network physics

controlling the cloud

chaos and the physics of the Internet

table of contents

This white paper is intended for decision-makers and industry influencers as an introduction to chaos theory, its role in understanding the physics of the Internet, and how Network Physics applies it to deliver a 30% improvement in response time and a 20-50% reduction in variability for all traffic across both public and private networks—using a data center appliance that does not require any modification of the network or cooperation from other network elements.

Introduction

Enterprise networks are becoming a hybrid of private, VPN, and public networks, driven by the lower cost, global reach, ease of use, and scalability of the Internet. Responding to that trend, network providers have spent billions of dollars on the Internet infrastructure. But the Internet is still not reliable, still plagued by unpredictable bursts of congestion and delay, and still does not supply the speed and predictability that businesses need to support mission-critical applications. Forced to over-provision or lease expensive private lines in an attempt to overcome these problems, enterprises are spending too much on bandwidth.

The problem is that the fundamental protocols that control the Internet are blind to its "hidden physics:" the collective behaviors that arise from the interaction of the myriad links, switches, and traffic flows that comprise complex IP networks. And what you cannot see, you cannot control, or even use efficiently. This is especially true of the Transmission Control Protocol (TCP), which operates on a flow-by-flow basis and so does not account for the effects of millions of other flows impacting the many routers along any path from server to client. It is also true of new and emerging protocols such as MPLS and IPv6.

One solution lies in a body of knowledge that has come to be called the science of complexity: the study of the collective behaviors that arise from the interactions between the elements of complex systems. Over the past 20 years, a branch of this science called chaos theory has been especially fruitful in describing and controlling complex systems as apparently disparate as lasers and the human heart. Now Network Physics is applying it to controlling the Internet, in the form of chaos control algorithms that, applied to aggregated flows by a single device in the data center, can significantly reduce the impact of Internet congestion on those flows.

The result is faster, less variable performance, without re-routing, and without the cooperation of any downstream elements or network, or any client application. And, given this improvement, an enterprise can reduce bandwidth costs by moving its applications from leased lines to the Internet, avoiding additional multi-homing, and/or reducing its reliance on premium ISPs.

These algorithms are but part of a body of knowledge Network Physics is developing about the real-time physics of the Internet. That knowledge can also be used to improve the performance of traffic flows that are not congested, delivering an overall 30% improvement in response time for all traffic, and reducing variability by 20-50%.

This white paper is intended to give decision-makers and industry influencers a basic understanding of the science behind these chaos control algorithms and their application to enterprise networks for improving performance and reducing bandwidth costs. It is designed for a broad range of understanding, so that those desiring a "view from 50,000 feet" may skip the more technical details while still gaining useful knowledge.

We'll start with a brief overview of complexity science and chaos theory, leading to a discussion of how chaos can be controlled by understanding the collective behaviors that underlie it. From there we'll look at one of the primary sources of chaotic behavior on the Internet and how it offers an opportunity for control to reduce congestion and improve end-to-end performance. We'll also briefly look at how a knowledge of the collective behaviors of the Internet can be used to improve the operation of TCP for uncongested flows, and its response to congestion.

Complexity and Chaos

While there is no general agreement on a definition of "complexity science," it is concerned with the collective behaviors of complex systems. Some theorists speak of "emergent order" as another way of pointing out that systems of simple elements can give rise to surprisingly complicated activities. In other words, the science of complexity maps fairly well to a familiar saying: "The whole is greater than the sum of its parts."

While this may seem like mere common sense, it actually flies in the face of three centuries of Newtonian determinism as famously stated by the 18th-century mathematician Laplace. He asserted that given sufficient information, the behavior of even the most complex system could be determined with whatever degree of accuracy was desired.[1] This implies that controlling a complex system is merely a matter of having sufficient information about the behavior of its individual elements; in other words, if you throw enough processing power at a problem you can solve it. It was chaos theory that finally destroyed this comfortable assumption.

By the early 20th century, cracks were beginning to appear in the Laplacian edifice, but it wasn't until 1963 that a meteorologist named Edward Lorenz at MIT blew it apart—and even then, what he had discovered wasn't recognized until a decade later.[2] Lorenz had adapted another scientist's work on convection in the atmosphere by simplifying it into a set of three differential equations. The behavior of this system is interesting: after a number of iterations the value of the y variable goes into apparently random oscillations.

The system's behavior, however, is not patternless. Plotting the equations in a certain way gives Figure 1 (page 3), which is now called the Lorenz attractor (more about that later). Neither is the system's behavior unpredictable. Plotting each instance of z

---

1. Pierre Simon de Laplace, *Philosophical Essays on Probabilities*: "An intellect which at any given moment knew all the forces that animate Nature and the mutual positions of the beings that comprise it, if this intellect were vast enough to submit its data to analysis, could condense into a single formula the movement of the greatest bodies of the universe and that of the lightest atom: for such an intellect nothing could be uncertain; and the future just like the past would be present before its eyes."
2. Good accounts of the development of chaos theory can be found in *Does God Play Dice: the Mathematics of Chaos* by Ian Stewart. (Blackwell Publishers, 1989) and *Chaos: Making a New Science* by James Gleick (Penguin, 1988).

versus the previous instance gives a precise curve with a spike in the middle: given any value of z in the evolution of this system, you can predict the next value

