

## Quasiperiodic Route to Chaotic Dynamics of Internet Transport Protocols

Jian-Bo Gao,<sup>1</sup> Nageswara S. V. Rao,<sup>2</sup> Jing Hu,<sup>1</sup> and Jing Ai<sup>3</sup>

<sup>1</sup>*Department of Electrical and Computer Engineering, University of Florida, Gainesville, Florida 32611, USA*

<sup>2</sup>*Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*

<sup>3</sup>*Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, New York 12180, USA*

(Received 23 November 2004; published 18 May 2005)

We show that the dynamics of transmission control protocol (TCP) may often be chaotic via a quasiperiodic route consisting of more than two independent frequencies, by employing a commonly used ns-2 network simulator. To capture the essence of the additive increase and multiplicative decrease mechanism of TCP congestion control, and to qualitatively describe why and when chaos may occur in TCP dynamics, we develop a 1D discrete map. The relevance of these chaotic transport dynamics to real Internet connections is discussed.

DOI: 10.1103/PhysRevLett.94.198702

PACS numbers: 89.75.Hc, 05.45.Ac, 87.23.Ge, 89.20.Hh

The Internet is one of the most complicated systems that man has ever made. In recent years, two types of fascinating multiscale behaviors have been found in the Internet. One is the temporal-domain fractal and multifractal properties of aggregated network traffic flows [1,2]. The other is the spatial-domain scale-free topology of the Internet (see [3] and references therein). Also, it has been observed that the failure of a single router may trigger routing instability, which may be severe enough to instigate a route flap storm [4]. Furthermore, packets may be delivered out of order or even get dropped, and packet reordering is not a pathological network behavior [5]. All these point out the rich and complex dynamics of various facets of the Internet. As the next generation Internet applications such as remote instrument control and computational steering are being developed, another facet of complex behavior is beginning to surface in the form of transport dynamics. In order to sustain the needed control channels over wide-area connections, it has become increasingly important to understand the dynamics of the widely deployed transport protocol, the transmission control protocol (TCP). Typically, the transport dynamics over the Internet connections are a result of TCP's nonlinear additive increase multiplicative decrease (AIMD) congestion-control mechanism interacting with the stochastic Internet traffic. This complex interaction leads one to naturally expect the transport dynamics to be highly complicated.

For more than a decade, it has been conceived that the dynamics of complicated communication networks could be chaotic [6]. More recently there have been several research efforts [7–13] to better understand this issue, and yet several fundamental questions remain. In particular, the observation of chaos in [9], based on the widely used ns-2 network simulator [14], could not be repeated, which leads to the suspicion that these observed motions may, indeed, be periodic albeit with very long periods [15]. Our main objective is to systematically carry out the simulations in order to critically examine whether TCP dynamics are chaotic and, if yes, to identify a route to chaos. We present

Internet measurements as well to complement the simulation results.

TCP provides a reliable data transmission mechanism over Internet connections by utilizing (explicit or implicit) acknowledgments and loss inferences of data segments sent to the destination. It relies on two mechanisms to set its transmission rate: flow control and congestion control. The flow control ensures that the sender does not overflow the receiver buffer. The congestion control ensures that the sender does not unfairly overrun the available connection bandwidth. We assume that the flow window is appropriately chosen and the main dynamics are due to its congestion control. TCP maintains the congestion-control window size  $w$ , and sends data at a rate  $w/T$ , where  $T$  is the round trip time (RTT) of the connection. It works as follows:  $w$  packets are sent consequently, and then no more packets are sent, until acknowledgments arrive or loss is inferred, then  $w$  is updated. For simplicity, we consider the Tahoe version of TCP that adjusts  $w$  based on an acknowledgment of a data segment or inferred loss due to a timeout; our overall conclusions, however, hold for more recent AIMD TCP versions (such as new Reno). Ignoring the transient slow-start phase, TCP dynamics alternate between two phases,  $R_1$  consisting of only acknowledgments and  $R_2$  consisting of only inferred losses. The rule for updating  $w$  is as follows [16]:  $w \leftarrow w + 1/w$  when an acknowledgment is received in  $R_1$ ;  $w \leftarrow w/2$  when loss is inferred in  $R_2$ . More concretely, by assuming the flow control window to be fixed, we have the following  $w$ -update map:  $[1, W_{\max}] \rightarrow [1, W_{\max}]$  [10]:

$$M(w_i) = w_{i+1} = \begin{cases} w_i + 1/w_i & \text{if } w_i \in R_1, w_{i+1} \in R_1, \\ w_i/2^{n_i} & \text{if } w_i \in R_1, w_{i+1} \in R_2, \\ w_i/2^{n_i} & \text{if } w_i \in R_2, w_{i+1} \in R_2, \\ w_i + 1/w_i & \text{if } w_i \in R_2, w_{i+1} \in R_1. \end{cases} \quad (1)$$

