

The Packet Switching Brain

Daniel Graham and Daniel Rockmore

Abstract

■ The computer metaphor has served brain science well as a tool for comprehending neural systems. Nevertheless, we propose here that this metaphor be replaced or supplemented by a new metaphor, the “Internet metaphor,” to reflect dramatic new network theoretic understandings of brain structure and function. We offer a “weak” form and a “strong” form of this metaphor:

The former suggests that structures and processes unique to Internet-like architectures (e.g., domains and protocols) can profitably guide our thinking about brains, whereas the latter suggests that one particular feature of the Internet—packet switching—may be instantiated in the structure of certain brain networks, particularly mammalian neocortex. ■

INTRODUCTION

Over the centuries, theories of how the brain works often reflect the zeitgeist of the era, articulated in terms of whatever is the current high technology. In Descartes’ time, the great advancements in plumbing, culminating with the water gardens at Versailles, inspired Descartes to imagine the nervous system as an intricate waterworks in which the plumbing of the brain was controlled by a master valve (the pineal gland) and his interpretations of brain function were guided by this metaphor. Later, Leibniz saw the brain as behaving like a mill. At the turn of the 20th century, telephones came into wide use and with them arose switchboard-inspired theories of the brain like those of the Connectionists.

The most recent technological paradigm to shape the language and approach of brain science is “the computer.” This goes back at least to Von Neumann, influenced by work of people like McCulloch and Pitts. In this context, processing in the brain is often cast as some kind of computation or execution of a program—to the point of mapping memory circuits in computers to memory circuits in the brain—even though few in the field believe there is a direct correspondence between these two architectures. Despite the differences and less than perfect correspondence, there are aspects of the analogy that can be useful: It is generally agreed that the brain, like a computer, translates input data (sensory data in the brain, and user input in the computer) into an internal language (neuronal spikes for the brain, and bits for the computer), to be processed by one or more function-specific systems (brain: navigation, eye movement, posture, locomotion; computer: word processor, image manipulation program, email client), that are finally expressed as some form of output (brain: speech, movement, memory trace; computer: monitor dis-

play, print out, memory allocation, Internet communication). Although it is important to be careful that such metaphors do not promote “elaborate fictions” (as noted by O’Reilly, 2006), there are, in fact, useful insights and interesting hypotheses generated by such analogies. Perhaps the most famous of these is the Marr (1980) model in which the brain is the architecture on which the algorithms of perception are executed.

Limitations of the Computer Metaphor

Many investigators would agree that the brain carries out computations, thus there is no need to forget the computer metaphor entirely. Such is true of many of the earlier metaphors (e.g., the plumbing analogy of Descartes is reborn in our current notion of ion pumps). However, new understandings have led us to believe that the computer metaphor, although perhaps still useful for certain aspects of brain function, is no longer a useful one, or at worst, a misguided one, for the next generation of brain science. The computational line of thinking centers on stimulus–response functions of neurons and brain regions, and it is assumed that what is important is determining the calculation carried out in each component. But if cortical neurons are performing computations, we have so far failed to understand most of the basic rules of the computations they perform. For example, outside of primary visual cortex (V1), a random sampling of cortical neural responses to natural stimuli in the visual stream will rarely be predictable. Even in V1, only about 15% of the variance in neural responses to natural stimuli is accounted for by the best current models (Olshausen & Field, 2005).¹ This is true because of fundamental factors, including response nonlinearities as well as for procedural reasons (neurons can only be recorded over a limited time period, and small neurons are difficult to record from). However, this failure may also be due to the fact that the computer metaphor

has misdirected our approach to deciphering the neural code. With respect to computers, the current approach is akin to measuring current at selected sites on a computer motherboard and attempting to deduce the function of the system as a whole by means of “decoding” these signals.

