).

As a postscript to this footnote in the annals of VLSI, commercial CAD firms, such as Mentor Graphics and Cadence Design Systems, rose to achieve what we (with only marginal success) had tried to do. Since it was for in-house use only, CHAS did not generate direct revenue. CHAS wound up being more difficult to implement than any of us imagined, and its klugey performance engendered resentment on the part of designers who were forced to used it. With picks and shovels, CHAS was gleefully de-commissioned by burying its magtape installation kits in a grave at the DEC semiconductor plant in Hudson, Massachusetts (photo at www.bershad.com/ds/).

Notwithstanding the awkward implementation that we served up with CHAS, I'm fairly certain we made the proper decision by foregoing a penny-wise but dollar-dumb commitment to curvy layouts, and instead sticking with the grid model of placement and routing. Fast forwarding to the present day, I welcome the initiative of [Andrews et al. 2008], and their overtures for preferring holistic network solutions to the diminishing returns scavenged by micro-optimizing node-to-node channels. As epitomized by the Air Force Research Laboratory anecdote thumbnailed in Section 2, signal processing specialists have historically not always been quite so sage.

Toward an *effective* theory, and along lines envisioned by [Andrews et al. 2008], link-centric models tend to neglect a stark reality: application-level throughput is not the same as network capacity. As quasi-planar and grid layout approximations have worked well for VLSI, modularizing routing as separate from topology (*cf.* Section 5) will work well for MANETs. Such modularization provides a computationally tractable, sufficiently accurate approximation to information-theoretic limits. Similarly, it behooves us to optimize flow and capacity separately.

To put a face on things, imagine embedded Connection Foundry algorithms, in conjunction with a fresh generation of link-level subsystems, that *effectively* optimize the channel capacity of an optimal topology, multi-connected in the sense of Menger. This, in turn, allows provisioning (no explicit OSI layer, but present in Internet routers) that optimizes multiple flows against workload demands, consistent with mission-level policies for prioritization, and despite competition for channel capacity. The optimization can be linear or, if fairness is an objective, quadratic.

For illustration, I'll synopsize a 2007 tutorial on video delivery "diamond lanes" that I whipped up for one of our clients, nuMetra (www.nuMetra.com).

Suppose that embedded Connection Foundry algorithms, in conjunction with a compatible link-level subsystem, have configured a *cycle* on a four nodes, such that each pair of adjacent nodes share a bi-directional channel whose message capacity is, say, 1 Mbps; *i.e.*, over all four links, the overall capacity is 4 Mbps. Suppose further that our objective is to maximize aggregate throughput (sum of all pairwise flows), subject to per-session minimum throughput of x Mbps per second between every two nodes. As a generalization of Ford-Fulkerson source-sink max-flow, we can express this problem as a linear program. If the LP is feasible, then its solution comprises six (OSI-Layer 4) sessions, each spanning one or two flows (since each node adjoins two links, we need no more than two flows per session).

It's instructive to write this linear program out in full, and (perhaps with the assistance of optimization software) I invite you to do so. Verify that values of x greater than ½ Mbps are infeasible, and that x = ½ Mbps maximizes the minimum constrained session throughput. At this maxi-min, the maximum aggregate throughput, over all sessions, equals 3 Mbps, *a full 1 Mbps less than the 4 Mbps aggregate channel capacity admitted by the links*. These numbers underscore the "good thing" prescience with which [Andrews et al. 2008] hit the nail on the head:

> *Ultimate Shannon limits on the performance will likely be extremely optimistic in networks, as opposed to links, even with arbitrarily good engineering many years into the future. Therefore, if placing (reasonable) constraints on the Shannon limit makes it easier to compute, this might actually be a good thing.* [Andrews et al. 2008], p. 97 (6)

Continuing our video-on-demand example, at a minimum session throughput of x = ½ Mbps, the solution for maximum aggregate throughput is *unique*: of necessity, and for exactly two of the six sessions, it spans two flows. As might be expected, the maximum aggregate throughput increases as we decrease the minimum session throughput. At x = 1/3 Mbps, for example, the LP admits an aggregate of 10/3 Mbps, or 5/9 Mbps per session. At the same time, the minimum unfairness, as expressed by variance, *increases* from 0 $(\text{Mbps})^2$ / session, when x = ½ Mbps, to 2/81 $(\text{Mbps})^2$ / session, when x = 1/3 Mbps. At x = 0 Mbps we have unconstrained, a.k.a. *best-effort*, traffic: here we can achieve the aggregate channel capacity of 4 Mbps, albeit with session starvation. For $x > 0$ Mbps, however, any optimum solution must include multi-flow sessions.

The preceding reveals how throughput optimization depends crucially on topology. In particular, it takes a multi-path topology to have a shot at a genuine multi-flow session. As we've illustrated, multiple flows are often the best solution, and sometimes the only solution. I should add that embedded algorithms for provisioning will need to approximate optimizations in a distributed fashion whose details depart from the classical, centralized LP by which I have framed things.

If you're keen on optimization, then you probably recognize that the provisioned session throughput minus the minimum session throughput x equals *slack*. Invoking a dual formulation, and where we seek to minimize channel capacity (hence power), subject to the same per-session minimums, our distributed session and flow manager can return slack (*e.g.*, 2/9 Mbps per session at $x = 1/3$ Mbps) to the Connection Foundry topology manager, which in turn instructs the link-level manager to actually reduce channel capacities, whence power. Such interaction exemplifies how the interdisciplinary field of computational connectivity enables *cross-layer optimization*.

Contributors to the fundamentals of session control (my own basic research is more related to topology) have made a credible case for why their algorithms should be folded into in an optimization solution mix that scales to large networks (*e.g.*, position paper: [Briscoe 2008]; in-depth paper: [Wang et al. 2005], co-authored by Steven Low, founder of FastSoft, www.fastsoft.com). Much as configuration of capacitated network topology can (and should) be modularized as separate from routing and link-level concerns (*cf.* Sections 1 and 5 of this essay review), session control and flow provisioning can (and should) be modularized as separate from capacitated topology. The respective domains of analytical discourse are, I submit, both sufficiently different and sufficiently formidable to warrant this partition.

To be sure, there are signal processing specialists (such as those at Engenium, *cf.* Section 2) who have resisted (and probably still resist) the approximations that go hand-in-hand with such modularization. Predictably, these signal processing specialists stridently object that partitioning inhibits us from squeezing every last drop of information-theoretic goodness from our MANET.

To such overspecialized, myopic naysayers, I point to the modeling successes of computer architecture, in general, and to the quasi-planar and grid layout models of VLSI, in particular. Better yet, read [Andrews et al. 2008], and take to heart what your signal processing brethren have written. When properly validated, layered-model approximations can be just fine. So, let's give ourselves permission to go ahead and optimize flow and capacity separately.

7. Metrized Packet Routing Can, and Will, Reduce MANET Overhead

Complementing a transistor-versus-CPU comparison, [Andrews et al. 2008] punctuate the case for a stratifiable theory of wireless networks (1) by noting how merely understanding the behavior of individual neurons does not predict brain function in-the-large [von Neumann 1958]. Along these biologically-inspired lines, consider the interplay between *anatomy* and *physiology* ([LaForge et al. 2006], Sec. 2.4; [LaForge 2008], Sec. 3). Tinker-Toyed together via link- and physical-layer technologies, network topology can be likened unto anatomy. Traffic that traverses our topological anatomy can be viewed as embodying network physiology.

It makes little sense to enter the Kentucky Derby by training the muscles of a horse whose skeleton is flawed. In the case of networks, physiologically superior traffic flow dovetails with anatomically superior topology.

As with most analogies, of course, mine is imperfect. In the case of MANETs, for example, topology morphs with time, sometimes rapidly. [Andrews et al. 2008] properly emphasize the need for reconfiguration under non-equilibrium conditions, such as those animated in our Connection Foundry Voronoi-Delaunay excerpt at www.The-Right-Stuff.com. Further – and as bones, skin, and other tissues are in reality dynamic subsystems of cells – dissection of a MANET reveals dynamic subsystems that continually (re)compute our Tinker Toy skeleton.

With these caveats, and with respect to muscular subsystems that manage traffic across topological MANET skeletons, let's amplify our authors' apt articulation of Roadblock 3:

> *Overhead is a much greater burden in MANETs and must be accounted for …*

> *... This is not simply an academic problem; current military prototype MANETs routinely experience overhead on the order of even 99 percent of the end-to-end packet transmissions ... In a dynamic network the cost of maintaining optimal communication routes may be severe at every layer of the network stack ...*
> [Andrews et al. 2008], p. 96 (7)

Right indeed. For illustration, consider the overhead of *packet routing*. *Given* an optimal (or approximately optimal) topological skeleton, we seek physiologically superior ways for MANET nodes to compute packet egress (5). But hold on. As long as we're revisiting foundational postulates along the lines of Roadblock 1 of [Andrews et al. 2008], might not we do *better* than merely insulating predominantly independent subsystems via well-defined, arm's-length interfaces? After all, and as architects of the VAX / VMS processor / operating system were wont to emulate ([Levy and Eckhouse 1989], [Kenah, Goldenberg, Bate 1988]), biological form and function have co-evolved to common advantage [Thomson 1917, 1992].

Beyond the flow provision / slack capacity adjustment thumbnailed on page 8 of this essay review, that is, can we craft a *synergistic architecture* whose partitioning *bolsters* interactions among MANET subsystems (*cf.* [Andrews et al. 2008] para 2, p. 96)? Reinforcing Roadblock 2's enjoinder for decompositions that "… account for nodal interactions over time and space …", can we plug into ready-made ways for cross-optimizing topology and routing, in a computationally tractable fashion, while at the same time honoring modularity, as advanced in Section 5?

The answer: *Yes*, *we can satisfy all these desiderata*. As a paradigm, I submit *metrized routing* that dovetails with *distributed topology synthesis*.

First consider distributed topology synthesis that computes the links between nodes. In place of the (Euclidean) range, an oversimplification to which [Andrews et al. 2008] properly object (para 2, p. 96), I model MANETs in a terrain or space that has been morphed by power × delay. Here it pays to draw from *computational geometry* [Toussaint 1992], a field which I include as part of computational connectivity. Notwithstanding that the "cell" in "cellphone" can be traced to Voronoi diagrams, computational geometry has been largely neglected by MANET researchers and practitioners (*cf. stochastic geometry*, *Ibid.* p. 99). In addition to my doctoral duties in the Microelectronics Laboratory at McGill, I was part of the Computational Geometry Research Group, and shared offices with three of the authors of what is generally regarded as the best textbook on the subject [de Berg et al. 2000]. Computational geometry opens the door to insights such as this: embedded algorithms for topology synthesis need *not* explicate the morphed power × delay terrain or space. For illustration, browse to www.The-Right-Stuff.com and play the 10-node Connection Foundry Voronoi-Delaunay animation.