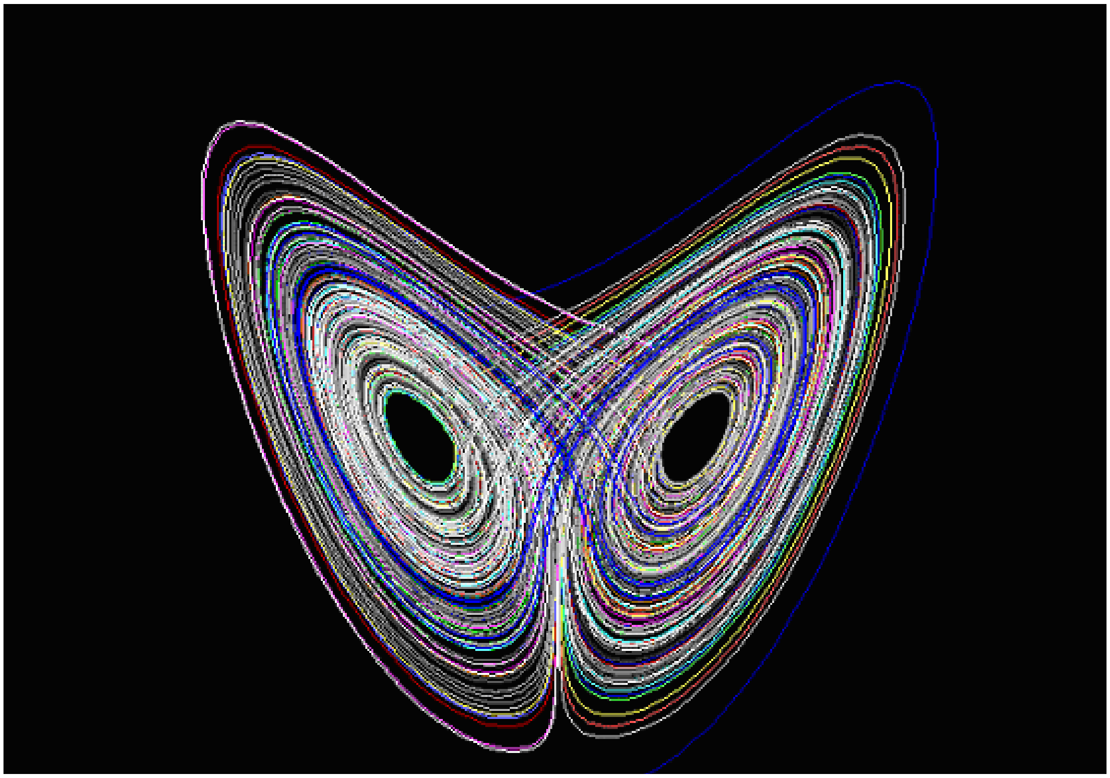

Figure 1: The Lorenz Attractor

But, as Lorenz found out, that doesn't mean that you can then string these predictions together to predict the system's long-term evolution. Instead, entirely by accident, Lorenz found that the system was highly sensitive to initial conditions: a vanishingly small difference in the initial values leads to wildly different results. Contra Laplace, no matter how precisely you specify the starting conditions of this system, after some period of time related to that precision, the state of the system is not predictable![3] In Newtonian terms, this was a very upsetting development: an entirely deterministic system (each state is entirely determined by the preceding state—there is nothing of chance in it) whose final state is indeterminate!

Given the impact this had on science and mathematics, it's perhaps no surprise that such systems have come to be called chaotic. For quite some time many scientists were skeptical that chaotic systems were anything more than mathematical curiosities without application to the real world. But by the beginning of the 1980s the reality of chaotic behavior in nature was confirmed, and it is now known that chaotic systems are common in nature and human affairs.

---

3. Stewart, p.140

The weather and rush hour traffic are examples that everyone is familiar with. In traffic, for instance, someone braking slightly on a freeway can create stop-and-go conditions a mile back. This is because some of the drivers behind the braking car slow down far more than is necessary, and some less, and so forth back up the line—their responses are dynamic, nonlinear, and add up in unpredictable ways. Despite this, simple adjustments in the flow of cars onto a freeway, e.g., using metering lights, can greatly improve traffic flow. Although on-ramp metering of this sort is not actually chaos-based, this example shows in principle that quite complex systems can be nudged in a favorable direction with very simple perturbations.

Controlling Chaos

"Controlling chaos" sounds like an oxymoron. One might reasonably ask: how does one control what one cannot predict? However, as you can see from the Lorenz system of equations above, when physicists talk about chaos, they are describing a certain kind of complex nonlinear system that is highly sensitive to initial conditions. They don't mean one that is random, lacks order, or is fundamentally unpredictable. The behavior of a chaotic system may be impossible to understand or control using traditional techniques, but it has an underlying order that can be exploited. And, because of its sensitivity, control can be accomplished using very small perturbations.

One way that order manifests itself is in chaotic attractors (or "strange attractors") such as the one illustrated above. Attractors are a feature of the qualitative theory of dynamical systems originally devised by Henri Poincare at the beginning of the 20th century. This theory enables physicists to model the behavior of dynamical systems by describing them in topological or geometric terms, a powerful technique that can render quite complex systems more amenable to description and control.

Poincare's scheme represents the state of a dynamical system at a given moment as a point in an n-dimensional phase space, where n is basically the number of variables necessary to specify the instantaneous state of the system. The trajectory of that point through the phase space models the evolution of the system over time, and allows many extremely powerful topological and geometric techniques to be applied to understanding it.

For example, the state of the system of equations studied by Lorenz, a "toy" simplification of atmospheric dynamics, is specified by three variables, x, y, and z. The Lorenz attractor is thus a highly complex trajectory in a three-dimensional phase space.[4] It's called an attractor because the state of the system—its trajectory—appears to be attracted to that part of phase space—the system tends to be in a state close to the attractor. That fact adds a number of powerful techniques from statistical mechanics to the physicist's tool box for dealing with chaos.

4. The illustration above is a two-dimensional view of the attractor seen in the x-z plane.

The attractor concept also hints at how the control of chaos is possible. The formal process can be stated with deceptive simplicity:

1. Find the attractor that represents the usual state of the system;
2. Find an attractor that represents a state closer to what you want;
3. Perturb the system appropriately to nudge the system from the first attractor to the second.

Easily said, hard to do. Beyond the difficulty of choosing the appropriate phase space (what variables really count, or play a role in the observed chaotic behavior?), a particular obstacle is the "fractal dimension" of the system. This does not refer to the size of the phase space but is more a measure of entropy—how much information you need to have about the system to accurately describe, predict, and control it. High-dimensional chaotic systems are not amenable to control. Examples of such systems include the weather and the stock market—their fractal dimension is thought to be approximately eight.

However, many chaotic systems exhibit what are called "collective behaviors" that have the effect of reducing their dimensionality. In collective behavior, the behavior of each unit in a complex system is dominated by the influence of its neighbors. Such systems exhibit a variety of ordering phenomenon that make them amenable to control, if you can decide the proper shape of the perturbation necessary to "nudge" the system into a favorable state. Dr. Duong-van's research into the collective behaviors of lasers and the human heart enabled him to demonstrate that these systems exhibit attractors much like the Lorenz, with a dimensionality near two, which makes them amenable to control.