The focus on the individual computations performed by neural components has perhaps come at the expense of deep consideration of the nature of the transmission of the information across the physical neuronal network. Again, this is perhaps an artifact of the origins of the Von Neumann style mainframe-based computer metaphor. We have, since then, come to an understanding of the power of networked computation, and within that paradigm, issues of communication and information transmission are paramount. It is thus, in taking into account exactly these aspects of neural information processing that we argue in this article, time to shift the metaphor once again, to the Internet. An Internet metaphor offers a fundamentally new approach to the basic problems of neuroscience and offers promising new directions for research, although we note that it is subject to many of the technical hurdles of earlier approaches. A key insight provided by the Internet metaphor lacking in the computer metaphor is a consideration of communication across networks, and more specifically, of routing.

Shifting to a New Metaphor

The year 2009 marks the 40th anniversary of the start of the Internet and this period has also seen the rise of network science (see Börner, Sanyal, & Vespignani, 2007, for a review). With respect to the study of the brain, there are obvious reasons for its relevance. An Internet metaphor not only reflects our technological zeitgeist but also, more importantly, highlights network theoretic frameworks guiding a good deal of new neuroscience research, most notably by Sporns, Tononi, and Kotter (2005), Sporns, Chialvo, Kaiser, and Hilgetag (2004), and Sporns and Kotter (2004). Advances in diffusion tensor imaging have helped reveal aspects of the physical network structure of neural architecture, so that it is ripe for quantifiable comparison with other communication networks, both in structure and, as we propose here, in function. For the former, Sporns et al. have calculated descriptive statistics of network structure of a variety of vertebrate and invertebrate brains. Their findings appear to support the notion (previously put forth by Cherniak and others) that brains are highly efficient in their network structure. A hallmark of the network’s efficiency is its “small-world” structure (Sporns & Zwi, 2004), which allows any node (neuron) to communicate with any other node over only a few “hops,” or synapses.²

The finding that brain networks are structured in this way has deep connections to the development of the Internet, and it beautifully illustrates the interplay between technology and brain science. At a time when computer access and communication systems were highly central-

ized, Paul Baran proposed that distributed network architectures (which include small world networks) would allow efficient transmission of information from any node to any other node, requiring only a few hops. Compared to centralized and de-centralized networks, distributed networks could also better survive the destruction of nodes (Baran studied nuclear missile command networks). The more distributed the network, the better the robustness. Although this work was sensitive, it was published in the open press in 1964, and the dissemination of this discovery to scientists in other fields was key to the development of the modern Internet (Gillies & Cailliau, 2000). We therefore believe that this new unifying metaphor will spur further insights into the fundamental nature of information transmission in the brain, and that it gives us a new lens through which to view the findings of Sporns and others in this area. We demonstrate that this shift in metaphor not only offers new ways of synthesizing many disparate fields of neuroscience but also provides a new viewpoint from which to attack the many fundamental outstanding questions in neuroscience. In particular, it may aid in the deciphering of the neural code.

Reminiscent of the distinction made between “strong AI” and “weak AI” (see Searle, 1980), we describe a “weak” and a “strong” version of the Internet metaphor. The weak form stresses the new modes of model building gained in reference to Internet-like systems. The strong form of the metaphor offers potentially testable hypotheses regarding brain organization, evolution, and development. In particular, we propose that the network architecture of mammalian cortex instantiates a key feature of Internet-like systems, namely, packet switching. Although we restrict our “packet switching brain” hypothesis to mammalian neocortex, the variety of submetaphors described in the weak version of the Internet metaphor—which include the notion of “domains” and “protocols”—is more widely applicable to a variety of brain systems, although the mapping is less precise. But regardless of the strength of the mapping, a shift to the Internet metaphor privileges communication over computation as a prime goal of brains.