The power × delay Delaunay triangulation properly contains a minimum power × delay spanning tree. Packet routing within this tree is preferable when, for example, the OSI Layer 5^{+} mission-level priority is to minimize power, as might well be the case in a MANET of remote sensors. For details, refer to Section 2.1 of my pending patent [LaForge 2008 Pat Pend].

Now suppose that, in the process of synthesizing a topology, embedded Connection Foundry algorithms store a label in each node, such that the *Hamming distance* between the labels on every two nodes p and q equals the minimum number of hops that a packet must traverse in a (bidirectional) path between p and q, said path within a minimum power × delay spanning tree. That is, suppose that, at the same time as they synthesize a topology, embedded Connection Foundry algorithms *Hamming label* it, in a distributed fashion.

A Hamming labeling encodes hop distances, which may be used to advantage in this *geodesic Hamming routing algorithm*: forward a packet to a neighbor whose label minimizes the Hamming distance to the destination label, as encoded in the packet [LaForge 2004].

Setting aside for the moment the issue of re-labeling as a topology morphs, our geodesic (a.k.a. *distance vector*) routing algorithm transports packets with minimum hop latency. The algorithm – *and* the overhead that enables it – are distributed, correct, streamlined, and simple. The animation at www.The-Right-Stuff.com illustrates the action of this algorithm for a packet traversing a 10-node minimum spanning tree contained in a Hamming-labeled topology, in turn induced by Delaunay triangulation of a power × delay terrain. The labels have dimension 9 and radix 2.

We can generalize Hamming labeling in a straightforward way. In the process of synthesizing a topology, embedded Connection Foundry algorithms (say) store a (mixed radix) label in each node, such that the Δ distance $| p, q |_\Delta$ between the label on every two nodes p and q equals the minimum (weighted) number of hops that a packet must traverse in a (bidirectional) path between p and q, and within the topology (which need not be a spanning tree). That is, suppose that, at the same time as they synthesize a topology, embedded Connection Foundry algorithms *Δ-label* it, in a distributed fashion. Thus, Hamming routing is but a special case of a *geodesic metrized algorithm*, in turn a type of distance-vector routing: forward a packet to a neighbor whose label minimizes the Δ distance to the destination label, as encoded in the packet.

In this way, we achieve an architectural form that facilitates function. *I.e.*, distributed synthesis of Δ-labeled topologies (homomorphically metrizable on the Δ distance) yields anatomical information that empowers the modularly separate physiology of metrized packet routing.

In terms of both storage and communication overhead, metrized distance-vector routing costs much less than legacy protocols (*e.g.*, OPSF) of the *path-state* variety. Imposing combinatorially explosive overhead, the latter tabulate the cost of (almost) every path to (almost) every destination, at every routing node (*cf.* matrix *H* in Sec. II.A, [Wang et al. 2005]). Further, and as my 2007 collaboration with Motorola uncovered, MANET-oriented alternatives to path-state routing, such as reactive-proactive hybrid protocols (4), are ineffective at suppressing this combinatorial explosion. With metrized distance-vector routing, on the other hand, the MANET *is* its own routing table, and overhead is scalably parsimonious ([LaForge et al. 2006], Sec. 2.4; [LaForge 2008], Sec. 3). Metrized packet routing is computationally tractable; it is to be preferred over path-state or alternative MANET-oriented protocols; and it serves as a paradigm for fleshing out our authors' reminder of the importance of understanding when …

> *... state information acquisition improves the observed capacity; that is, when is overhead messaging justified by a net capacity increase?* [Andrews et al. 2008], p. 96 (8)

Metrized packet routing on the *Hamming* distance is a natural starting point, for historical as well as technical reasons.

My own discovery of Hamming topologies (or *graphs*, in the proper language of mathematics) began when, as a NASA Fellow, I launched into an investigation of connectivity architectures for high performance, fault tolerant flight computers aboard deep space probes ([LaForge 1999 JPL], [LaForge 2000 NASA]). I proved that that every Hamming graph is induced from a smallest hypermesh, or *ideal*, with the identical labeling radix, non-binary in the general case. The *anatomically attractive* Hamming ideals: 1) maximize throughput, in a Ford-Fulkerson sense; 2) maximize fault tolerance per unit cost, in the sense of Harary and Hayes; 3) minimize latency, in that the ratio of quorum diameter to the best-possible Moore Bound, mentioned on page 3 of this essay review, approaches one with increasing node count (details: [LaForge 2008], Sec. 3).

It wasn't until later on, when sponsored by the Department of Defense, that it occurred to me how Hamming routing [LaForge 2004] could physiologically exploit Hamming labeling [LaForge et al. 2006], as sketched in this section. While I had independently latched onto how Hamming labeling enables distance-vector routing, however, my subsequent research revealed that others had done so, some three decades prior, when working on the so-called *Loop Switching Problem*.

This group included prominent researchers from Bell Laboratories, as well as (unbeknownst to me, until 2006) my own PhD advisor, David Avis (literature survey in [LaForge 2008], Sec. 3.2). Many results about Hamming graphs and labelings have been independently rediscovered, and will quite possibly be rediscovered again, often with new insights. Similarly, many of the insights articulated by [Andrews et al. 2008] have been independently introduced elsewhere, and will quite possibly be re-introduced again, with cumulative benefits. In the meantime, metrized packet routing stands ready and able to reduce MANET overhead.

8. Signal Processing Specialists and Computer Architects Should Get to Know Each Other

Writing, by and large, from the viewpoint of a computer architect, I explicitly acknowledge the influence of the signal processing camp on my research in Hamming graphs (*cf.* tribute to Richard Hamming in [LaForge 2004]). To advance a viable theory of wireless networks (1), moreover, I will need to continually broaden and deepen my understanding of the results and problem-solving mindsets of signal processing specialists. So it is with computer architects everywhere.

Similarly, to further a viable theory of wireless networks (1), signal processing specialists will benefit from broadening and deepening their understanding of the results and problem-solving mindsets of computer architects (*e.g.*, study what's behind the three reasons, synopsized in Section 7, why Hamming ideals are anatomically attractive). [Andrews et al. 2008] are doing just that. So are a handful of signal processing specialists whom I know personally.

As early as 2003, and as a program manager for the Air Force Research Laboratory, Alan Lindsey voiced views bearing uncanny similarity to those expressed by [Andrews et al. 2008]. While Alan's signal processing specialty is Ungerboeck-type, a.k.a. *trellis*, codes ([Alder, Lindsey, Dill 2001]), he has actively reached out for a rigorous cross-layer theory of computational connectivity (*cf.* acknowledgement, p. 28 of [LaForge 2008]; also *Transmission Hypercube Project* remarks in Sections 2 and 13 of this essay review). In August of 2008, Alan rang me up to chat about how he and his DoD contacts had concluded that the time might be ripe for revamping networks as we know them. On hearing this, I referred Alan to the Internet2 and 100 × 100 Clean Slate projects (www.internet2.edu *resp.* www.100x100network.org/institutions.html).

Upon reading my work on laser-tunable topologies [LaForge et al. 2006 VCSELs], Alan introduced me to Stuart Milner, Director of the Center for Networking of Infrastructure Sensors at the University of Maryland. While he specializes in signal processing aspects of optical interconnect, Stu also has an excellent grasp of how topology takes top prize as the most under-tapped facet of cross-layer optimization [LaForge with Milner 2007]. His opinions about the need to overhaul the theory of wireless networks are consistent with those of [Andrews et al. 2008].

Formerly with Sandia Laboratories, and now a professor of electrical engineering at the University of Illinois at Urbana-Champaign, Kent Choquette specializes in fabrication of opto-electrical devices. Immediately following my presentation at the Photonic Technologies session at the 2006 IEEE Aerospace Conference [LaForge et al. 2006 VCSELs], Kent spoke on the steerable VCSEL that he had just invented [Choquette et al. 2006]. Kent immediately realized the potential for cross-layer synergy: his device had found a killer app in architectures of the sort I had proposed; *i.e.*, multi-processors with tunable connectivity [LaForge et al. 2006]. Alas, our proposals for capitalizing on this device-and-architecture synergy have, so far, been rebuffed (see Section 13 of this essay review for why).

In summary, the theory and practice of computational connectivity and cross-layer optimization constitute a *contact sport*. Reaching across the boundaries of our respective communities, computer architects and signal processing specialists would do well to get to know each other's work, *and* each other. *Sportsmanship* will prove pivotal in forging a fresh, comprehensive, practical theory of wireless networks.

9. Random Graphs Will (Almost) Surely Play Key Roles in a Theory of Wireless Networks

Page 99 of [Andrews et al. 2008] recommends earnest application of random graphs to a fresh theory of wireless networks. Although it is difficult to disagree, the authors' illustrations of random modeling are decidedly tilted toward signal processing. For example, contrast models of link behavior among nodes whose spatial distribution is Poisson (or Gaussian, for that matter [Miller 2001]) with graph-theoretic randomness as commonly employed by computer architects:

> *... Traditionally, the latter seeks the fewest IID edges such that, with high probability 1 – o(1), a random graph of order* n *exhibits some particular property; e.g., connectedness, Hamiltonicity, or the emergence of a clique of prescribed order. By contrast, [I] ask for some graph having fewest edges, such that stochastic fault injection – modeled by the deletion of an IID fraction of vertices – induces a subgraph which almost surely possesses beneficial properties; e.g., connectedness or low eccentricity.* [LaForge et al. 2006], p. 19 (9)

Elaborating the passage above [*Ibid.*], I distinguish *random graphs* (edges randomly present on a fixed number of nodes; *i.e.*, the model of Erdös and Rényi [Bollobás 1998] p. 232) from *stochastic graphs* (subgraphs resulting from stochastic deletion of nodes or edges). As a paradigm, I invoke Blough's $\Theta(n \cdot \omega[n])$ counterexample to Scheinerman's claim of an $\Omega(n \log n)$ lower bound on the number of tests for almost sure diagnosis ([Blough 1988], [Scheinerman 1987]). This serves as segue to my construction for a class of stochastically optimal topologies; *i.e.*, ones requiring the fewest number $2 \cdot \log_{1/p} [n/\omega(n)]$ of links, despite debilitating faults distributed with Bernoulli probability p among the n nodes, and such that, with high probability 1 – o(1), all healthy nodes remain Menger connected (*Ibid.* Theorem 3). My construction is based on the VLSI redundancy architecture known as *local sparing* (a.k.a. *cross-strapping* at NASA). For my PhD, I derived analytic formulae for the minimum redundancy and maximum fault tolerance of local sparing in the worst and stochastic cases, and in multivariate detail ([LaForge 1991], [LaForge 1999]).