A laser is a nonlinear system that turns energy into highly organized, coherent light. However, if too much energy is pumped into a laser, its output loses coherency and it becomes nothing more than a very bright searchlight. While at Lawrence Livermore Labs, Dr. Duong-van developed a theoretical technique involving precise perturbation of the light input that significantly increased laser efficiency and output. Likewise, the human heart can lose "coherency," a phenomenon known as fibrillation where, instead of all its parts beating in synchrony, they oscillate out of step. A person in ventricular fibrillation has only minutes before permanent brain damage or death ensue—all too often not enough time to find an external defibrillator. Dr. Duong-van's techniques, used in an implantable heart defibrillator, enable it to automatically administer a small jolt of electricity at exactly the right moment to shock the heart back into proper operation.

The Chaotic Internet

At the beginning of the 90s, researchers proved that the Internet exhibits fractal, or self-similar behavior.[5] That is, no matter what time scale you look at—minutes, hours,

---

5. K. Park and W. Willinger. "Self-similar network traffic: An overview" . Chapter 1 of Self-Similar Network Traffic and Performance Evaluation, Wiley-Interscience, 2000.

days, weeks, months, or longer—the traffic varies unpredictably, unlike the phone system, where traffic averages out over time, which would be the case were packet arrival rates random

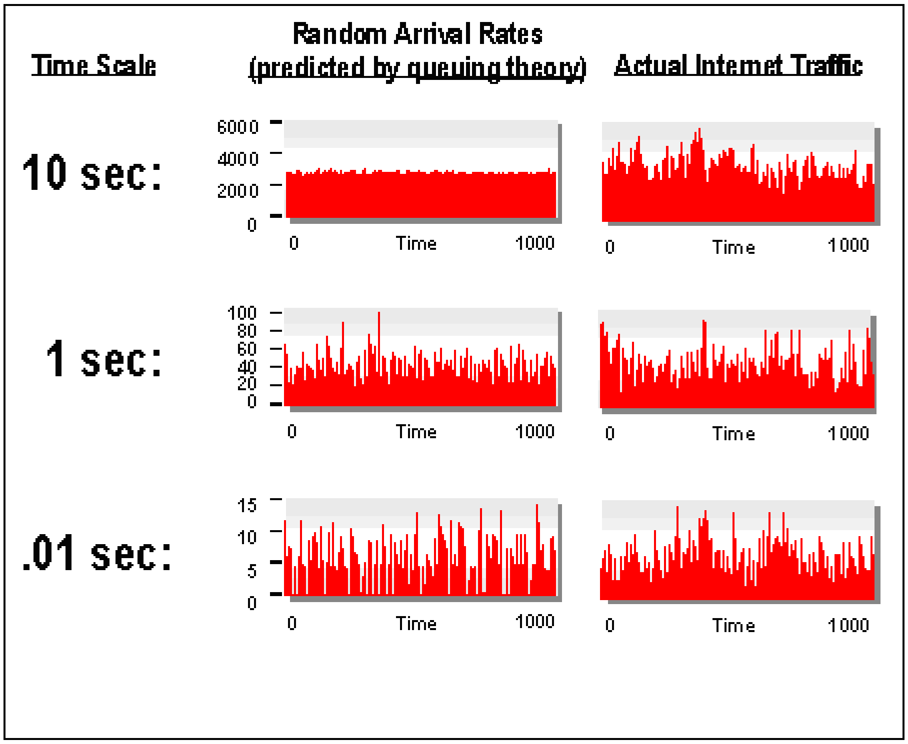

Figure 2: Random and Self-Similar Traffic

Although few people realized it at the time, this is a clue pointing to the chaotic nature of Internet dynamics.[6] By the end of the decade, researchers demonstrated that the Internet's fractal nature is due to the chaotic behavior of the Transmission Control Protocol (TCP)—that the Internet is a chaotic system.[7]

If that were all there is to it, that would be the end of the story. If you take the traditional, packet-oriented view of the network, the Internet is a chaotic system of a high fractal dimension, and decidedly not amenable to control. But an understanding of the collective behaviors that are the foundations of chaos on the Internet reveals, as with lasers and the human heart, an attractor with a dimensionality of about two. These collective behaviors arise from the correlation imposed on multiple TCP flows, each under the control of the non-linear TCP algorithms, by the router queues through which they pass.

6. A chaotic attractor is called a strange attractor in part because it is fractal—there is an intimate association between fractals and chaos. For instance, as noted above, the fractal dimension of a chaotic system determines its controllability.
7. A. Veres and M. Boda.. "The Chaotic Nature of TCP Congestion Control". Paper presented at IEEE Infocom 2000. (http://www.ieee-infocom.org/2000/papers/74.pdf)

Although the mathematics involved are complex, it's fairly simple to sketch out an intuitive picture of how this correlation of TCP's non-linear dynamics happens.[8] What follows here is a brief overview of the operation of the Internet to set the scene for chaos and control.

Packets and Routers

Most data intended for transmission across the Internet is broken up by the originating machine into packets in accordance with the TCP/IP protocols. IP governs where the packets are going, while TCP guarantees that the packets will get there. All packets are IP-based (have an IP header), and there are delivery protocols other than TCP that do not guarantee delivery (e.g., UDP). However, most of the traffic on the Internet is composed of TCP flows, exchanges of packets between two computers whose rate of transmission is controlled by the non-linear TCP flow-control algorithms.

The Internet infrastructure is a mesh of "pipes" or links of different traffic capacities (bandwidth) connected by routers, packet-switching devices that generally have multiple inputs and outputs. Each possible path through a router, from input *x* to output *y*, has a queue associated with it, which holds packets until the output link can accept it. When a packet appears at one of a router's input ports, the router decides which output port to switch it to—in effect the packet's next "hop" across the Internet-based on an internal routing table, which is kept current by one of various routing protocols, of which the most important is the Border Gateway Protocol (BGP). Thus, the packets in a flow do not follow a predetermined path through the Internet; they may not even all follow the same path. This hop-by-hop forwarding process makes the Internet highly robust and scalable.