Given the emergent nature of Internet-like complex systems and the inherently nonlinear nature of their dynamics, it is natural to subsume theories of dynamical cognition (Spivey, 2008; see Friedenber, 2009, for an overview) within the Internet metaphor, and we note that researchers are already using the idea of routing to understand dynamics in embodied cognition models (Zhang & Ballard, 2001; Ballard, Hayhoe, Pook, & Rao, 1997). Likewise, information theoretic notions of neural coding efficiency also fit well within the same metaphor. Information theory, being the quantitative study of communication channels, has exerted enormous influence over the past half century of brain science (Field, 1994; Barlow, 1961; Attnave, 1954; see Reinagel, 2000, for a review). The advantage of collecting these diverse ideas under the Internet metaphor is to give theoreticians and experimentalists a common framework on which to project their models.

THE WEAK FORM OF THE INTERNET METAPHOR

We stress that we are interested here in applying only the analogy of Internet-like network routing to the brain. That is, we focus exclusively on communication protocol in Internet-like networks. However, useful analogies at other levels of the Internet are also possible: For instance, one could view cortical networks as possessing “content addressing” which, by linking related content, could function in a similar manner to the Google search algorithm (see, e.g., Griffiths, Steyvers, & Firl, 2007). We encourage further work into such extensions of the Internet metaphor.

On the Internet, the various communication protocols such as http, ftp, and email ultimately devolve to a form of information transfer called packet switching that define it as a “packet-switched network” (PSN). Unlike a traditional telephone network, where whole messages are sent from one node to the next in their entirety, all along the same “wire,” PSNs chop messages into small pieces. Each is addressed to the destination, allowing every “packet” to take the individually most efficient route to the destination, at which the packets are all recombined to give the sent message. It should be noted that the difference between PSNs and other networks is fundamentally one of communication protocol (i.e., the method of sending messages), not of connectivity (see Box 1 and Figure 1).

The basic unit of Internet-like networks is the node, which could be any device that transmits messages on the network (e.g., modem, router, or switch). Such nodes have an appealing correspondence to neurons and, therefore, serve as a better analogy than do transistors. However, we note that a node in a brain network may correspond to a collection of neurons, or a “cell assembly.”

The Internet metaphor subsumes a new and flexible collection of submetaphors, many of which may be found to align with brain function.

- (1) One advantage of the Internet metaphor compared to the computer metaphor is that it has a hierarchy of function built in. Rather than viewing cortical brain systems, such as language production, navigation, or object recognition, as stand-alone applications, more rightly they could be seen as domains (i.e., .com, .edu, etc.). In other words, cortical systems could be viewed as members of broad domains that share general properties, rather than as task-specific, software-like applications. These domains could encompass subdomains, in the same way that the .edu top-level domain subsumes the domains of each US university.
- (2) The array of neural coding strategies employed by the brain (gap junctions, synaptic transmission, neurotransmitters) can be seen as different layers of protocol (TCP/IP, ftp, etc.). The standardized, hierarchical protocol “stack,” which allows different operating systems, hardware, and applications to communicate, is one of the prime technological advances that allowed the de-

Kleinrock (1976) defines three general classes of network structure (see also Figure 1).

Circuit switching: This kind of network is typified by traditional telephone switchboards. Herein an initial “call” seizes all links on a complete path of connections between sender and receiver. When the “call” is “answered,” the message is sent along this path. The path remains dedicated to this call, whether or not data are being sent, until the sender frees the connections along the circuit. At this point, other senders may lease unused channels, and the process begins again. This structure is most efficient for transmitting long, dense, continuous messages in real time.