Reminiscent of the code-theoretic Hamming graphs discussed in Section 7 of this essay review, [LaForge et al. 2006] illustrates a straightforward probabilistic variant that prioritizes latency, rather than fault tolerance and connectivity. Suppose that we seek a binary tree topology for minimizing the expected number of hops from a server to any of n client nodes, and that we know the frequency distribution of messages. Then we can forge links to realize the corresponding *Huffman tree*: the server is the root, and each of n clients is a leaf. This tree effectively minimizes expected latency to and from the server, albeit at the expense of $n – 2$ extra links and $n – 2$ routing nodes, in addition to the nominal n links needed for any tree that spans the server and its n clients.

Besides Blough, Scheinerman, and myself, other researchers have, of course, adapted probabilistic graph theory to wireless networks. Fred Holt, whom I met at Veljko Milutinovic's IPSI Conference in Montenegro in 2004, is one of them. The SWAN computer program that Fred developed with his colleagues makes use of random walks in an existing quorum to determine how to reconfigure links when a new node joins the quorum. Simulations show that this yields attractively tight quorums, with diameters logarithmic in the number n of nodes in the quorum [Bourassa and Holt 2003]. For fixed number of links per node, this is within a constant factor of the best possible Moore Bound on hop latency, first mentioned on page 3 of this review (logarithmic diameters exceed optimal when link degree scales with n). Beyond simulations, Fred and his collaborators have recently published results reflecting a melding of random and stochastic graphs [Greenhill et al. 2008]. Signal processing specialists would benefit from reading the writings of Fred Holt. If you get stuck, I'd suggest ringing him up or emailing him. Fred is very gracious about answering my questions, and I bet he'd be delighted to collaborate with *you*.

[Pinninghoff et al. 2008] probabilistically synthesize networks in a manner quite unlike other approaches mentioned in this section: the randomness manifests as mutations in a *genetic algorithm*. While I cannot see clearly how embedded genetic algorithms could synthesize MANET topologies in realtime, perhaps others can. [Pinninghoff et al. 2008] is notable as well because of its focus on maximizing the throughput of *cars and trucks* traversing networks of *highways*. Although communications scientists and engineers pay lip service to the theory of vehicular traffic flow, its rich body of results (*cf.* www.mtreiber.de/MicroApplet) remains under-tapped by wireless network researchers and practitioners. As with computational geometry, vehicular traffic flow should be welcomed into the tent of computational connectivity.

In summary, while random graphs will (almost) surely play key roles in the development of an interdisciplinary theory of wireless networks, at least some of these roles will probably be quite different from those conceived by signal processing specialists and computer architects.

10. Network Traffic Shaping Could Indeed Benefit from Robust Control

[Andrews et al. 2008] rightly recommend that we invigorate the theory of wireless networks with robust control:

> *... There has been some work in this direction in information theory, perhaps best characterized by considering the capacity of channels when the channel distributions are uncertain .... In general, however, the consideration of robustness to assumptions and modeling is not a prominent aspect of contemporary information theory ...*
> [Andrews et al. 2008], p. 100 (10)

While the authors' thrust is well-taken, their example is, somewhat understandably, tilted toward signal processing. Another application of robust control, certainly under-tapped and perhaps more compelling: *network traffic shaping*. Using feedback provided by observed quantities, this process seeks to proactively throttle the rate at which packets are allowed to enter channels. Related to discussions of capacitated flow and vehicular traffic management in Sections 6 and 9 of this essay review, traffic shaping manifests as stoplights gating the on-ramp entrance rate of freeways in metropolitan cities, and is used mitigate rush-hour congestion. Ski areas employ rudimentary traffic shaping by limiting the number of lift tickets that they sell on any day.

Notwithstanding endpoint transmission control protocols (*e.g.*, TCP, http://en.wikipedia.org/wiki/TCP_congestion_avoidance_algorithms), traffic shaping has not yet had a fair chance at proving its worth, either to the Internet, or to MANETs. The distributed control effected by traffic shapers faces inherent uncertainties. As convolution encoders realize (Z-transform) transfer functions (http://en.wikipedia.org/wiki/Convolutional_codes), control-theoretic traffic shapers can be cast in terms of transfer functions over an appropriate transform domain.

In reality, and as [Andrews et al. 2008] correctly point out, however, the exact values of the coefficients (a.k.a. *plant parameters*) in the transfer functions for these shapers cannot be known. This, in turn, calls for shaping algorithms that are robust; *i.e.*, that manage to satisfactorily throttle traffic within target tolerances, despite such uncertainties.

The advent of Kharitonov quotient algorithms brings an attractively simple means of realizing such robust shapers. I am partial to this particular branch of so-called quantitative feedback theory (QFT) as a result of my research collaborations with Sami Fadali, a professor at the University of Nevada, Reno. In [Fadali and LaForge 2001], for example, Sami – who is an expert in control theory – set up the problem in terms of feasible regions in the compensator plane, in turn determined by Minkowski (set) quotients of Kharitonov regions of uncertainty.

Using ideas from computational geometry, I composed an algorithm for efficiently explicating the boundaries of these quotients. Sami's student, Assem Sonbol, implemented the algorithm as a computer program [Fadali et al. 2001]. It remains to translate our results into a first commercial application.

Whether we wind up using Kharitonov quotient algorithms, QFT, or some other approach, the time is certainly ripe for buttressing the operation of wireless networks with robust control.

11. The Timescale Coherence Problem Underscores the Need for Experimental Testbeds

Table 1 of [Andrews et al. 2008] showcases what I believe to be the most pressing technical challenge to cross-layer optimization: *how much time do subsystems have to perform their respective functions*? Adopting the authors' terminology, I'll call this *timescale coherence*.

For example, consider again the Connection Foundry Voronoi-Delaunay animation at www.The-Right-Stuff.com. How fast can a single node can move through a power × delay terrain, such that the MANET's distributed topological control can keep pace, and reconfigure quickly enough to maintain a Delaunay-triangulated-topology? Is it ten, a hundred, a thousand, or ten thousand kilometers per hour? The answer depends on a number of factors, including, but not limited to, the time it takes for link-level subsystems to sense and respond. This, in turn, depends on the compute power available at each node, as well as the runtime efficiency of link-level algorithms.

The maximum node speed that the MANET can tolerate also depends on our topology of choice. For example, a single reconfiguration event in a Delaunay-triangulated terrain is spatially limited, hence updates to topology (or metrized labelings) are $\Theta(1)$ localized. On the other hand, if we simplify the topology to a minimum power × delay spanning tree (which is properly contained in the Delaunay triangulation), then we get a counterintuitive result: link reconfiguration can propagate across the MANET. Such non-localized reconfiguration may take time $\Theta(n)$ [LaForge 2008 Pat Pend], a prospect that is distinctly unattractive from the standpoint of scalability. Timescale coherence gets even less predictable as more and more nodes move.

More generally, every time I set out to propose or analytically predict the impact of some new subsystem on overall MANET behavior, I am stymied by an inability to confidently estimate (much less pinpoint) timescale coherence. This, in turn, inhibits my ability to forecast conditions under which the MANET will deliver services as advertised or required. My conversations with other scientists and engineers doing MANET research and development confirm that the timescale coherence problem is both broad and deep.

As a result of the timescale coherence problem in particular, and because MANETs in general are dauntingly complex systems of subsystems, I suggest that *experimental testbeds* are our best springboards for advancing a comprehensive, practical theory of wireless networks. Such testbeds would a) provide controlled baselines, benchmarks, and workloads; b) standardize figures of merit and measurement methods; c) validate or dispel assumptions; and d) assuage the provincial practice of researchers and practitioners to discard objectivity in favor of their own solutions.

One of my former digital engineers, Kirk Korver now works for a "smart home" company, Control4 (www.control4.com). Over lunch a couple of months back, he shared with me that Control4 is the world's most voracious consumer of chips for ZigBee wireless technology. Kirk's firm also makes extensive use of its own ZigBee testbed, one that spans some *five hundred* nodes.

Under many conditions, a multi-hop path makes more sense than a single node-to-node link. As there are VLSI sweet spots for trading signal propagation, integrity, and the spacing and size of buffers on-chip [Lin and Mead 1984], there are communication sweet spots for throughput, message rate, distance, embedded computational power, and store-and-forward buffer size.

Holding all other factors fixed, suppose, for illustration, that the link message rate decreases inversely as distance r raised to some constant $c > 1$; at unit distance the rate is k/m, and at distance r it is $(k/m)r^{-c}$. Then a two-hop path can deliver a better rate than a single link. To see this, hold the endpoint nodes apart at fixed distance $2R$, so that they communicate at rate $(k/m)(2R)^{-c}$. Along the line segment between the endpoints place a third node at distance $x \geq R$ from one of the endpoints. As a "restoring logic" store-and-forward repeater, the intervening node enables a rate that is half the minimum rate to either endpoint node. The factor of one-half accounts conservatively for each link in the two-hop solution holding silent while the other is busy. Then, using the intervening node, the end-to-end rate $\frac{1}{2}(k/m)x^{-c}$ is greater than the single link rate, as long as $2^{1 - 1/c} \geq x/R$. In particular, when $x = R$ the two-hop rate achieves its maximum $\frac{1}{2}(k/m)R^{-c}$, which is always greater than the single-hop rate. More generally, multi-hop paths make more sense than single-hop paths under many conditions. Experimental MANET testbeds are essential to validating and refining our predictive understanding of these conditions.

People walk around to get better cellphone reception, so why shouldn't MANET nodes do the same sort of thing? Cross-layer aspects of computational connectivity lend themselves to experimental testbeds that fold *motion planning* and *robotics* into the mix. Taking the exercise of the preceding paragraph one step further, imagine the intervening node optimizing message rate by commanding its locomotive subsystem to move to distance $x = R$ from one endpoint node.

In summary, our building and measuring activities lag too far behind our analytics and simulations. Timescale coherence, motion planning, and robotics give a glimpse of but a few of the reasons why both academia and industry should make it their priority to emplace quality, large-scale testbeds for experimenting on MANETs.

12. Ensure That Network Theory Puts a Premium on Practical News We Can Use

A specialist in signal processing, Alan Lindsey confesses that information theorists do it. Sami Fadali, an expert in QFT, freely admits that control theorists do it. Based on my experience doing research in computer fault tolerance, I'd have to concede that professionals who publish in the reliability literature do it with reliable regularity. I have been guilty of it, and not just once.

So just what is *it? It* is the tendency of researchers to become so engrossed with theory (*e.g.*, [Kantabutra 1991]) that they neglect to make themselves useful to the rest of the world.