In region  $R_1$ , there is no packet loss,  $n_i = 0$ , and the TCP is in additive increase (AI) phase. In  $R_2$ , there are  $n_i > 0$

packet losses, where  $n_i$  may be a complicated function of time, and the TCP is in multiplicative decrease (MD) phase. Note that when there are many competing TCP flows, each flow is described by Eq. (1).

We now examine TCP dynamics by systematically carrying out ns-2 simulations. In order to examine the temporal evolution of a dynamical system, one often measures a scalar time series at a fixed point in the state space. The Takens embedding theorem [17] ensures that the collective dynamic behavior of the dynamical variables coupled to the measured ones can be conveniently studied by the measured scalar time series. To illustrate the qualitative aspects of TCP dynamics by measuring only a scalar time series, it is most effective to consider a single bottleneck link, which tightly couples all flows through it. For simplicity, we shall follow the setup of Veres and Boda [9] consisting of long-lived TCP flows on a single link. This setup has four parameters: the link speed  $C$  in Mbps, the link propagation delay  $d$  in ms, the buffer size  $B$  in units of packets of 1000 bytes, and the number of competing TCP flows  $N$ . We have found that  $N$  acts as a critical bifurcation parameter in the route to chaos via quasiperiodic motions, as explained below.

For each of the  $N$  competing TCP flows, we record the time series  $W(i)$  corresponding to its congestion window size  $w$ . An example of quasiperiodic  $W(i)$  is shown in Fig. 1(a), with  $C = 0.1$  Mbps,  $d = 10$  ms,  $B = 10$ ,  $N = 3$ . Its power spectral density (PSD) is shown in Fig. 1(b). We observe many discrete sharp peaks when  $10^5$  points are used, but until then PSD has a flat spectrum. Thus, one might be tempted to interpret  $W(i)$  to be chaotic on smaller data sets. Therefore, the  $W(i)$  time series is not ideally suited for the study of (quasi)periodic motions with long periods. This motivates us to define a new time series,  $T(i)$ , which is the time interval between the onset of two successive MD (also called exponential back-off) regions as indicates in Fig. 1(a). The start of an exponential back-off indicates triggering of a loss episode. Thus  $T(i)$  is closely related to RTT since it is equivalent to the time interval between two successive loss bursts. The  $T(i)$  corresponding to  $W(i)$  of Fig. 1(a) is shown in Fig. 1(c), together with its PSD in Fig. 1(d). We observe that this  $T(i)$  is periodic, and, hence,  $W(i)$  is quasiperiodic. When  $N$  is increased to 17, the  $T(i)$  becomes very irregular, as shown in Fig. 1(e). While its broad spectrum of Fig. 1(f) suggests that  $T(i)$  is chaotic, this is indeed so, as will be explained soon.

Before we proceed, we note that it is possible that  $T(i)$  could be quasiperiodic. An example is shown in Fig. 1(g), with  $C = 0.5$  Mbps,  $d = 10$  ms,  $B = 10$  packets,  $N = 3$ , together with its PSD in Fig. 1(h). We have further extracted two new time series,  $T^{(k)}(i)$ , which is the interval time series between the successive local maxima of the  $T^{(k-1)}(i)$  time series, for  $k = 1, 2$ , and  $T^{(0)}(i) = T(i)$ . We have found that the  $T^{(2)}(i)$  is simply periodic, and hence  $W(i)$  is quasiperiodic with four independent frequencies. We also note that the ‘‘chaotic’’ congestion window size