Unfortunately, the packet-switched nature of the Internet that makes this routing technique possible also makes the Internet subject to bottlenecks-load or traffic mismatches—because the routers in such a network don't handle flow control. That's the responsibility of the computers that are communicating via TCP. If more packets come into a router than go out, due to changing patterns of demand, a bottleneck occurs: congestion. Congestion first manifests itself as queuing delay, as packets spend more and more time in the router's buffers waiting for the packets ahead of them to be sent on. When the queue is full, the router starts throwing packets away and actual packet loss occurs.

Packet loss in particular has tremendous impact on web application performance. Figure 3 (page 8) shows the network response time measured over a five-hour period on a very large web site. The blue trace is the response time for all flows not experiencing any packet loss; the red trace is all flows experiencing any packet loss at all. The difference is striking: without packet loss performance is much faster and

8. There are other dynamics on the Internet, as well, one of which is discussed in Appendix B.

highly consistent. With packet loss, performance degrades appreciably and is highly variable. Both consequences are unacceptable for mission-critical applications.

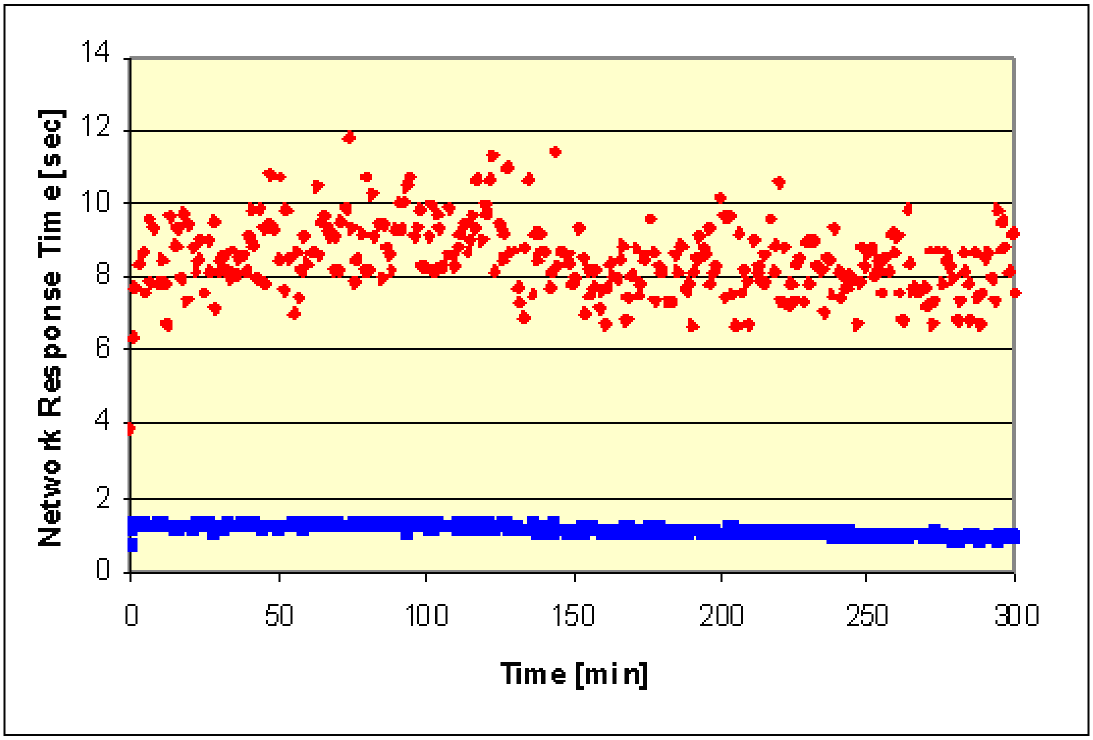

Figure 3: Network Response Time for Flows with and without Packet Loss

Congestion can happen for two reasons, and is sometimes predictable, sometimes not. First, whenever traffic is switched from a higher bandwidth input to a lower bandwidth output, there will be a bottleneck. This is common, for instance, where a regional or local ISP gets traffic from the backbone. But even if all the inputs and outputs to a router are of equal capacity, a bottleneck will still result whenever the net output to a given pipe exceeds its capacity because traffic from multiple input pipes is being switched to it.

Some of these bottlenecks are well known, such as the ones at major peering points where the backbones exchange traffic. Others are less predictable, or entirely unpredictable and can arise anywhere. More importantly, the overwhelming majority of congestion episodes are only seconds or minutes in length, as shown in Figure 4.[9]

9. This is the primary reason that route-selection technologies (so-called "intelligent routing) can only address a small part of the Internet performance problem.

No amount of network engineering can design them out because they arise from the chaotic nature of the Internet.

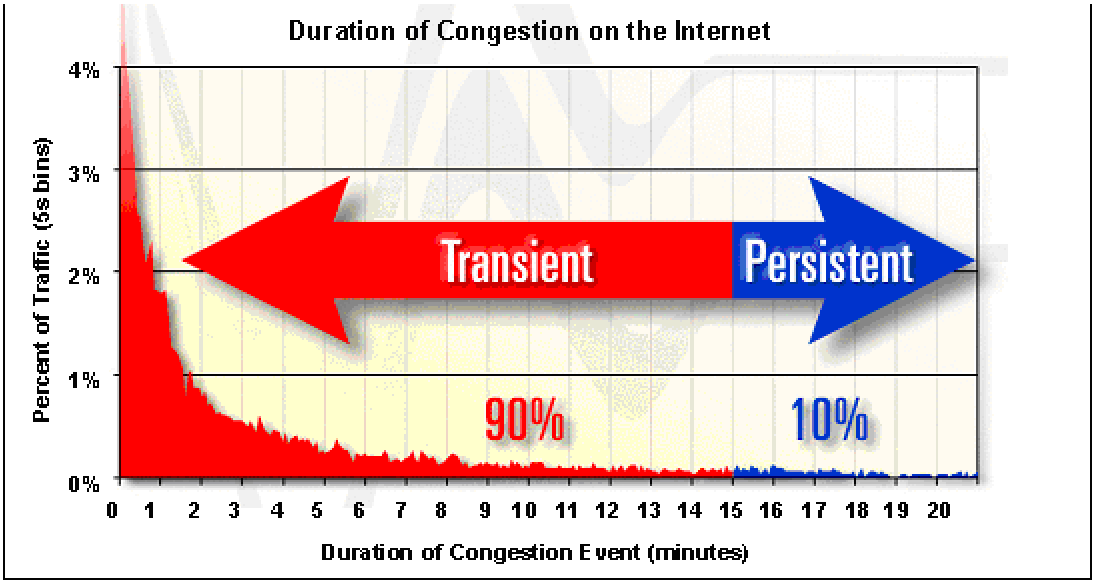

*Source: Yahoo Traffic (one hour sample)