Stalwart editor of the *IEEE Transactions on Reliability* from 1969 until 2004, Ralph Evans captured *it* with his own pet acronym: "CBIM", short for "Correct But Irrelevant Mathematics" (I was relieved when Ralph told me that CBIM did not apply to my own paper [LaForge 2000]).

According to Pete Stoddard, a first-rate engineer if there ever was one, *it* stems from "LOLO".

In 1985, and riding the tide of Moore's Law [G. E. Moore 1970], large-screen workstations featuring individual processors were beginning to supplant dumb terminals attached to CPUs. In an early application of windowing software, Pete and I were lead developers for METALS, virtual emulators of the functionally rich (but dumb) CAD terminals leased by McDonnell Douglas (since acquired by Boeing), and key to that company's licensing for their Unigraphics mechanical engineering package (now known as NX, http://en.wikipedia.org/wiki/NX_5). Working in Dick Anderson's mechanical CAD group at Digital Equipment Corporation's vintage facility (a.k.a. *The Mill*) in Maynard, Massachusetts, our job was to port the look, feel, and functional interface of these Unigraphics terminals onto the then-fledgling VAXstations [LaForge 1985], thus saving DEC, a paying user of Unigraphics, about ten thousand dollars per seat.

When burdened with an especially silly decision that would come down from upper management, the no-nonsense Pete would turn to me and growl, "LOLO". Translation: "Lock On, Lock Out".

LOLO is exactly what we MANET researchers and practitioners need in order to do our jobs well. LOLO also inhibits us from stretching beyond our areas of concentration. One example, threaded throughout this essay review: the tendency of signal processing specialists and computer architects to lock out most everything (and everyone) else, in order to lock onto whatever here-and-now problem they deem to be important, and so that they can muster the focus to solve it.

The antithesis of LOLO: "Mile Wide, Inch Deep", "MWID" for short. To forge a functional theory of wireless networks, and picking up where Section 8 of this essay review left off, it pays to balance LOLO against MWID. To get there demands acknowledging our human frailty: we are apprehensive that we'll never know all of what we're supposed to know – in depth or in breadth. This leads to fear of plunging into yet another domain of concentration (after all, doing so might cause our brains to explode, like the *Enterprise* computer did in *Star Trek*). The upshot: we are reluctant to explore beyond our comfort zone, in anything but a perfunctory way.

In light of our finite energies and lifespans, it *does* make sense to exercise caution against overreaching. Often as not, however, we err on the side of too much caution. Forging a viable theory of computational connectivity requires extraordinary willpower and effort, even for the authors of [Andrews et al. 2008], who have already proved they can do extraordinary things. For guidance about how to handle challenges such as this, I turn to two world-class researchers whose inspirational writings live on: Richard Hamming [Hamming 1986], and a former professor of mine at MIT, Gian Carlo-Rota [Brehm 1998]. In particular, this reminder from Rota is spot-on relevant: "You are never going to catch up, and neither is anyone else" ([Rota 1997], Lesson 8).

I choked down a heaping helping of Rota Lesson 8 when, in 1983, I signed on as the youngest member of a team of otherwise-seasoned computer architects in Digital Equipment Corporation's Large VAX Engineering group (LVE). DEC had invested in, and partnered with, Trilogy Corporation. Founded by computer pioneer Gene Amdahl, Trilogy was freshly infused with $180M in capital. At that time, IBM mainframes executed the lion's share of application software. Trilogy planned to emulate the instruction set(s) of IBM's System/370 family of CPUs, and at superior price-to-performance. Amdahl's new company would implement the micro-machine in a 2 micron bipolar process, and with integration that was wafer-scale for the era.

As part of an overall DEC effort known as Trident, LVE aimed to apply Trilogy's technologies in order to realize a high-performance VAX: *Antares*. My boss for *Antares*, Fernando Colon-Osorio, tasked me with learning Trilogy's technologies, from soup to nuts, so that I could serve as a living knowledge resource that he and the rest of the architects on his team could tap. Trilogy required that proprietary details about its technologies be kept under lock and key. It's interesting to re-read notes that I jotted down when studying Trilogy documents and video tapes, in a windowless room in the rear of the DEC library in Marlboro, Massachusetts [LaForge 1983 Trilogy].

Glancing through those crinkled, yellowed pages, I see, for example, that each 28 cm^2 Trilogy wafer was 62 mm on an edge, ran hot at 900 watts (50 watts / cm^2), made use of voting logic with triple-module-redundancy (TMR, or *tri-logic*, hence *Tri*logy), and featured 2000 vertical input-output pins realizing – via triple and double redundancy – 780 logical inputs and outputs. The hermetically-sealed wafer was backed by molybdenum (coefficient of thermal expansion about that of silicon) for extracting heat to water at 22° C. Trilogy was to package its 30-MIPS machine, sporting 1 GB of memory, as a 70 *kilo*watt dual-processor. Retailing for about $5M, and at a profit margin of about 25%, it would be cooled to an on-wafer junction temperature of 55° C by way of water flowing at 50 gallons minute, and pressurized to 50 pounds per square inch (yes, Trilogy engineers mixed metric and *avoirdupois* units). Within a wafer, latched combinational logic was partitioned to fit 250×500 micron^2 reticles, and synchronized via hierarchical clocks at 5, 10, and 20 nanosecond ticks. Trilogy's instruction-level architecture was ambitious: two concurrent streams, each pipelined over 18 stages.

At first inspection, the Trilogy documents and video tapes were overwhelming. I felt completely out of my league. I had worked with MOSFETs, but was much less familiar with bipolar circuits of the sort that Trilogy was using. I was not fluent in IBM instruction sets, and had never even heard of triple-module-redundancy. (At that time, TMR enjoyed staunch proponents in the literature of fault tolerance [Siewiorek and Swarz 1982]; later, as a doctoral student, I established conditions under which other redundancy techniques are much more meritorious than TMR.)

Figuring that I'd better fess up to my boss that I might not be up to the job, I slinked over to the Fernando's office, where he happened to be chatting with Yale Patt. Now at the University of Texas (http://users.ece.utexas.edu/~patt/), Yale was at that time a consultant to LVE. Sympathetic to my trepidations, Yale and Fernando counseled me with enduring advice: as computer architects, we must learn, analyze, and build like architects, in depth and in breadth. We must *try* to learn as much we need (or can or want) about any subject. We cannot be hindered by doubts (internal or external) that our efforts will not be good enough. This shot-in-the-arm of encouragement spurred me to reinvigorate my study of the Trilogy technologies.

At things turned out, it was all for naught. Trilogy lost its bet that TMR would overcome the obstacle of wafer yield [Sunday Camera 1984]. Along with its partnership with DEC, Trilogy faded into the sunset of high-performance computing efforts that never bore fruit (also see survey in [LaForge et al. 2006 VCSELs], Sec 1.2). Even before it became clear that Trilogy's technology would not work, however, *Antares* was scuttled. In an urgent call for all hands on deck, LVE vice-president Bob Glorioso impressed Fernando's team into service to help rescue a long-overdue ECL-based 32-bit computer, internally designated as *Venus*. I still sip coffee from a blue mug (my favorite) whose *Venus, The VAX Mainframe* gold foil is nearly worn off. I wound up as a utility player writing addressing mode diagnostics, improvising the microcode release system, and pitching in with sundry other chores [LaForge 1983]. Fernando, on the other hand, stepped into a star leadership role for *Venus*. With amazing perseverance (or would that be LOLO?), he turned things around, and spearheaded delivery of what DEC eventually sold as the VAX 8600.

The wise words of Fernando Colon-Osorio and Yale Patt shine light on ingredients – unspoken by [Andrews et al. 2008] – that are *essential* to computational connectivity. For our theory to be comprehensive, rigorous, and practical, we must enact a willing suspension of disbelief that we shall make it so, tempered in turn by realistic, intellectual honesty and – not least – hard work.

To make computational connectivity *practical*, we theoreticians need to frequently (not occasionally) come down off our high horses, reach out to practicing engineers charged with building wireless networks that actually work, and *actively* make our research *accessible*.

By analogy, consider the results derived in [Leighton and Leiserson 1985] for discretionary wiring of one- and two-dimensional arrays. As a doctoral student at McGill, I spent several weeks digesting that one paper, arguably the single most potent contribution to the theory of VLSI fault tolerance. Under the stochastic fault model described on page 12 of this essay review, discretionary wiring delivers $\Theta(\log \log n)^2$ redundancy, ratioed as the minimum area of the fault tolerant layout divided by the nominal area of a two-dimensional array. This is somewhat better than the $\Theta(\log n)$ redundancy of local sparing, also described on page 12 of this essay review. (Both discretionary wiring and local sparing yield $\Theta(\log n)^{1/2}$ maximum wirelength.) However, the reconfiguration algorithm for discretionary wiring is complicated, while that for local sparing is straightforward. In addition, the number $\Theta(1)$ of switches per array element in the case of local sparing is much less than that $\Theta(\log \log n)^2$ for discretionary wiring, a scaling penalty that casts doubt on assumptions of fault-free switching. These theoretical results underscore the need for experimental wafers, complete with the ability to inject faults into array elements and switches, and for the purpose of objectively assessing the predictive accuracy of the analyses of [Leighton and Leiserson 1985] and myself. That is, which architecture is better in practice?

In 1991, I put this very proposition to Vinod Agarwal, Northern Telecom Professor of Electrical Engineering at McGill, and my supervisor in the Microelectronics Laboratory. (Vinod went on to found LogicVision, a Silicon Valley company which, late in 2008, rejected a buyout from Virage Logic www.viragelogic.com). I mentioned to Vinod that Jack Raffel's group at MIT's Lincoln Laboratory had built proof-of-concept arrays, using discretionary wiring and fuse-type switches [Raffel 1988]. I further suggested that Northern Telecom might fund a McGill team to fabricate dynamically-switched array architectures, from which we could glean credible, comparative data.

Vinod replied that a study group commissioned by him had already rejected making such wafers. The reason: the group could not figure out what [Leighton and Leiserson 1985] were saying ought to be built. To electrical engineering circuit specialists unschooled in *computational* VLSI [Ullman 1984], the exposition of [Leighton and Leiserson 1985] was incomprehensible. I was about to graduate, and was unable to persuade Vinod to undertake the project without having me around to interpret the paper. To my knowledge, controlled experiments on array architectures, like those I proposed to Vinod, still have not been carried out. Even today, such experiments would constitute a valuable contribution to the body of VLSI research.