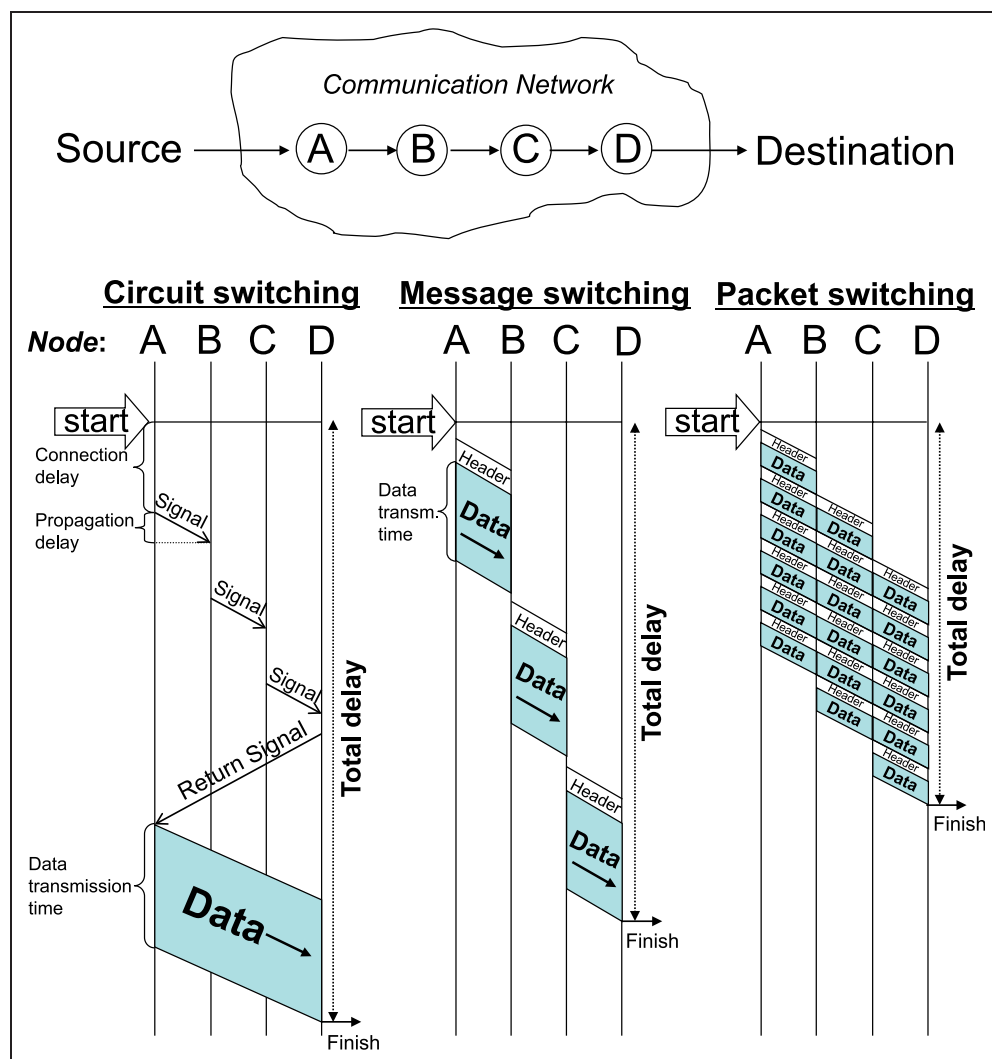
Message switching: Under this network structure, the entire message is sent from one node to the next along a circuit-switched path. That is, the entire message travels from node to node only when connections between the predefined set of nodes are free. The major difference of this scheme compared to circuit switching is that information is passed in a store-and-forward manner at each node, rather than in a continuous fashion. If a given connection in this path is in use, the message is stored at that node until the connection is no longer busy. Although this network structure allows more efficient sharing of channels, the wait-times experienced by the message at busy nodes make this structure ill-suited to fast communication. And unlike packet switching (see below), routing is static, that is, it is determined prior to sending the message.