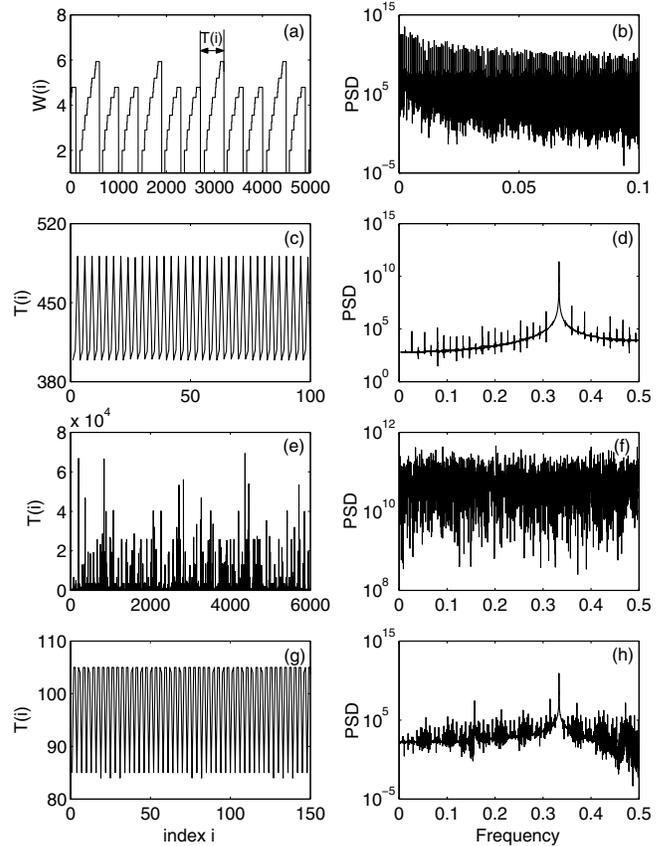


FIG. 1. (a)–(f) A quasiperiodic route to chaos, with  $C = 0.1$  Mbps, delay  $d = 10$  ms,  $B = 10$  packets, where a packet is of size 1000 bytes. (a) For  $N = 3$ , a quasiperiodic congestion window size  $W(i)$  time series sampled by an equal time interval of 10 ms. (b) The power spectral density (PSD) of (a). (c) The  $T(i)$  time series corresponding to (a). (d) The PSD of (c). When  $N = 17$ , the  $T(i)$  becomes irregular (e), with a white-noise-like PSD (f). (g),(h) show a quasiperiodic  $T(i)$  time series and its PSD, when  $C = 0.5$  Mbps,  $d = 10$  ms,  $B = 10$ , and  $N = 3$ .

data studied in [9] are, in fact, quasiperiodic with two independent frequencies. Furthermore, the probability distributions for the constant, periodic, and quasiperiodic  $T(i)$  are composed of only a few discrete peaks. This suggests that the observed quasiperiodic  $T(i)$  constitutes a discrete torus.

Let us now determine whether the irregular time series of Fig. 1(e) is chaotic or not. Note that  $T(i)$  often can be of the order  $10^4$  to  $10^5$ , which further establishes that a shorter  $W(i)$  time series is even more ill suited for the study of the underlying irregular dynamics. To study the chaotic nature of irregular  $T(i)$  time series, we employ the direct dynamical test for deterministic chaos developed in [18]. This is one of the more stringent tests for chaos, and has found applications in the study of the effects of noise on dynamical systems [19] and experimental time series [20]. To employ this method, we first construct vectors of the form  $V_i = [T(i), T(i + L), \dots, T(i + (m - 1)L)]$ , where  $m$  is the embedding dimension and  $L$  the delay time.

Since  $T(i)$  are interval time series, below we have simply chosen  $m = 6$ ,  $L = 1$  (very similar results are obtained when  $m = 8, 10$ , etc.). We then compute the  $\Lambda(k)$  curves defined by

$$\Lambda(k) = \left\langle \ln \left( \frac{\|V_{i+k} - V_{j+k}\|}{\|V_i - V_j\|} \right) \right\rangle. \quad (2)$$