The preceding vignette from my graduate school days reinforces Section 8 of this essay review: theorists and practitioners ought to make more of an effort to relate to each other. Better yet, why not get our hands dirty, and implement theories ourselves? Are we reluctant to try this because it requires so much effort? Maybe, but isn't that why they call it "work"? For a telling example of the persistence so often required to turn theory into reality, please take the time to read R. Stanley Williams' refreshingly candid account of his group's discovery of the memristor, at Hewlett Packard. A few excerpts:

> *... To this day, I have no idea how Greg [Snider] happened to come across that [brilliant but underappreciated 1971 paper by Leon Chua]. Few people had read it, fewer had understood it, and fewer still had cited it. At that point, the paper was 31 years old and apparently headed for the proverbial dustbin of history ...*
>
> *I wish I could say I took one look and yelled, "Eureka!" But in fact, the paper sat on my desk for months before I even tried to read it. When I did study it, I found the concepts and the equations unfamiliar and hard to follow. But I kept at it because something had caught my eye ...* [Williams 2008], p. 33 (11)

To forge a comprehensive, rigorous, practical theory of computational connectivity, we must muster the work ethic of architects, and pursue whatever subjects we require, in depth and in breadth. We also need to adopt a holistic perspective that tempers our tendency to lock onto the *problème du jour*, and lock out everything else. On this account I recommend this aphorism from David Avis, my doctoral thesis advisor at McGill: *Good theory isn't necessarily hard theory*.

Or, in the reputed words of Albert Einstein, we should keep things "as simple as possible, but no simpler" [Brainy Quote]. Where the simplest possible theory is nevertheless complicated or difficult – and it frequently is – we should bend over backwards to convey it clearly.

Not least, the theory that we serve up should be of practical benefit to the rest of the people in the world. After all, they pay our wages. And this fact serves as fitting segue to Section 13.

13. The Agenda of [Andrews et al. 2008] Requires Substantial Money. Let's Raise It

The biographical thumbnails at the end of [Andrews et al. 2008] speak to the ability of the authors to attract research funding; for example, eight out of the eleven are listed as NSF CAREER recipients. I heartily congratulate them, and wish them continued success in this important facet of research (college administrators nowadays say that it is the only facet that really matters).

That said, I wonder what *percentage* of U.S.-based professionals, accomplished in one or more of the disciplines mentioned in this essay review, would say that their MANET-related R&D has been unfairly starved for funds? I bet that it's most of them. Stu Milner, introduced in Section 8 of this essay review, is among them. I know this because he told me so [LaForge with Milner 2007].

Money for MANETs is a topic whose scope and importance warrants thorough exposition under separate cover. Nevertheless, funding – or lack thereof – for computational connectivity definitely merits inclusion among the fourteen points headlined in this essay review. Let's hit the highlights.

I'll reserve exhausting my laundry list of transgressors and transgressions until some other time. For our purposes here, funding problems with MANETs unfurl as a tragedy in three interrelated acts: I) incompetent, irresponsible, myopic, lazy, and – yes – corrupt managers of purse strings at U.S. Government agencies, or at companies, such as Boeing or DRA, who are in bed with the government; II) misappropriation of funds to spurious busywork, to the politically correct or programmatically anointed, or to projects whose technical execution or managerial framework doesn't pass the smell test; III) failure of signal processing specialists and computer architects to achieve solidarity, and to stem the debilitating pestilence unleashed in Acts I and II.

In but one scene within Act I, enter the Small Business Innovation Research program, SBIR for short. Even though my shop has on several occasions won SBIRs or SBIR-type awards, the program is, in reality, a waste of taxpayer money. It is so antithetical to integrity that we have given up on it. For one thing, it is clear that evaluators are not seriously interested in innovation. In violation of procurement rules, they frequently make up and substitute their own criteria *ex post facto*. An examination of debriefings from our own in-house sample of submitted proposals, dating back ten years, reveals that only about forty percent of the evaluations came within spitting distance of applying the criteria set forth in the respective solicitation. In many cases (including two where we won) it is painfully clear that the evaluators did not read even most of the proposal.

When my shop proposes a project that emphasizes computational connectivity (or a functional theory of wireless networks, in the lexicon of [Andrews et al. 2008]), evaluators typically fail to grasp even the bare rudiments of technical merit (yes, we *have* wordsmithed our submissions [Lucky 2005]). Evaluators routinely fail at this crucial responsibility, despite technical merit being stated as the primary factor. Here's what Kent Choquette, introduced in Section 8 of this essay review, says about the grants-and-contracts hypocrisy: "My experience with proposals these days is that technical issues are not what matters to get funded" [Choquette to LaForge 2006].

So what *does* matter to get funded? My contacts at SAIC say that it is *relationships* and *image*.

Fair enough. After all, and in line with the topics of Section 12, it *is* up to us to *reach out* and sell our wares (and ourselves), whether we are signal processing specialists or computer architects.

But hold on: a U.S. Government whose grants and contracts *systematically prioritize* relationships and image at the expense of technical merit is prone to fraud, waste, and abuse. Which is exactly what we have in Act I of our tragedy. In 2006, for example, the Air Force Research Laboratory Sensors Directorate at Wright Patterson Air Force Base (WPAFB) awarded an SBIR contract, under solicitation AF06-198, to Defense Research Analysts (DRA). However, DRA was already an embedded contractor at WPAFB. Ronald Clericus, the principal investigator for this particular award, was a DRA employee who worked in the Integrated Demonstrations and Applications Laboratory (IDAL). Mr. Clericus was even so brazen as to list a WPAFB telephone number and email handle in his submission. The Air Force itself provided this information openly on the Internet, along with the abstract of his proposal. As a poster child for corruption hidden in plain sight, the latter states, "… DRA will use the Sensors Directorate Virtual Combat Laboratory (VCL) and Integrated Demonstrations and Applications Laboratory (IDAL) testbeds to demonstrate the [proposed] architecture's feasibility …". Beyond the obvious impropriety of

DRA, an in-house contractor, benefiting from an erstwhile level-play-field solicitation, this is a direct violation of rules that prohibit use of a U.S. Government facility for a project funded under the SBIR program, without prior approval from the Small Business Administration (*cf.* Section 9(f)(3) of the SBA's SBIR Policy Directive, www.sba.gov/SBIR/SBIR-PolicyDirective.pdf).

I know this for a fact from the Freedom of Information Act response I received from Mr. Edsel Brown, Jr., Assistant Administrator of the SBA's Office of Technology. In 2007, Mr. Brown furnished me with a list of the 68 waivers granted since 2003 (Section 9(f)(3) took effect near the end of 2002). Not a single one of these waivers makes mention of DRA. Oops, busted.

What's worse, from 1995 to 2006, the Air Force expended at least $13,309,829 on at least 19 SBIRs for DRA to work on or in VCL or IDAL. Eight of these commenced work after Section 9(f)(3) took effect in late 2002; all eight involved DRA working on or in VCL or IDAL. In light of these egregious violations, and as detailed in my firm's agency protest, did anyone in the U.S. Government admit to malfeasance in enriching their DRA bedfellows, and then set about enacting a remedy? Of course not. In an arrogant wink-and-nod variation on a theme of corruption, Contracting Officer David Shellabarger ignored the rules and clear evidence of wrongdoing, pronounced no foul, and rewarded the Air Force and DRA for their *relationship*.

Escalating the matter to the DoD Inspector General, to then-Secretary of the Air Force Michael Wynne (later fired by Secretary of Defense Robert Gates), to Senator John Ensign, to Congressman Dean Heller, and to Phil Neel of the SBA IG, my letters have all been met with a "Thank you for your correspondence, but we don't really want to fix this" rebuff. These Pontius Pilate impersonators are what we get when we let *relationships* and *image* trump technical merit.

Complementing the preceding Cliffs Notes of a small-business scene from Act I of our MANET funding tragedy, consider this *big-business* variant. From 1993 to 2002, Darlene Druyun was Deputy Assistant Secretary of the Air Force for Acquisition and Management. Caught trading sweetheart treatment of Boeing's airborne tanker program in exchange for employment favors benefiting her daughter, son-in-law, and herself, she pled guilty to fraud in 2004. Released after nine months in jail, Druyun nonetheless collects her U.S. Government pension [Colarusso 2006].

The decade-long collusion between DRA and the Air Force demonstrates that the jailing of Darlene Druyun was the exception, not the rule. With pension intact, she is laughing all the way to the bank. Having *effectively* beat the rap, perhaps she is toasting Ronald Clericus, DRA, and Druyun-wannabe David Shellabarger. With the Air Force, grants-and-contracts crime *does* pay.

Mind you, the Navy is neck-and-neck with the Air Force for incompetent, irresponsible, corrupt mismanagement of MANET research and development. In 2008, I discovered, quite incidentally, that Izhak Rubin, head of IRI Computer Communications Corporation, had violated SBIR Rule 1.3. The latter prohibits anyone with fulltime employment with any firm, other than the small business selected for award, from serving as principle investigator. Since Rubin is – and for some years has been – a fulltime professor at UCLA, he is – and has been – ineligible to serve as PI. However, an email I received from SPAWAR technical program manager Doug Hulbert plainly revealed that Rubin was the PI for an about-to-be-awarded contract under solicitation N07-193. When I pointed out the pending violation, Hulbert and cohorts responded by allowing substitution of a puppet PI; figuring they had closed the loophole, SPAWAR issued the contract! This despite Rubin having perpetrated the identical fraud in at least ten previous SBIRs or STTRs, enriching himself to the tune of $3,969,583. Only when I filed a third-party complaint with the SPAWAR IG (my shop did not bid on any of these proposals) did the Navy terminate the contract with IRI. Even then, the IG's response was toothlessly lame. Where's the punishment of Hulbert and other SPAWAR officials? Why not put Rubin in jail, or at least make him repay the fraudulently-obtained $3,969,583? While the final line in the script for this scene has yet to be written, one thing is clear: with the Navy, as with the Air Force, procurement crime *does* pay.

Moving on to Cliffs Notes from Act II, it's instructive to storyboard the Transmission Hypercube Project. Launched in 2003 with a laudable agenda quite similar to that of [Andrews et al. 2008], TH fizzled fast. A major reason: higher-ups at the Air Force Research Laboratory substituted a myopic signal processing specialist in place of the project founder, Alan Lindsey. Introduced in Section 8 of this essay review, Alan is himself a signal processing specialist. However, he kicked off TH with a novel, *holistic* approach. Alan strove to integrate expert contributions from five small R&D firms into an embedded, cross-disciplinary solution for using the electromagnetic spectrum much more efficiently than ever before. Alan is now with Austral Engineering and Software (www.australengineering.com). His replacement, Mike Gans, gutted TH by relegating all five companies to the busywork of competing against one another (*cf.* Section 2 of this essay review). The benefit to the taxpayer? A $500K Quixotic quest in pursuit of a narrow, contrived benchmark, borrowed and rewarmed from a program long-beleaguered: JTRS.