Packet switching: Like message switching, packet switching does not appropriate the complete path to the destination. But rather than sending the entire message from node to node, the message is broken up into chunks or “packets,” whose size is limited. Each packet contains a “header” that contains the network address of the receiver, as well as information about what part of the message that packet contains. At each node, there is the possibility of dynamic routing, whereby the router at that node chooses the most efficient path for a given packet at the moment it arrives. By sending each packet on the most efficient path to the destination (via unused channels), this structure can greatly increase the speed of transmission compared to message-switched networks (Davies, 1968; Roberts, 1967; Baran, 1964). Indeed, PSNs can now function in real time in a way indistinguishable from circuit-switched networks, as evidenced by the rise of voice-over-IP and streaming video on the Internet.

Box 1. Network structures.

- (3) As Internet technology advances, new solutions that increase efficiency may also prove to be analogous to brain network properties. For example, it is now understood that high-bandwidth applications, such as real-time video, are transmitted most efficiently over peer-to-peer networks (e.g., CNN’s Pipeline service). Such networks store vast amounts of data in a highly distributed fashion, but thanks to sophisticated addressing and network control, entire files (even live video) can be accessed without disruption, and without large, dedicated memory allocations. Given that the dynamics of brain networks are poorly understood, the Internet metaphor could inform future approaches to neural coding by expanding the range of possible

Figure 1. This figure shows three major types of switching architecture. At the top of the figure is a schematic diagram of a network: A message (data) is sent across the network (consisting of Nodes A, B, C, and D) and on to the destination. In a circuit-switched system (below left), much of the time spent sending the message is used in the initial “call” that seizes the entire path over which the message will travel. The data are then sent all at once along this path without additional information. A message-switched system (below middle) eliminates the “call,” thereby reducing the total time of transfer, but this system suffers from significant delays at each node, where the entire message is stored until the next link is free. A “header” takes the place of the call: It contains the network address of the sender and destination, as well as a prespecified path, and it accompanies the entire data stream across the network. A packet-switched system (below right) sends a continuous stream of data along channels that are shared with other signals, thereby making the most efficient use of network bandwidth. Each packet’s header contains the sender and destination address, but routing is done dynamically and independently at each node depending on current network traffic (see Box 1 for further details on these switching architectures). Redrawn from Kleinrock (1976).



models, including highly distributed peer-to-peer frameworks. Indeed, the distributed nature of cortical coding is now widely accepted (Haxby et al., 2001; Field, 1994; Felleman & Van Essen, 1991).

Applying the Internet Metaphor

PSNs are a relatively recent network architecture which, by necessity, are much more flexible than simple message switching networks (somewhat akin to postal systems) or circuit-switched networks (e.g., telephone networks). Although it is appealing to expect that the brain compromises between the extremes of circuit and packet switching by employing message switching, we believe this is

unlikely due to the likelihood of delay. Nevertheless, such a tradeoff is possible.

However, many neural networks do exhibit properties of circuit-switched networks: The retina, for example, appears to use “leased lines” that transmit continuous streams of data to other parts of the brain. For example, retinal data about the visual scene are sent continuously along a dedicated path to “receivers” in the thalamus and the tectum. Although individual retinal cells often encode multiple dimensions of the scene simultaneously, there is considerable specificity in the division of labor, such that some cells transmit color information, whereas others encode motion, mean luminance, and so forth. In other words, retinal signaling does not appear to exhibit the dynamic routing

typical of PSNs (but see Gollisch & Meister, 2010). However, there is suggestive evidence that a circuit-switched network is a poor model of cortical network architecture. The Internet metaphor and, in particular, the notion of packet switching permit new explanations for a number of current questions for which a circuit switching model is inadequate.