Figure 4: Distribution of Internet Congestion Episode Durations

TCP Flow Control and Chaos

These episodes arise from the many non-linearities of the Internet, ranging from the unpredictable nature of user demand to the Transmission Control Protocol's non-linear flow control algorithms, which is what we'll focus on. TCP operates in the computers (e.g. web or application servers and browsers) that are communicating across the Internet. It has two responsibilities: to reassemble IP packets that it receives into the proper order using sequence numbers, and control the rate at which the packets are sent: flow control. Flow control is designed to control congestion along whatever end-to-end path results from the switching decisions of the various routers. It does so by gauging how much traffic that path can handle and controlling the transmission rate of the sending computer accordingly.

Flow control starts once a client contacts the server (for instance, when a browser user clicks on a URL) and TCP performs a "handshake" exchange of special packets to set up a two-way connection. This enables the client to signal when it has received a packet (or a burst of packets) using a special ACK packet that references the received packet's sequence number. The sender's TCP stack uses these ACKs to detect and respond to congestion by changing the number of packets it sends at one time. It can roughly gauge queueing delay as router buffers fill up by timing ACKs to detect changes in round trip time (RTT). And, if the sender doesn't get an ACK within a specified time (the timeout period, which is adjusted in accordance with the RTT), or gets a repeated ACK (something the receiver does if the series of packets it's getting skips a sequence number), it assumes the packet or group of packets it last sent was lost and retransmits them.

TCP also has to gauge the capacity of the end-to-end path from server to client. However, the TCP flow-control algorithms apply only to individual flows. TCP is largely blind to the collective behavior of TCP flows as correlated by router queues.[10] For instance, it can't see the side traffic that competes for queue space in the routers along the route from server to client, nor is it aware of the "congestion resonances" that indicate a congested router queue, as discussed below. Without this knowledge, TCP's only way to gauge path capacity is to actually try to, in effect, cause congestion by increasing the number of packets transmitted in each burst until packet loss occurs. This process is called Slow Start: first transmitting two packets, waiting for the ACK, then four packets, wait, eight packets, and so forth. When packet loss finally occurs, TCP abruptly reduces the number of packets it sends and then begins the Slow Start process all over again as it begins to retransmit the discarded packets.[11] The rate of transmission from a given computer thus looks a little like the edge of a saw, but more irregular (Figure 5).

This is classical non-linear behavior, and is one reason the Internet is fractal: that no matter on what time scale you consider Internet traffic, the irregular saw-edge pattern looks pretty much the same: bursty and unpredictable. It is also a fundamental cause of chaos on the Internet.

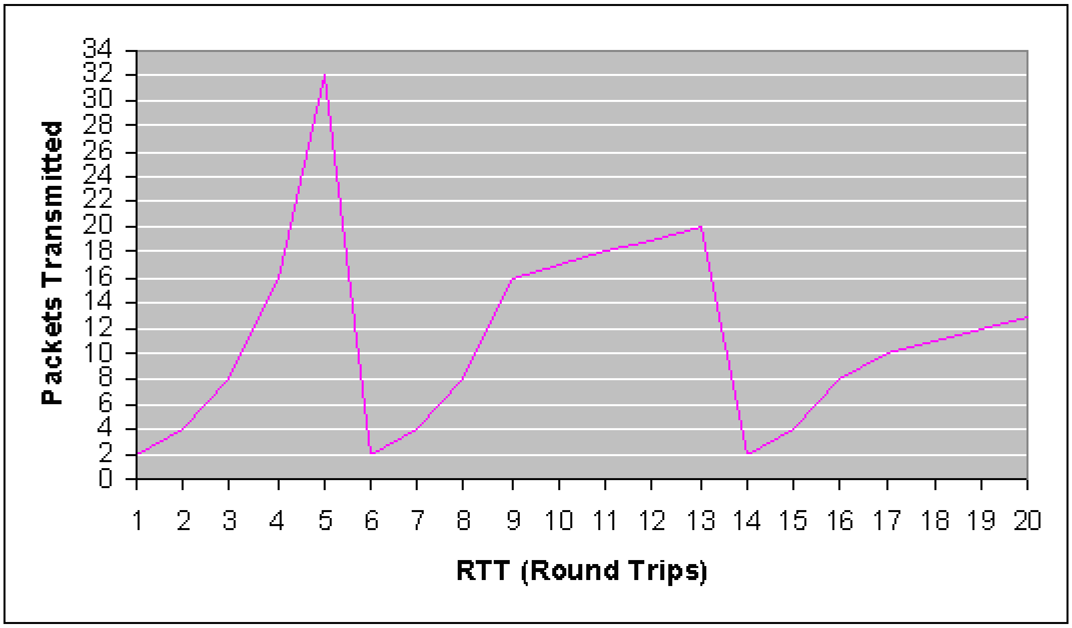

Figure 5: TCP Flow Control

10.Some TCP implementations see/impose certain types of collective behavior, such as TCP Pooling, discussed below.

11.The actual operation is rather more complex: see Appendix A for a more detailed description.

Collective Behaviors on the Internet

This pattern, of course, affects every other TCP flow using the same router queue, because the flows compete for buffer space. This is why correlation occurs between flows. Correlation thus tends to synchronize the TCP back-off-and-resend behavior of all those flows, piling the peaks in traffic in each flow on top of each other to create greater and greater variations in traffic. The TCP flows are thus said to "burst together."

This selfsame correlation (as well as correlations between routers, as discussed in Appendix B) also creates collective behaviors and ordering phenomena that lower the dimensionality of the system, revealing a chaotic attractor like those of the Lorenz equations, lasers, or the human heart, making the system controllable. Just as the heart can go into useless fibrillation where the various parts fail to cooperate, so it is with the Internet, and just as with the heart, it's possible to "tune" the Internet for better performance using very small changes.