The computation is carried out for a sequence of shells,  $r_i \leq \|V_i - V_j\| \leq r_i + \Delta r_i$ , where  $r_i$  and  $\Delta r_i$  are prescribed small distances ( $\Delta r_i$  is not necessarily a constant). The angle brackets denote the ensemble average of all possible  $(V_i, V_j)$  pairs, and  $k$  is called the evolution time. For true low-dimensional chaotic systems, the  $\Lambda(k)$  curves for different shells form a common envelope, and the slope of the envelope accurately estimates the largest positive Lyapunov exponent. For random systems, the  $\Lambda(k)$  curves corresponding to different shells do not form a common envelope, and hence the system under study cannot be interpreted as chaos [19]. The  $\Lambda(k)$  curves for  $T(i)$  of Fig. 1(e) are shown in Fig. 2(a), where the four curves, from bottom to top, correspond to shells of sizes  $(2^{-(i+1)/2}, 2^{-i/2})$ ,  $i = 14, 15, 16, 17$ . We observe a very well defined linear common envelope at the lower left corner of Fig. 2(a). The existence of a common envelope guarantees that a robust positive Lyapunov exponent can be obtained no matter which shell is used in the computation. Hence, the  $T(i)$  time series of Fig. 1(e) is, indeed, chaotic.

We have noted that the probability distributions for the  $T(i)$  time series of regular motions are composed of a few discrete peaks. What are the distributions for the chaotic  $T(i)$  time series such as shown in Fig. 1(e)? They are power-law-like, namely,  $P(T \geq t) \sim t^{-\beta}$ , for almost 2 orders of magnitude in  $t$ , as shown in Fig. 2(b). The exponent  $\beta$  is not a universal quantity, however.

To qualitatively explain TCP dynamics let us first modify Eq. (1) to a form that is amenable to analysis. This can be achieved by noticing that when TCP is in the AI phase, the window size must not exceed the maximally allowable window size. When this condition is not met, then TCP switches to the MD phase. These verbal descriptions can be expressed by the following map [21]:

$$M(w_i) = w_{i+1} = \begin{cases} w_i + 1/w_i & \text{if } \sum_{i=1}^N w_i \leq W_{\max}, \\ w_i/2^{n_i} & \text{if } \sum_{i=1}^N w_i > W_{\max}. \end{cases} \quad (3)$$

In Eq. (3), if  $W_{\max}$  is interpreted as the maximal allowable window size it can account for the effect of nonconstant flow window. Therefore, the modified map is more general than the map of Eq. (1).

When there are a large number  $N$  of competing TCP flows, it is often more convenient to lump the effect of the  $N - 1$  other TCP flows as background traffic. In such a scenario, if we focus on  $w_1$ , then  $\sum_{i=1}^N w_i \leq W_{\max}$  may be rewritten as  $w_1 \leq W'_{\max} = W_{\max} - \sum_{i=2}^N w_i$ . Therefore,  $W'_{\max}$  is determined by the background traffic, and typically is a complicated function of time, since the bottleneck bandwidth can vary with time considerably due to the dynamic background traffic. We have carried out the simulation of Eq. (3) by assuming  $W'_{\max}$  to be periodic and quasiperiodic, and have, indeed, observed chaos.

We are now ready to understand why chaos cannot occur when the number of competing TCP flows is small, such as 2. Let us assume that  $w_i$  and  $w_i + \Delta w_i$  as well as  $w_{i+1}$  and  $w_{i+1} + \Delta w_{i+1}$  are smaller than  $W'_{\max}$ . It is then easy to see that  $|\Delta w_{i+1}| \leq |\Delta w_i|$ , that is, small disturbance decays. This means during the AI phase of TCP, nearby trajectories contract instead of diverge. Hence, the dynamics are stable. In order for the dynamics to be unstable so that the Lyapunov exponent is positive, two conditions have to be met: (i) the transition from the AI phase to the MD phase has to be very frequent, and (ii) the time spent in the MD phase has to be significant. To satisfy these conditions, the number of competing TCP streams has to be necessarily large, and thus the network scenarios considered in [9] are not conducive to chaotic dynamics.