- The rapid functional reallocation observed in many areas (e.g., tactile processing in V1 following blindness) suggests that a circuit-switched network structure is unlikely. Indeed, functional reorganization of somatosensation has been observed within hours of chemical blocking of peripheral nerve signals, suggesting that new routes for neural signals arise very rapidly (Weiss, Miltner, Liepert, Meissner, & Taub, 2004). If every connection is fully dedicated to sending only one kind of message, and only to its nearest neighbors on the network, how can the system as a whole quickly reorganize to process a different sensory stream entirely? Dynamic routing in PSNs could, in theory, allow this reorganization.
- Beyond V1 and V2 in the visual stream, the fan out of neural information to the vast, Web-like network of higher brain areas is so rapid and directed to so many disparate locations that some addressing system appears necessary. PSN structure could provide a solution to this problem. It is possible that packet reformation (i.e., the successful assembly of packets at the destination) is signaled by synchronized activity. This view aligns with proposals by Gray, König, Engel, and Singer (1989). However, it should be noted that PSNs succeed, in large part, because they are asynchronous methods of communication, and therefore, other means of signaling reformation may be at work.³
- Consider that a prefrontal cortex neuron involved in decision-making, which likely receives spike trains from many areas of the brain dealing with signals from multiple modalities, must have a mechanism for knowing where the input signals arise, and what they “mean.” A similar problem is faced at the other end of the network, by motor outputs, which likewise receive inputs from a great variety of areas. The notion of packet reformation provides a novel way of conceptualizing the signaling taking place over the entire extent of the network.
- Bursting activity, the common, but little understood, firing pattern characteristic of many cortical neurons, may be related to the fact that these bursts carry information-dense “packets” of information, which are distributed in a temporally sparse fashion. Indeed, sparse, bursty communications are precisely the type of signal for which PSNs are most efficient (Kleinrock, 1976). Spike trains are theoretically capable of transmitting a great deal of information beyond what is possible for rate codes (see Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997), and timing codes have been argued to be necessary and/or advantageous for many brain functions (Van Rullen, Guyonneau, & Thorpe, 2005) (see below for further discussion of neural “data packets”).

- Feedback is another widespread phenomenon in brains whose function few agree upon. In the visual system, it is estimated that over 90% of inputs to the thalamus (lateral geniculate nucleus, LGN) arise from higher areas such as V1 (primary visual cortex). Because the LGN is also the main terminus for axons of the optic nerve, feedback clearly plays an important role even in the earliest stages of processing. Traditionally, feedback has been thought to serve to “adjust the weights” of thalamic signals to cortex, although much remains unknown about what function this might serve. Alternatively, feedback could be seen as a return message from higher areas, one similar to the “acks” used in packet switching networks to acknowledge receipt of information. An example will help illustrate the idea. After one sends an email, if the machine sending the message (which has been chopped into packets) repeatedly does not receive a timely “ack” (a small return message) in response confirming receipt of the packets at their destination, the sending machine will try again. If it still fails to receive timely acks, it can look for a different route for unsuccessfully delivered packets. Likewise, thalamo-cortical feedback could be seen as a feature of a communication network: Data from LGN are sent to visual cortical areas for processing of spatial information, color, motion, and so forth. Feedback signals could then function as a way to acknowledge successful receipt or processing of sense data (e.g., object segmentation). One could speculate that thalamic channels that do not receive expected feedback could attempt to send their “message” to cortex again, or find another route. This change in viewpoint does not necessarily invalidate the computational or “neural network” picture of feedback. Rather, it suggests a new way of framing questions about it.

Advantages of PSNs

The idea of a packet of neural information has echoes in physics, where the notion of a wave packet connotes the fact that a photon possesses both particle- and wave-like properties. Likewise, a neural packet could be thought to carry concurrent forms of information (e.g., message content and “address”) that travel together across the network. As noted above, a neural “data packet” could correspond to a spike train or burst of spikes, with spike rate carrying message “content” and spike timing (e.g., first-spike time differences, or relative timing) carrying addresses. Additionally, cell connectivity could implicitly signal what portion of the “message” each packet contains: This would obviate the need to explicitly encode such information in the packet itself, as is required on Internet-like systems. It could be the case that a neural “data packet” is composed of an ensemble of signals from more than one neuron, and more complex forms of encoding are also possible.