Chaos Control on the Internet

The foregoing review has put all the pieces in place needed to understand, on a basic level, the operation of Network Physics' chaos control algorithms in terms of the formal process outlined earlier.

1. Find the attractor that represents the congested state of the system. In this case the system is a set of TCP flows from server(s) to client(s) and the routers they pass through. Only a single device in the data center is needed to control Internet performance because the TCP flows themselves, aggregated appropriately, contain all the information (the resonances) needed to characterize their collective chaotic behavior.
2. Find an attractor that represents a state closer to what you want. Along with Step One, this is the heart of our intellectual property, described by a number of pending patents.
3. Perturb the system appropriately to nudge the system from the first attractor to the second. This is simplicity itself. Based on the real-time information gathered in Step One, our adaptive control technology times the transmission of individual packets to minimze the probability that they will be lost in a congested router queue.

The chaos-theoretical view of this operation in a simple test network is illustrated below.

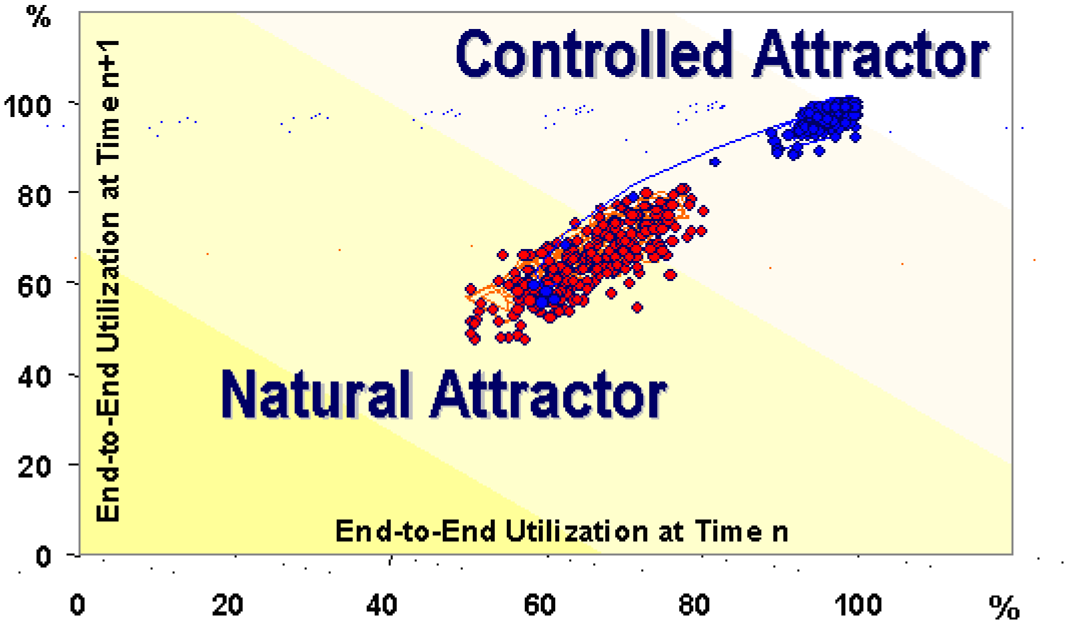

Figure 6: Chaotic Attractors in a Simple Network

The two axes of the graph are both time: the x-axis is time n, and the y-axis is time n+1. Thus, the large red natural attractor represents highly variable and slower performance, while the smaller blue controlled attractor represents less variable and faster performance.

An example of the effects of this careful perturbation of the system in terms of network performance can be seen in the illustrations below, this time taken from actual operation at a large hosting center. Overall we see a 30% reduction in response time and a 20-50% reduction in variability: faster, more consistent performance for web-enabled applications.

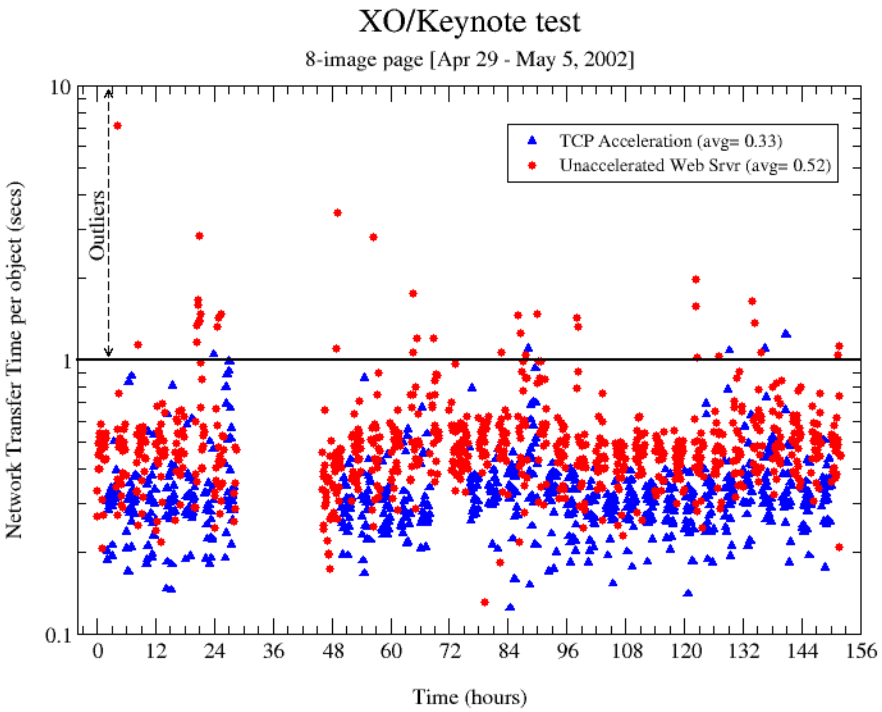

Figure 7: Chaotic Flow Control Improvements

Making TCP Smarter

The real-time data exploited by the chaos control algorithms can also be used to optimize the performance of flows not encountering congestion by modifying the operation of TCP's flow control algorithms. Network Physics does this using one of two techniques, one chaos-based and several based on IETF standards and RFCs.

We use chaos control theory to safely accelerate TCP's Slow Start behavior for uncongested flows[12] by increasing the initial congestion window, and within that window, applying the chaos control algorithms to pace individual packets to avoid congestion. In many cases, especially for the HTTP traffic characteristic of web-

12.Known by remembering state in a process notionally similar to TCP Pooling as discussed below.

enabled applications, this means that an entire file (TCP flow) can be transmitted in one RTT following the handshake, making better use of available bandwidth.