DoD's Joint Tactical Radio System illustrates only all-too-well a big-business scene from Act II of our MANET funding tragedy. Jim Van Buren, a colleague of mine at Draper Laboratory, has firsthand experience with JTRS. At the outset, JTRS Cluster 5 was an $8.7B program led and administered by the Army. With $341M for research and development alone, General Dynamics squandered opportunity, treasure, and schedule from the get-go [Aero Daily 2006].

While the preceding snippets from Acts I and II of our MANET tragedy punctuate funding transgressions of the Air Force, Navy, and Army, culpability is by no means confined to these branches of our military. Instead of David Shellabarger, Doug Hulbert, or Mike Gans, I could have introduced you to their miscreant counterparts at the Missile Defense Agency, at the Office of the Secretary of Defense, or at DARPA. Nor has DoD cornered the market on grants-and-contracts screw-ups. For academic as well as business-oriented solicitations, my shop has been subjected to the fraud, waste, and abuse that permeates Homeland Security, the National Science Foundation, the Department of Energy, the Small Business Administration, and even the United States Agency for International Development. This basic storyline is echoed by many of my colleagues in industry and academia. If you've bid on a MANET grant or contract then, likely as not, you could rewrite Acts I and II, substituting your own encounters of fraud, waste, and abuse.

Which bring us to Act III, and the fast-fading possibility of transforming our MANET funding tragedy into one that ends on an upbeat note. Here I'd suggest tapping the wisdom of Robert Lucky, IEEE Fellow and former Communications Sciences Research Director at AT&T. Lucky's treasure trove of pieces appeared regularly in *IEEE Spectrum*, through 2008, and under his "Reflections" column. Writing in 2001, for example, Lucky effectively reinforces differences between signal processing specialists and computer architects – a theme kicked off in Section 1, and woven throughout this essay review. Feast on this morsel from [Lucky 2001 Giants]:

> *Where are the Claude Shannon's of today, I wonder?… Where once there seemed to be only a handful of colorful personalities creating pioneering technology, now there is a cast of millions, doing such important, but forgettable, things as creating small variations on the details of protocols. … Individualism and risky technological endeavor have been largely replaced by adherence to the business plan.* (12)

By extension, if signal processing specialists and computer architects are to execute the agenda of [Andrews et al. 2008], then perhaps a *subset* of Lucky's cast of millions might self-organize into a *virtual* Shannon (colorful personalities go with the territory), and set about pioneering computational connectivity. The investment powers behind Lucky's anonymous, generic *business plan* ([Buckman 2009]) have neglected to reseed the bounty of innovation inherited from communications Giants Who Walked the Earth. Pursuant to the demise of Bell Labs and other U.S. industrial research laboratories [Lucky 2006], investment powers – and we, by our inaction – are starving the geese that lay the golden eggs of wireless networking. Or, as Lucky writes:

> *The theory is that a commodity business doesn't support research, because there is no potential for added value. Can this be so in telecom? This question puzzles me, and as an engineer I worry about our own culpability. Where is our vision for what we could accomplish with technology? Do we have a dream, or are we at the mercy of cost-conscious business managers? ... We need a vision to take to the business leaders and to the governments – a brief elevator speech:* give us your blessing, and this is what we can do for you. [Lucky 2004] (13)

Acts I and II of our MANET tragedy underscore how U.S. Government funding is a lottery fraught with unwritten, situational ethics (hmm … make that *lack* of ethics). To boost our individual chances, we can buy more lottery tickets. We can rub elbows with program managers at conferences. We can bootstrap our image by hitching onto projects headed by notables whose reputations are entrenched. We can pad our publication count. (Come on, now, how many of your papers say something *really* profound? How many are just drivel? [Lucky 2001 Giants]). In the immortal words of Dr. Phil, "How's that working for ya?" Most of us would admit, "Not so well".

Allowing that *some* relationship- and image-building is necessary, just where on the slippery slope of self-promotion do we cross the Nixonian River Styx into the Clintonian Hades of delusional rationalization? (*It's a must for tenure and promotion. It's to pay my employees. Everyone does it.*) While the founders of DRA or IRI almost certainly did not start out with the intention of exploiting across-the-board shabby adherence to grants-and-contracts rules, their companies wound up doing it anyway. Furthermore, at what point do such rules, and our efforts to enforce them, materially impede getting the right things done? Or, as Robert Lucky recounts:

> *Before long every employee of the company [AT&T] would need to be occupied with the task of writing and distributing rules, and no one would be left to do the actual work ...* [Lucky 2005 Rules] (14)

No matter what we signal processing specialists and computer architects do, we will never completely exterminate the incompetent, irresponsible, myopic, lazy, corrupt malefactors at U.S. Government agencies. Nor will we exorcise the demons of malfeasance at companies who are in bed with the government. Nor can we count on Inspectors General, whose erstwhile charter is to fight fraud, waste and abuse, to do their job. It is *un*realistic to fantasize that we will ever be liberated from the diversion of funds to spurious busywork, to the politically correct, or to the programmatically anointed. Such is the perennial muck through which fate has condemned us to wallow, scraping as much as possible from our shoes prior to engaging polite company.

Still, we can and, and should, make it our business to keep the muck at neck level or below. Through our own example, of course, but also by giving the boot to grants-and-contracts polluters who – like Ronald Clericus, David Shellabarger, Darlene Druyun, Izhak Rubin, and Doug Hulbert – claim as inalienable their birthright to rip other people off. For this reason, and to paraphrase the reputed words of Thomas Jefferson, the price of *good theory* is eternal vigilance (http://en.wikiquote.org/wiki/Thomas_Jefferson).

On that note, we must also make it our vigilant business to restore sensible balance, and reverse the systematic prioritization of relationships and image at the expense of technical merit. Alas, this is by itself a huge undertaking, one that goes hand-in-hand with re-stoking the pipeline of quality, technically educated professionals in the United States:

> *So we engineers are largely responsible for creating the flat earth, and we're seen as the key to mitigating or capitalizing on its effects. But as Norman Augustine says* [Augustine 2007]*, we in the U.S. are in danger of falling off this flat earth. The trends are bad, and look to be almost irreversible, as the pool of engineers shrinks and the innovation ecology stagnates.* [Lucky 2008] (15)

One reason that submissions from my shop to U.S. Government agencies get such short shrift is that proposal evaluators go to work dumbed down, a predicament that gives rise to many of the same symptoms as street drugs. In rare instances where they *are* competent, evaluators tend to be nearing retirement, overloaded with work, or both. Be careful about gratifying the urge to dumb things down. Exacerbated by an economic downturn, we are now suffering the consequences of that wish come true. Aptly forecasted by [Augustine et al. 2007], a storm of technological deficiency has indeed descended upon us. In a furious crescendo, it has washed out to sea our *Sputnik*-era pride in scientific and engineering prowess. What remains? Lazy, smug ineptitude.

While advocating a plethora of spending for education and young researchers, [Augustine et al. 2007] ignore the obvious: we have let too many links in the chain of scientific and engineering knowledge stay broken for too long, and do not have enough competent evaluators to properly award monies based on technical merit. Until we remedy this problem, in large part at least, agency-sponsored R&D will continue to be squandered. Is it any wonder that taxpayers mistrust the U.S. Government's fiscal prudence? In disagreement with Norman Augustine, I instead say, "Uncle Sam, we will *reduce* your allowance each time you spend unwisely. Instead of sitting on your duff, get to work. Instead of dumbing down, smarten up. *Or else*." The *or else* matters.

Alas, and in a panicked herd response to the economic downturn of late, the U.S. Congress is *stampeding* in the direction of dumbing down: too much spending too soon just goes to waste (*but at least we're doing something!*) In desperation, we seem to crave unabashed socialism. That being our poison of choice, money to bail out banks would be better spent on startups [Markman 2009], such as those that have sprouted to fill the post-Lucent telecom void [Lowenstein 2005]. But wait … isn't this is a Catch-22 that brings us full circle back to Acts I and II?

Maybe, maybe not. Perhaps we U.S.-based signal processing specialists and computer architects should abandon any nationalistic notions of science and technology, let the globalized chips fall where they may, and pedal our MANET wares to the highest bidders. The relatively robust economies of countries such as China and South Korea might be more friendly to wireless network research and entrepreneurship. Alternatively, we can hunker down like medieval monks, and try to implement the agenda of [Andrews et al. 2008] by the flickering light of tallow candles.

14. Computational Connectivity Is Poised to Step Up With a Functional Network Theory

In this essay review I've shared cross-sectional views of a fresh theory of wireless networks, views which are complementary – and usually complimentary – to those of [Andrews et al. 2008]. From modular architectures, to collegial collaboration, to practical utility, to funding, I cordially welcome the new theory into the big tent of computational connectivity.

Computational connectivity seeks solutions to variations on the anatomical problem of joining endpoints and waypoints as a least-cost network of node-to-node links, the aggregate of which minimizes latency, and maximizes capacity and fault tolerance. Computational connectivity also seeks to optimize, approximately if not exactly, the physiological flow of traffic through the network. Such traffic – including but not limited to information packets and automobiles – may be treated either as discrete units or as continuous fluids.

Spanning multiple levels of abstraction, computational connectivity is akin to computer architecture in general, computational VLSI in particular [Ullman 1984]. While the objectives and constraints of computational connectivity are multivariate, we can stave off the Curse of Dimensionality by applying just a very few judiciously-selected figures of merit. To be sure, the subject domains within the tent of computational connectivity are replete with open challenges. I list two of these in Section 11 of this essay review, nine in [LaForge et al. 2006]. That said, and bolstered by recent advances in research, the state-of-the-art is easily good enough to begin building cross-layer optimization into serious testbeds of the sort I advocate in Section 11.

While mobile *ad hoc* networks stand out as a killer application for computational connectivity, there are others. In the case of *tethered* networks, the endpoints and waypoints are nodes at rest, and the links are static or quasi-static (realizations of) Shannon channels. If traffic consists of automobiles and trucks then links correspond to highways, and the waypoints to interchanges. With MANETs, there can (or should) be little or no distinction between endpoints and waypoints; *i.e.*, we have *edgeless* topologies. Especially where nodes are in motion, the links prescribing the instantaneous topology can – and generally will – be opened, closed, and adjusted much more rapidly, at much less cost, than for wired networks. This adds more degrees of freedom to the control of links. It also poses greater challenges to the subsystem that manages the links.

Does computational connectivity include information theory? Yes, especially *computational* information theory, *e.g.*, modulation algorithms and their implementations on signal processors. Does it include other subject domains? Yes, especially computational graph theory and computational geometry [Toussaint 1992]. Computational connectivity also ties together queuing theory, the theory of capacitated flows, the theory of optimization, and computational electromagnetics. This list is not exhaustive, and will undoubtedly be augmented as time goes on.