The main advantage of a PSN is its ability to efficiently reroute network traffic around faulty nodes, something that is not possible on message-switched or circuit-switched

networks. On PSNs, dynamic routing is accomplished by sending each packet constituting a “message” on a potentially unique path. This is done on computer networks by means of hierarchically structured arrays of routers that store lists of routes to a great number of hosts, usually in a corresponding level of the hierarchy. Although the brain has been described as a hierarchical system, it is not obvious in what way path independence is at play in the brain. However, here too there is something to be learned from the PSN metaphor. Real PSNs typically send each packet along the same path to the destination. That is, once a successful route is found for some packets, rerouting is relatively rare. This can be viewed as a form of learning or memory instantiated in the path taken by a given packet in a “message.” When messages get through to their destination, a stable path is established. When messages fail to make it, a new path is chosen. Stability is achieved in the path through this routing scheme, and this may correspond to the stability of a memory or a long-term planning strategy.

Dynamic routing is a central feature of one prominent model of early visual processing (see Wiskott, 2006, for an overview). In Olshausen, Anderson, and Van Essen’s (1993) model, dynamic routing circuits normalize the view of an object to a canonical size and location, obviating the need for a complete tiling of retinal space with identical sets of feature detectors. Perhaps, in light of this model, cortical cells can be seen to act as routers, “reading” the address of the sender, the address of the receiver, and the message itself. In this view, all the advantageous properties of PSNs would be available to the cortical network as early as V2 or V4 (following Olshausen’s model). Moreover, plausible mechanistic models of neural routing have recently been proposed (Vogels & Abbott, 2009; Möller, Lücke, Zhu, Faustmann, & von der Malsburg, 2007). Given the brain’s small-world architecture, addressing may only be necessary over a handful of “hops.”

A consideration of “noise” can provide evidence in favor of the Internet metaphor as well. Noise (i.e., corrupted messages) on PSNs is minimal, as packets are of such small size. Moreover, components have multiple ways of redressing faulty connections: Corrupted packets can be detected midstream and dropped, and lost packets will be sent again by means of “acks.”

THE STRONG FORM OF THE INTERNET METAPHOR

The strong form of the Internet metaphor is intended to demonstrate an instance where the computer metaphor offers little in the way of hypothesis testing, but where the Internet metaphor provides a natural array of basic structures for comparison. Determining which basic switching structure (see Box 1 and Figure 1) is employed in cortex could greatly advance our understanding of the neural code. Here, we propose a way to answer this question by appealing to brain evolution and development, specifically of neocortex (or more appropriately, isocortex). In a general sense, we

also suggest that switching is a fundamental aspect of brains, one that could shape brain evolution and development.

It is now understood that mammalian cortex obeys general scaling properties, which play some role in the cognitive capabilities of a given species. Finlay and Darlington (1995) showed that small changes in the time course of brain development lead to predictable differences in the size of brain components. This notion is summarized by the phrase “late equals large,” that is, the later a developing brain begins forming out of a pool of precursor cells, the larger it will become.

Here we argue that increases in neocortical volume are constrained by a fundamental property of PSNs—one not shared by other types of networks—and that this fact can provide important insights into brain evolution. Moreover, we argue that other types of networks would be in danger of reaching sharp limits on scaling. In this section, we suggest that data from comparative neurology can be used to provide evidence regarding what kind of switching system is in use in mammalian cortex.

Constraints on Switching Architecture

A packet switching system, although providing the advantages described above, also imposes constraints on the number of cells that can communicate simultaneously. That is, because the benefits of fault tolerance, speed, noise minimization, flexibility, multi-band architectures, and so forth accrue only up to some limit, there will be a cost—and perhaps an evolutionary one—to neuronal networks that exceed this limit. Although metabolism and evolutionary history certainly play a role in constraining brain enlargement over evolution, the organization of cortical cell connectivity could be influenced by structural constraints on overall network efficiency. This idea suggests a potential way of determining whether mammalian neocortex employs packet switching architectures.