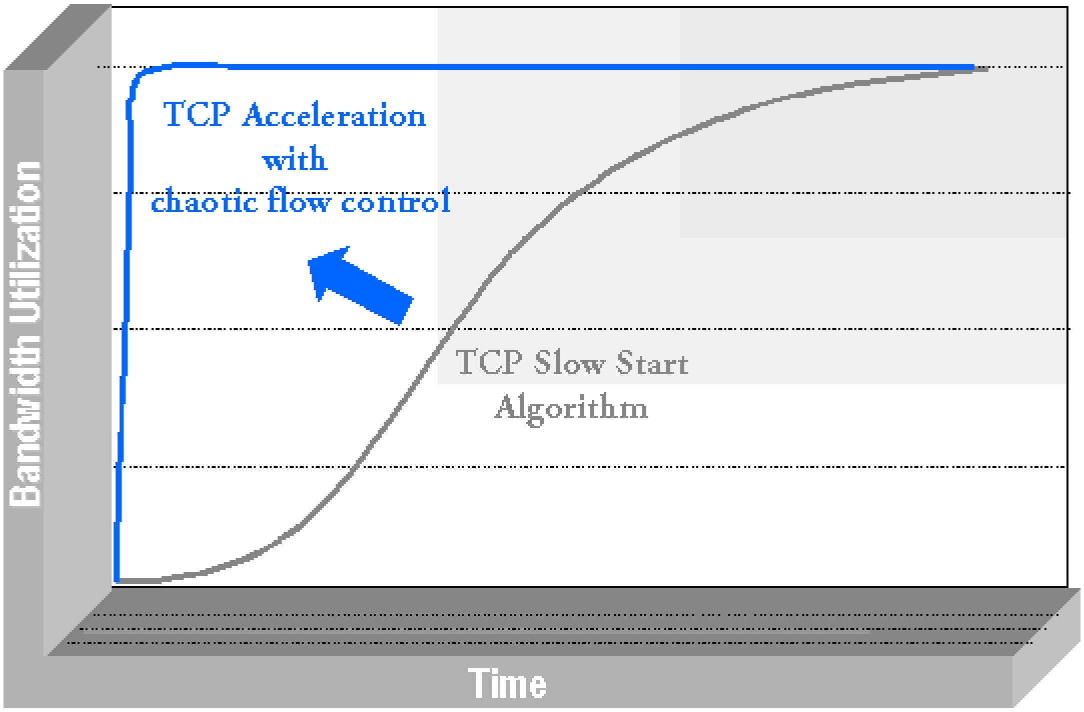

Figure 8: Chaotic Flow Control Improves Bandwidth Utilization

The data is also applied using standards-based TCP algorithms. These include SACK and TCP Pooling. Selective Acknowledgement (RFC2018 et al) improves TCP's response to and recovery from congestion episodes. SACK is supported by a growing number of clients: all current Windows versions and many others. It is automatically disabled for non-compliant clients. TCP Pooling (RFC2140) improves TCP's knowledge of flow conditions and requires no client support.

Conclusion

A brief document such as this white paper cannot hope to do more than outline the potential a physics-based approach to improving Internet performance holds for realizing the true potential of the global network. Perhaps the greatest strength of this approach is that it will always yield improvements in performance, regardless of advances in hardware or protocols, for the Internet, as a complex network, will always be chaotic, and will always exhibit collective behaviors that can be exploited.

What is required is an unflagging dedication to uncovering and understanding the physics of the network so that we can apply the knowledge gained in other areas of science to improving its operation and make it truly a platform for performance as well as connectivity.

This is not the task of Network Physics alone, but of many different researchers throughout the world in the years to come. We are but pioneers in the application of complexity science to the Internet, so it seems fitting to close with the words of another pioneer, one of the greatest physicists the world has ever known, speaking of the experience that is common to all scientists:

> If I have seen further, it is by standing on the shoulders of giants.
>
> *Sir Isaac Newton*

Appendix A: How TCP Flow Control Works

There are a number of parameters tracked by TCP in the sender, of which the most important for the purposes of this brief description are:

1. The *congestion window*, which is the maximum amount of un-ACKed data sent (i.e. the maximum amount before getting an ACK packet back from the receiver);
2. The *slow start threshold*, which is where slow start's geometric or multiplicative increase in the congestion window ceases and TCP goes into what is called *congestion avoidance*, in which the increase in the congestion window is linear.
3. The *maximum advertised window*, which is set by the receiver during the set-up of the TCP connection to advise the sender of how much data it is the receiver is willing to buffer while re-ordering the packets (in accordance with the sequence numbers) and waiting for the application to accept the data.