How relevant is the present study to the dynamics of the real Internet? Is the irregularity of  $T(i)$  shown in Fig. 1(e) an artifact of the Tahoe version of TCP we used, or more generally an artifact of the ns-2 simulator? To find an answer, we collected  $W(i)$  measurements between Oak Ridge National Laboratory (ORNL) and Louisiana State University (LSU) at millisecond resolution using the net 100 instruments using new Reno TCP. We computed  $T(i)$  as shown in Fig. 3(a), whose profile is qualitatively quite similar to Fig. 1(e). In fact, it also follows a (truncated) power law, as shown in Fig. 3(b). The exponent for the power law part is even smaller than the time series of Fig. 1(e). The truncation in the power law could be due to the shortness of the  $T(i)$  time series. As we have pointed out, the  $T(i)$  time series is related to RTT. Cottrell and Bullock [22] have been experimenting with many advanced versions of TCP, and also observed RTT time series very similar to these. We also have observed from RTT and loss data measured on geographically dispersed paths on the Internet (with a resolution of 1 s) by researchers at the San Diego Supercomputer Center that the probability distributions for the time interval between successive loss bursts typically follow a power-law-like distribution, with the exponent  $\beta$  also smaller than 1. Thus, we have good reason

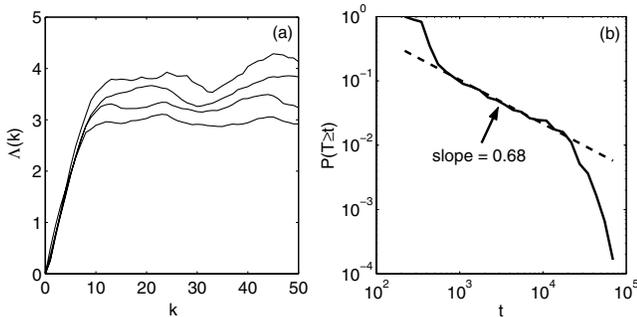


FIG. 2. (a) The  $\Lambda(k)$  curves and (b) the complementary cumulative distribution functions (CCDFs) for data of Fig. 1(e).

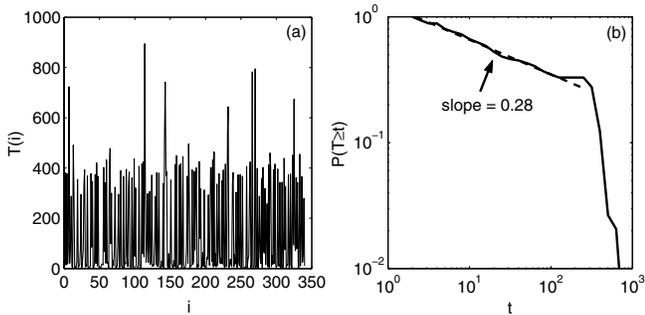


FIG. 3. (a) A  $T(i)$  time series extracted from a  $W(i)$  time series collected using net 100 instruments over the ORNL-LSU connection. (b) The CCDF for the  $T(i)$  time series.

to believe that the observed irregular  $T(i)$  time series is not an artifact of the ns-2 simulator, but reflects reality to some degree. In fact, the complicated dynamics described here is much simpler than the actual dynamics of the Internet, considering that there are all kinds of randomness in the Internet; in particular, typical TCP flows are not long lived, but vary considerably in durations determined by specific applications. Indeed, a detailed analysis of Internet measurements in [13] reveals that  $W(i)$  dynamics contain a dominant chaotic part but there is also an additional stochastic component due to the competing traffic.

There are a number of important implications of this study. (i) Conventionally, congestion control focuses mainly on steady state analysis and convergence to a steady state. Now that TCP dynamics can often be chaotic, one should also focus on the “transient” nonconvergent dynamics. (ii) Its dynamics being often chaotic, TCP is not the best choice for Internet control applications, for which transport dynamics have to be kept stable. One solution may be to guide the unstable dynamics back to a stable one, and here chaos control may be an effective strategy. (iii) To fundamentally eliminate unstable transport dynamics, new transport protocols (perhaps non-AIMD approaches such as stochastic approximation methods of [23]) have to be developed and their stability has to be carefully studied using chaos theory.

This research is sponsored by Defense Advanced Research Projects Agency under MIPR No. K153, National Science Foundation under Grants No. ANI-0229969 and No. ANI-0335185, and by Engineering Research Program and High-Performance Networking Program of Office of Science, U.S. Department of Energy under Contract No. DE-AC0500OR22725 with UT-Battelle, LLC.