Computation is key to computational connectivity. Thus, for example, we are interested in offline algorithms and software – such as Connection Foundry, OPNET, and ns-2 – for automating network design and flow optimization. Computational connectivity is also about online (perhaps realtime) algorithms, be they distributed or centralized, and embedded into nodes as software, firmware, or hardware. Yet, how many signal processing specialists write about, know about, or care about computational models, or computational efficiency? Meet your new partners, folks: computer scientists have a rightful place, next to you in the tent of computational connectivity.

Not long ago Robert Lucky lamented, "If systems engineering is so valuable, why is it so seldom practiced?" [Lucky 2006] With MANETs, this trend must be – and will be – reversed. For reasons set forth in Section 12 of this essay review, signal processing specialists and computer architects will be reborn as systems engineers. Emerging together from the tent of computational connectivity, we will forge a fresh, comprehensive, *practical* theory of wireless networks. Give us your blessing – and a few breadcrumbs – and we will do good things for you.

Index of Names: People, Places, Projects, Products, Organizations, Acronyms, Allusions
This index covers text through Section 14. It does not cover References beginning on page 29.

References

[Aero Daily 2006] GAO assessment of Joint Tactical Radio System Cluster 5. *Aerospace and Defense Daily Report.* 1-Jun-2006. p. 5. Online at www.aviationweek.com. Cited p. 21.

[Akella 2004] A. Akella. *On the Scaling of Congestion in the Internet Graph*. PowerPoint presentation. Dec-2004. Online at www.100x100network.org. Cited p. 5.

[Alder, Lindsey, Dill 2001] F. A. Alder, A. R. Lindsey, and J. C. Dill. Performance of simplex signaling for a trellis-based block code in high dimensions. *Digital Signal Processing*. **11** (2), Apr-2001. pp. 159 – 167. Cited p. 11.

[Augustine 2007] N. R. Augustine. *Is America Falling Off the Flat Earth?* National Academies: 2007. Summary essay of [Augustine et al. 2007]. Online at www.nap.edu/catalog/12021.html. Cited p. 22.

[Augustine et al. 2007] N. R. Augustine, C. R. Barrett, G. Cassell, S. Chu, R. M. Gates, N. S. Grasmick, C. O. Holliday, S. A. Jackson, A. K. Jones, J. Lederberg, R. Levin, C. D. More, Jr., C. Murray, P. O'Donnaell, Jr., L. R. Rayond, R. C. Richardson, P. R. Vagelos, C. M. Vest, G. M. Whitesides, R. N. Zare. *Rising Above the Gathering Storm: Energizing and Employing America for a Brighter Economic Future*. National Academies: 2007. Online at www.nap.edu/catalog/11463.html. Cited p. 23.

[Blough 1988] D. M. Blough. *Fault Detection and Diagnosis in Multiprocessor Systems*. Ph.D. dissertation, Baltimore: Johns Hopkins University, 1988. Cited p. 12.

[Bollobás 1998] B. Bollobás. *Modern Graph Theory*. New York: Springer-Verlag, 1998. Cited p. 12.

[Brainy Quote] Brainy Quote. *Make Everything as Simple as Possible, But Not Simpler*. Albert Einstein Quotes. Online at www.brainyquote.com. Cited p. 18.

[Brehm 1998] D. Brehm. *Rota reflects on math and mathematicians*. Cambridge, MA: MIT News Office, 28-Oct-1998. Online at http://web.mit.edu/newsoffice/1998/rota-1028.html. Cited p. 16.

[Briscoe 2008] B. Briscoe. A fairer, faster Internet: TCP – the way we share bandwidth – needs a makeover. *IEEE Spectrum*. **45** (12). Dec-2008. pp. 43 – 47. Online at www.spectrum.ieee.org/dec08/7027 and via ieeexplore.ieee.org. Cited p. 8.

[Bourassa and Holt 2003] V. Bourassa and F. B. Holt. Load-balancing in a Swan network of true peers. *Proceedings, 2003 IPSI: Internet, Processing, Systems for E-education/E-business, and Interdisciplinaries*. 4-Oct-2003. Cited p. 12.

[Buckman 2009] R. Buckman. Venture capital's coming collapse. *Forbes*. 12-Jan-2009. Online at www.forbes.com/forbes/2009/0112/066.html?partner=daily_newsletter. *Lottery slogans with an Ivy League veneer.* Cited p. 21.

[Choquette et al. 2006] K. D. Choquette, J. J. Raftery, Jr., and A. C. Lehman. Beam steering in photonic crystal vertical cavity. *Proceedings, 2006 IEEE Aerospace Conference*. 4-Mar-2006. Cited p. 11.

[Choquette to LaForge 2006] K. D. Choquette to L. E. LaForge. Re: STTR Phase I Debrief – Tunable Multi-Processor for Realtime Simulation of Tactical Mobile Ad Hoc Communication Networks. Email correspondence. The University of Illinois at Urbana-Champaign *resp.* The Right Stuff of Tahoe. 31-Jul-2006. Cited p. 19.

[Colarusso 2006] L. M. Colarusso. Druyun still gets pension. *FederalTimes.com*. 24-May-2006. Online at http://www.federaltimes.com/index.php?S=1813458. Cited p. 20.

[Conti et al. 2004] M. Conti, G. Maselli, G. Turi, and S. Giordan. Cross-layering in mobile *ad hoc* network design. *Computer*. Feb-2004. pp. 48 – 51. Cited p. 4.

[Cruz and Santhanam 2003] R. L. Cruz, A. V. Santhanam. Optimal routing, link scheduling and power control in MANETS. *Proceedings of the INFOCOM '03*. 22nd Annual Joint Conference of the IEEE Computer and Communications Societies. 2003. pp. 701 – 711. Cited p. 5.

[de Berg et al. 2000] M. de Berg, M. van Krevald, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. 2nd ed. Berlin: Springer-Verlag. 2000. Cited p. 9.

[Endlich 2004] L. Endlich. *Optical Illusions: Lucent and the Crash of Telecom*. New York: Simon and Shuster. 2004. Cited p. 32.

[Fadali and LaForge 2001] M. S. Fadali, L. E. LaForge. Linear time computation of feasible regions for robust compensators. *Int'l Jnl Robust Nonlinear Control*. **11**, Jun-2001. pp. 295 – 305. Cited pp. 2, 13.

[Fadali et al. 2001] M. S. Fadali, L. E. LaForge, and A. Sonbol. Linear time computation of robust stability margins. *Proc. 2001 American Control Conference*. pp. 3469 – 3474. Cited p. 14.

[Ford and Fulkerson 1956] L. R. Ford, Jr. and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*. **8**, 1956. pp. 399 – 404. Cited p. 4.

[Greenhill et al. 2008] C. Greenhill, F. B. Holt, N. Wormald. Expansion properties of a random regular graph after random vertex deletions. *European Journal of Combinatorics*. **29** (5), Jul-2008. pp. 1139 – 1150. Cited p. 12.

[Hamming 1986] R. W. Hamming. *You and Your Research*. Transcription of a talk delivered at a General Research Colloquium. J. F. Kaiser, ed. Morristown, NJ: Bell Communications, 7-Mar-1986. Cited p. 16.

[Katz 1985] R. H. Katz. *Information Management for Engineering Design*. New York: Springer-Verlag. 1985. Cited p. 2.

[Kantabutra 1991] V. Kantabutra. A geometric proof of Shannon's result using continuous mathematics. *Proc. 25th Annual Conf on Information Sciences and Systems*. The Johns Hopkins University: Baltimore. 1991. p. 186. Cited p. 15.

[Kawadia and Kumar 2003] V. Kawadia and P. R. Kumar. Power control and clustering in *ad hoc* networks. *Proceedings of the INFOCOM '03*. 22nd Annual Joint Conference of the IEEE Computer and Communications Societies. 2003. Online via http://black.csl.uiuc.edu/%7ekawadia/papers/clustering.pdf. Cited p. 5.

[Kenah, Goldenberg, Bate 1988] L. J. Kenah, R. E. Goldenberg, and S. F. Bate. *VAX/VMS Internals and Data Structures, Version 4.4*. Newton, MA: Digital Press. 1988. Cited p. 9.

[LaForge 2008 Pat Pend] L. E. LaForge. *Embedded System for Distributed Wireless Network Optimization and Topology Synthesis*. U.S. Patent Pending, Apr-2008. Cited pp. 5, 14.

[LaForge 2008] L. E. LaForge. A Cayley-Prüfer type formula for the number of Hamming-labeled trees. *Discrete Applied Mathematics*. Submitted Mar-2008. Online at www.the-right-stuff.com/A%20Cayley-Prufer%20Type%20Formula%20for%20the%20Number%20of%20Hamming%20Labeled%20Trees%20-%202008-03-17.pdf. Cited pp. 2, 3, 4, 8, 10, 11.

[LaForge 2004] L. E. LaForge. Hamming graphs. *Proc. 2004 IPSI: Internet, Processing, Systems for E-education/E-business, and Interdisciplinaries*. 9-Oct-2004. Cited pp. 9, 11.

[LaForge 2004 MANETs] L. E. LaForge. *Clique-Factorized Optimum MANET Throughput Enabled by Directed, Power-Controlled Antennas*. Technical report, U.S. DoD contract FA8750-04-C- 0020. Reno, NV: The Right Stuff of Tahoe. 14- Apr-2004 revised 7-Aug-2004. Cited p. 3.

[LaForge 2000] L. E. LaForge. What designers of microelectronic systems should know about arrays spared by rows and columns. *IEEE Transactions on Reliability*. Special Issue on Very Large Scale Integrated Systems. **39** (3), Sep-2000. pp. 251 – 272. Cited pp. 2, 15.

[LaForge 2000 NASA] L. E. LaForge. *Architectures and Algorithms for Self-Healing Autonomous Spacecraft*. Grant report, NASA Institute for Advanced Concepts, 9-Jan-2000, revised 28-Feb-2000. Online at www.niac.usra.edu/studies/347LaForge.html. Cited p. 10.

[LaForge 1999 JPL] L. E. LaForge. *Fault Tolerant Physical Interconnection of X2000 Computational Avionics*. Pasadena, CA: Jet Propulsion Laboratory, document JPL D-16485. Revised 18-Oct-1999. Online at http://faculty.erau.edu/laforgel. Cited p. 10.

[LaForge 1999] L. E. LaForge. Configuration of locally spared arrays in the presence of multiple fault types. *IEEE Trans. on Computers*. **48** (4), Apr-1999. pp. 398 – 416. Cited pp. 2, 12.

[LaForge 1991] L. E. LaForge. *Fault Tolerant Arrays*. PhD dissertation. Montreal: McGill University, 1991. Cited p. 12.