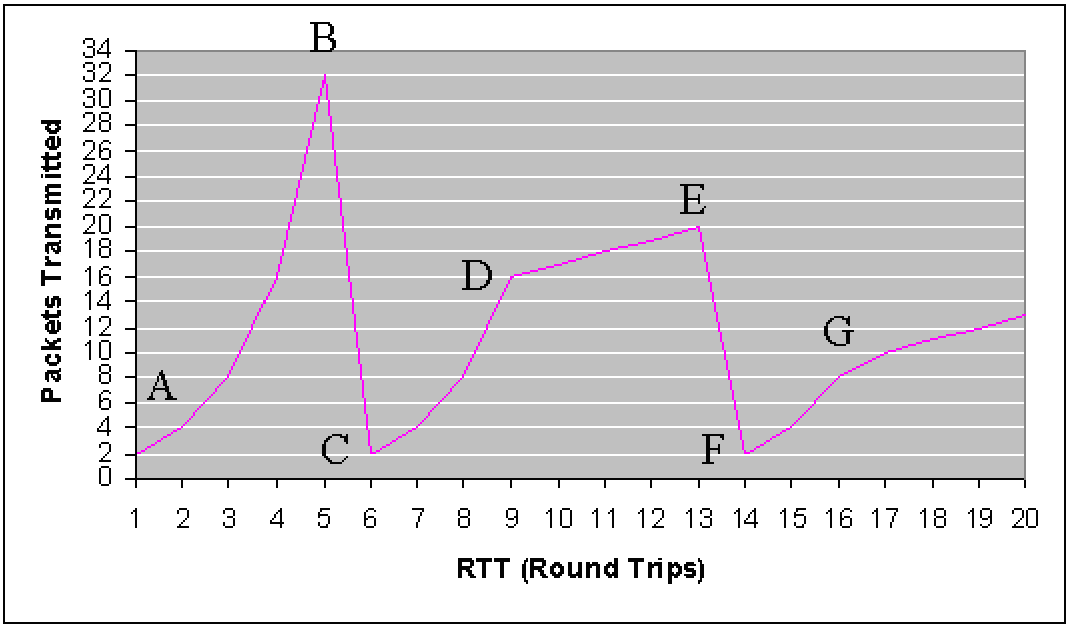

Figure 9: TCP Flow Control

A. After setting up a connection, TCP in the sender sets the initial congestion window. Generally this is between one and four packets, although the determination is actually somewhat more complex. (In this example we assume an initial window of two packets.) After each ACK is received, the congestion window is increased by one packet; depending on the ACK behavior of the receiver, this generally leads to a geometric increase in the window size until a packet is lost.
B. Packet loss occurs, detected by the sender either from the absence of an ACK by the expiration of the Timeout period (set by a complex algorithm with reference to RTT), or by receipt of a duplicate ACK, indicating that the receiver has detected a missing sequence number.[13]

C. The congestion window is reset to one, and the slow start threshold is set to one-half of what the congestion window was when packet loss occurred. (The slow start threshold is eventually reset by an algorithm if no further packet loss occurs.)
D. The connection enters Congestion Avoidance.
E. Packet loss occurs.
F. Congestion Window returns to the minimum, and the Slow Start Threshold is reset.
G. The connection enters Congestion Avoidance.

13.This latter, which results in a fast retransmit, is far more prevalent.

Appendix B: Collective Behavior on the Internet—an Example

Most research into the properties of networks has concentrated on the behavior of individual packets. However, by looking at the network as a *physical system* (rather than an abstract system of packet and queues) one can enable capabilities that are difficult or even impossible from a simply packet-oriented view—such as our chaos-control technology. Thus our name: Network Physics.

There is a great deal of research going on into the physical behavior of complex networks and the collective behaviors that arise from it, which is revealing phenomena with startling similarities to widely divergent areas of physics. For instance, in a recent paper, researchers showed that a sufficiently complex network acts like a Bose-Einstein gas, a collection of identical quantum particles that can occupy the same quantum state, such liquid helium at low temperatures.[14]

One such collective behavior that Network Physics has investigated is interference. Just as in optical interference (or quantum physics in general) a system may be observed as particles interacting or waves interfering, so with the Internet, or any packet-switched network. The correlation enforced on traffic by router queues creates strong interference effects between two congested routers along a data flow path.

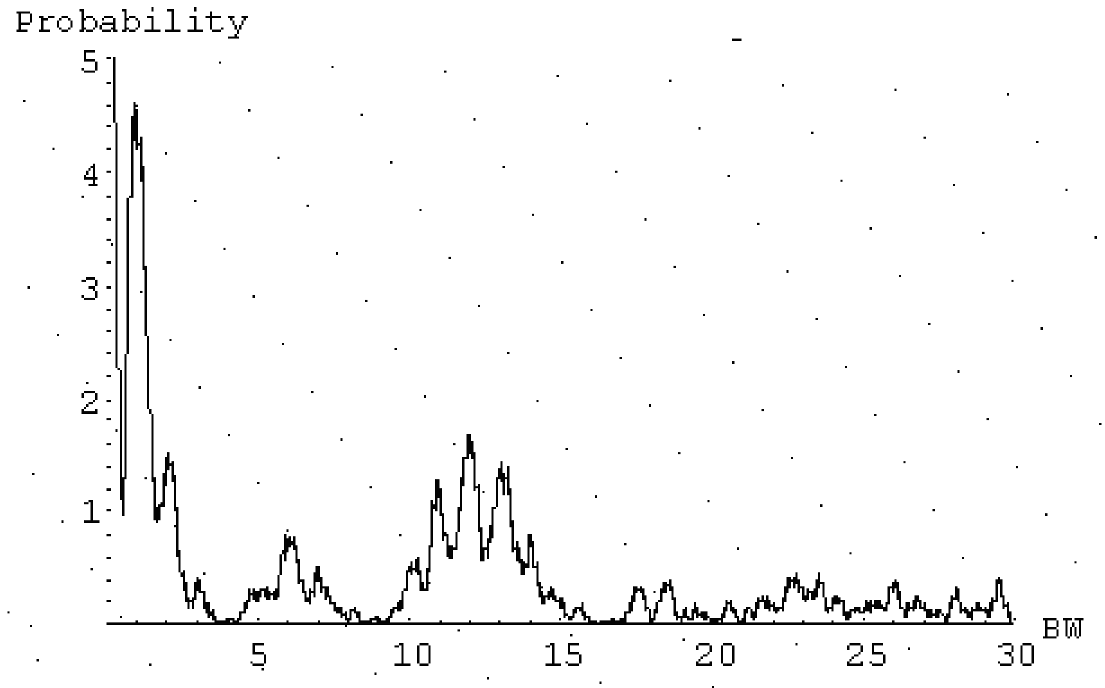

Figure 10: Interference between Routers

This figure shows "queuing activity" on the y axis and bandwidth in Mbps on the x axis. It reveals congestion at two routers, one associated with an effective bandwidth

14."The Physics of the Web", Physics World, July 2001, http://physicsweb.org/article/world/14/7/9 of approximately 1.5 Mbps, and one at approximately 12 Mbps. One can see the interference fringes (caused by "beats" between the resonances from the router at 1.5 Mbps and the one at 12 Mbps) at approximately 10.5 and 13.5 Mbps.

These resonances come from the interaction between a router and its output link. For instance, consider a router feeding 1500-byte packets into a T1 line (1.5 Mbps). Such a line can accept one such packet every 8 ms, a frequency of 125 Hz. Other size packets produce different frequences or resonances. Now imagine another router downstream from this one, also producing resonances by a similar process. Since the routers are coupled via the TCP flows that pass through both of them, these resonances can “beat against” each other, creating an interference effect like that revealed by the graph above.