[1] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, *IEEE/ACM Trans. Netw.* **2**, 1 (1994); V. Paxson and S. Floyd, *IEEE/ACM Trans. Netw.* **3**, 226 (1995);

- M.E. Crovella and A. Bestavros, *IEEE/ACM Trans. Netw.* **5**, 835 (1997); W. Willinger, M.S. Taqqu, M.S. Sherman, and D.V. Wilson, *IEEE/ACM Trans. Netw.* **5**, 71 (1997).
- [2] J.B. Gao and I. Rubin, *Comput. Commun.* **24**, 1400 (2001); *Nonlinear Anal.* **47**, 5765 (2001); *Int. J. Commun. Syst.* **14**, 783 (2001).
- [3] R. Albert and A.L. Barabasi, *Rev. Mod. Phys.* **74**, 47 (2002).
- [4] C. Labovitz, G.R. Malan, and F. Jahanian, *IEEE/ACM Trans. Netw.* **6**, 515 (1998).
- [5] J.C.R. Bennett, C. Partridge, and N. Shectman, *IEEE/ACM Trans. Netw.* **7**, 789 (1999).
- [6] A. Erramilli and L.J. Forays, in *Proceedings of the ITC-13, Copenhagen, 1991* (North-Holland, Amsterdam, 1991).
- [7] R. Srikant, *The Mathematics of Internet Congestion Control* (Birkhäuser, Boston, 2004).
- [8] C. Li, G. Chen, X. Liao, and J. Yu, *Chaos Solitons Fractals* **19**, 853 (2004).
- [9] A. Veres and M. Boda, in *Proceedings of the IEEE INFOCOM 2000* (IEEE, Piscataway, NJ, 2000), p. 1715.
- [10] N.S.V. Rao and L.O. Chua, in *Proceedings of the Workshop on Signal Processing, Communications, Chaos and Systems*, Newport, RI, 2002 (unpublished).
- [11] P. Ranjan, E.H. Abed, and R.J. La, *IEEE/ACM Trans. Netw.* **12**, 1079 (2004).
- [12] N.S.V. Rao, J. Gao, and L.O. Chua, in *Complex Dynamics in Communication Networks*, edited by L. Kocarev and G. Vattay (Springer, Berlin, 2005).
- [13] J. Gao and N.S.V. Rao, *IEEE Commun. Lett.* **9**, 4 (2005).
- [14] See <http://www.isi.edu/nsnam/ns/>.
- [15] D.R. Figueiredo, B. Liu, Vishal Misra, and D. Towsley, University of Massachusetts, Report No. 00-55, 2002.
- [16] L.L. Peterson and B.S. Davie, *Computer Networks* (Morgan Kaufmann, San Francisco, CA, 2003).
- [17] F. Takens, in *Dynamical Systems and Turbulence*, Lecture Notes in Mathematics Vol. 898, edited by D.A. Rand and L.S. Young (Springer-Verlag, Berlin, 1981), p. 366.
- [18] J.B. Gao and Z.M. Zheng, *Phys. Lett. A* **181**, 153 (1993); *Phys. Rev. E* **49**, 3807 (1994); *Europhys. Lett.* **25**, 485 (1994).
- [19] J.B. Gao, S.K. Hwang, and J.M. Liu, *Phys. Rev. A* **59**, 1582 (1999); *Phys. Rev. Lett.* **82**, 1132 (1999); J.B. Gao, C.C. Chen, S.K. Hwang, and J.M. Liu, *Int. J. Mod. Phys. B* **13**, 3283 (1999); K. Hwang, J.B. Gao, J.M. Liu, *Phys. Rev. E* **61**, 5162 (2000).
- [20] R. Bandyopadhyay and A.K. Sood, *Europhys. Lett.* **56**, 447 (2001); R. Bandyopadhyay, G. Basappa, and A.K. Sood, *Phys. Rev. Lett.* **84**, 2022 (2000); S. Venkadesan, M.C. Valsakumar, K.P.N. Murthy, and S. Rajasekar, *Phys. Rev. E* **54**, 611 (1996).
- [21] J.B. Gao and N.S.V. Rao, “*Synchronized Oscillations, Quasi-Periodicity, and Chaos in TCP*,” Department of Electrical and Computer Engineering, University of Florida, Technical report, 2002.
- [22] U.S. Department of Energy Science UltraNet Kickoff Meeting, <http://www.csm.ornl.gov/ultranet/package/Cottrell/Cottrell.htm>.
- [23] N.S.V. Rao, Q. Wu, and S.S. Iyengar, *IEEE Commun. Lett.* **8**, 66 (2004).