[LaForge 1985] L. E. LaForge. *Minimization of Serial Data Overrun*. Digital Equipment Corporation internal report. Maynard, MA: 13-Dec-1985, revised 4-Apr-1986. Counterintuitive conclusions about communication parameters that made the METALS/Unigraphics terminal emulator really work. Cited p. 15.

[LaForge 1983 Trilogy] L. E. LaForge. Notes about Trilogy technology. Marlboro, MA: Digital Equipment Corporation. 11-Nov-1983 through 27-Dec-1983. Under non-disclosure agreement that has long since expired. Cited p. 16.

[LaForge 1983] L. E. LaForge. *Mal-adjustment of Working Sets*. Digital Equipment Corporation internal report. Marlboro, MA: 23-Jan-1983. How to accelerate the productivity of VAX 8600 designers by properly balancing resources consumed by logic simulators. Cited p. 17.

[LaForge 1982] L. E. LaForge. *CAD at DEC and at MIT*. Digital Equipment Corporation internal report. Hudson, MA: 13-Jan-1982. Comparison of DEC's CHAS (CHip ASsembler) with VLSI design tools used in academia. Cited p. 6.

[LaForge et al. 2006] L. E. LaForge, J. R. Moreland, and M. S. Fadali. Spaceflight multi-processors with fault tolerance and connectivity tuned from sparse to dense. *Proceedings, 2006 IEEE Aerospace Conference*. 4-Mar-2006. Online at http://faculty.erau.edu/laforgel/Publications/Refereed/Tunable-Spaceflight-Multi-Processors-Small-Filesize.pdf. Cited pp. 4, 8, 10, 11, 12, 23.

[LaForge et al. 2006 VCSELs] L. E. LaForge, J. R. Moreland, R. G. Bryan, and M. S. Fadali. Vertical cavity surface emitting lasers for spaceflight multi-processors. *Proceedings, 2006 IEEE Aerospace Conference*. 4-Mar-2006. Online at http://faculty.erau.edu/laforgel. Cited pp. 11, 17.

[LaForge et al. 1997] L. E. LaForge, H. Bolouri, D Sciuto, and S. K. Tewksbury, eds. *Proceedings, 1997 IEEE International Conference on Innovative Systems in Silicon*. Los Alamitos, CA: IEEE Computer Society Press, 1997. Cited p. 3.

[LaForge, Huang, Agarwal 1994] L. E. LaForge, K. Huang, and V. K. Agarwal. Almost sure diagnosis of almost every good element. *IEEE Transactions on Computers*. **43** (3), Mar-1994. pp. 295 – 305. Cited p. 2.

[LaForge and Turner 2006] L. E. LaForge and J. W. G. Turner. Multi-processors by the numbers: mathematical foundations of spaceflight grid computing. *Proceedings, 2006 IEEE Aerospace Conference*. 4-Mar-2006. Online at www.the-right-stuff.com/Multi-Processors%20by%20the%20Numbers%20-%202006-03-04.pdf. Cited pp. 2, 3, 4.

[LaForge and Lin 1981] L. E. LaForge, S. L. Lin. *Tau-2: a VLSI Array Processor for Generate and Test*. Digital Equipment Corporation internal report. Hudson, MA: 14-Dec-1981. Signal processing chip that I co-designed: a pipelined, two-dimensional block parity architecture implemented in 3 micron NMOS. Depicted at www.The-Right-Stuff.com. Cited p. 3.

[LaForge with Milner 2007] L. E. LaForge with S. D. Milner. Telephone conversation. The Right Stuff of Tahoe *resp.* Center for Networking of Infrastructure Sensors, University of Maryland (www.cnis.umd.edu/muri/Dr_milner.htm). 14-Feb-2007. Cited pp. 11, 19.

[Leighton and Leiserson 1985] T. Leighton and C. E. Leiserson. Wafer scale integration of systolic arrays. *IEEE Trans. on Computers*. **C-34** (5), May-1985. pp. 448 – 461. Cited pp. 17, 18.

[Leiserson 1985] C. E. Leiserson. FAT-trees: universal networks for hardware-efficient supercomputing. *IEEE Trans. on Computers*. **C-34** (10), Oct-1985. pp. 892 – 991. Cited p. 3.

[Levy and Eckhouse 1989] H. M. Levy and R. H. Eckhouse. *Computer Programming and Architecture: The VAX*. Newton, MA: Digital Press. 1989. Cited p. 9.

[Lin and Mead 1984] T-m Lin and C. A. Mead. Signal delay in general RC networks. *IEEE Transactions on Computer-Aided Design*. **3** (4), Oct-1984. pp. 331 – 349. Cited p. 14.

[Lowenstein 2005] R. Lowenstein. How Lucent lost it. *Technology Review*. Feb-2005. pp. 78 - 80. Review of [Endlich 2004]. Cited p. 23.

[Lucky 2008] R. W. Lucky. Engineers and the flat earth. *IEEE Spectrum*. **45** (3). Mar-2008. p. 19 Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. Cited p. 22.

[Lucky 2006] R. W. Lucky. Unsystematic engineering. *IEEE Spectrum*. **43** (9). Sep-2006. p. 84. Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. Cited p. 24.

[Lucky 2006] R. W. Lucky. Lab for sale. *IEEE Spectrum*. **42** (9). Sep-2005. p. 92. Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. Lamenting Lucent's sale of its Bell Labs facility at Holmdel, N.J., the author recounts some of the technological history of the grounds and nearby establishments. Cited p. 24.

[Lucky 2005 Rules] R. W. Lucky. Robert's Rules. *IEEE Spectrum*. **42** (3). Mar-2005. p. 80. Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. Cited p. 22.

[Lucky 2005] R. W. Lucky. Wordsmithing. *IEEE Spectrum*. **42** (1). Jan-2005. p. 96. Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. *So many words. So little meaning*. Cited p. 19.

[Lucky 2004] R. W. Lucky. Where is the vision for telecom? *IEEE Spectrum*. **41** (5). May-2004. p. 72. Online at www.boblucky.com/spectrum.htm and via ieeexplore.ieee.org. Cited p. 22.

[Lucky 2001 Giants] R. W. Lucky. When giants walked the earth. *IEEE Spectrum*. **38** (4). Apr-2001. p. 104. Not at ieeexplore.ieee.org. Online at www.boblucky.com/spectrum.htm and at http://psyche.ee.auth.gr/adds-odds/texts_/whengiantswalkedtheearth.pdf. Cited pp. 3, 21, 22.

[MacWilliams and Sloane 1998] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error Correcting Codes*. Amsterdam: Elsevier Science. Tenth printing, 1998. Cited p. 2.

[Markman 2009] J. Markman. Why the bank bailouts are doomed. *MSN Money*. 29-Jan-2009. Online at http://articles.moneycentral.msn.com/Investing/SuperModels/why-the-bank-bailouts-are-doomed.aspx.Cited p. 23.

[Mead and Conway 1980] C. A. Mead and L. A. Conway. *Introduction to VLSI Systems*. Menlo Park, CA: Addison-Wesley. 1980. Cited p. 2.

[Miller 2001] L. E. Miller. Probability of a two-hop connection in a random mobile network. *Proceedings, 35th Conf. on Information Sciences and Systems (CISS)*. Baltimore, MD. Mar-2001. Online at http://w3.antd.nist.gov/pubs/2hpaper2.pdf presentation at http://w3.antd.nist.gov/pubs/ciss01.pdf. Cited p. 12.

[G. E. Moore 1970] G. E. Moore. What level of LSI is best for you? *Electronics*. **43** (6), Feb-1970. pp. 126 – 130. Often attributed as the original, published articulation of Moore's Law. Cited p. 15.

[Moore and Shannon 1956] E. F. Moore and C. E. Shannon. Reliable circuits using less reliable relays, part I. *Journal of the Franklin Institute*. **262**, Sep-1956. pp. 191 – 208. Early, perhaps first, use of *quorum* on p. 202. Cited p. 3.

[Pinninghoff et al. 2008] M. Pinninghoff, R. Contreras, and J. Atkinson. Using genetic algorithms to model road networks. *Computer*. **41** (12), Dec-2008. pp. 60 – 67. Cited p. 13.

[Raffel 1988] J. I. Raffel. The applications of wafer-scale technology to neuromorphic systems. In *Advanced Research in VLSI: Proceedings of the Fifth MIT Conference*. J. Allen and F. Thompson, eds. Cambridge, MA: MIT Press. Mar-1988. Cited p. 18.

[Rota 1997] G.-C. Rota. *Ten Lessons of an MIT Education*. MIT Alumni Association Family Weekend. Fall, 1997. Online at http://www.math.tamu.edu/~cyan/Rota/mitless.html. Cited p. 16.

[Scheinerman 1987] E. R. Scheinerman. Almost sure fault tolerance in random graphs. *SIAM Journal of Computing*. **16** (6), Dec-1987. pp. 1124 – 1134. Cited p. 12.

[Siewiorek and Swarz 1982] D. P. Siewiorek and R. S. Swarz. *The Theory and Practice of Reliable Design*. **16** (6), Dec-1987. pp. 1124 – 1134. Cited p. 17.

[Sunday Camera 1984] Trilogy abandons 'Superchip' plans. Boulder, CO: *Sunday Camera*. 12-Aug-1984. p. 10D. Cited p. 17.

[Toussaint 1992] G. T. E. Toussaint. Scanning the issue: computational geometry. (Introduction to the special issue). *Proceedings of the IEEE*. **80** (9), Sep-1992. pp. 1347 – 1363. Cited p. 9.

[Thomson 1917, 1992] D. W. Thompson. *On Growth and Form*. Cambridge University Press. Canto edition. 1992, J. T. Bonner, editor. Original edition published 1917. Cited p. 9.

[Ullman 1984] J. D. Ullman. *Computational Aspects of VLSI.* Rockville, MD: Computer Science Press. 1984. Cited pp. 6, 18.

[von Neumann 1958] J. von Neumann. *The Computer and the Brain.* New Haven and London, CT: Yale University Press. 1958. Eleventh printing: 1979. Cited p. 8.

[Wang et al. 2005] J. Wang, L. Li, S. H. Low, J. C. Doyle. Cross-layer optimization in TCP/IP networks. *IEEE / ACM Transactions on Networking*. **13** (3), Jun-2005. pp. 582 – 595. Online at http://netlab.caltech.edu/pub/papers/tcpip-ton2004.pdf. Cited pp. 8, 10.

[Williams 2008] R. S. Williams. How we found the missing memristor. *IEEE Spectrum*. **45** (12). Dec-2008. pp. 28 – 35. Online at www.spectrum.ieee.org/dec08/7024 and via ieeexplore.ieee.org. Cited p. 18.