Because the number of cortical neurons scales with cortical volume, larger and larger brains are subject to basic limits in their ability to pass messages from one neuron to another. These limits are imposed by switching architecture. If each neuron passes messages of a given type over a dedicated connection—as in a circuit-switched network—this limit is reached abruptly. Adding an additional pair of neurons to a system that has reached capacity will result in that pair being unable to communicate. Up until that point, however, adding new pairs is cost free: Each pair secures a dedicated connection, and there is no loss in performance from adding another simultaneously communicating pair.

PSNs are different: As each new pair is added to a network under load, some performance is lost (performance can be seen as the speed with which a given pair on the network is able to pass a message). It is this incremental cost of adding new pairs of communicating nodes (i.e., new cortical neurons) that could limit neocortical volume over the course of brain evolution.

Testing the Hypothesis: Sparseness and Scaling

As such, there are ways to test the hypothesis that cortex exhibits properties of PSNs. For example, the metric of sparseness is now well established as a powerful tool for understanding neural codes, one that gauges how many neurons are active at a given time for a particular class of stimulus.⁴ Increased sparseness corresponds to a network where fewer nodes are active at the same time. Although a number of statistics have been developed to gauge sparseness of neural response distributions, (kurtosis, activity fraction, simple threshold measures; see Willmore & Tolhurst, 2001), all appear to describe the same property of a system (see Graham & Field, 2006, for an overview).

With regard to brain evolution, we hypothesize that sparseness of neural populations, in response to natural stimuli, will scale in proportion to the number of neurons in a given species' brain. For example, for a given brain size, the fundamental architecture of mammalian neocortex is such that a fixed fraction of neurons is able to communicate simultaneously, a constraint imposed by packet switching. It is also possible that sparseness scales in proportion to the number of synapses: In bigger brained animals, there are proportionally more synapses, but fewer neurons (see Changizi, 2006, for an overview). Indeed, the relationship between relative neuron number, synapse number, and sparseness is little understood, and could be subject to emergent dynamics due to the complex, adaptive nature of brains. The goal here is to suggest two basic alternatives under which models can be built to test for such relationships. Is it the case that more neurons (or synapses) mean less sparseness, or more sparseness? Further, under what conditions would sparseness be a limiting factor in brain evolution? This would be the case if cortex was packet switched: The number of simultaneously communicating neurons could grow only as a fixed function of neuron and/or synapse number.

The Packet Switching Brain

If such scaling is found, and assuming brains are equally efficient (i.e., highly “fit”) at performing core functions, we argue that this would constitute evidence in favor of packet switching, and against circuit switching. It is implausible that no constraints due to simultaneous activity exist—such as would be the case in a purely circuit-switched brain—given the danger of an abrupt “overload,” where one pair of neurons has too many attempts to communicate, but cannot.⁵ We know of no large-scale tests of sparseness across species, or even in multiple brain regions of the same species. Such information—although, no doubt, more difficult to collect than anatomical data—could open new doors in the study of neural coding and brain evolution, given the insights provided by the strong form of the Internet metaphor.

If sparseness was simply constant for all mammalian brains, this result would require prompt explanation. A

further possibility is that sparseness has no simple relationship to brain size. Lacking scaling, it is possible that large, nonlinear differences in sparseness across species could reflect fundamentally different switching architectures—this conclusion would appear to support Holloway's arguments in favor of fundamental neural reorganization as a substrate for increased cognitive ability (see, e.g., Holloway, 1996). However, we believe packet switching is employed in all mammalian brains, and possibly other brains, and therefore, the mammalian lineage will be subject to a uniform constraint on sparseness in proportion to brain size. Whatever the relationship, we argue that a major constraining factor in the evolution of bigger and bigger brains is switching, and that bigger brains (with proportionally more synapses) that are packet switched must contend with limits on simultaneous activity.

Of course, it is possible that human cortex is still very far from this limit. The argument presented here is that *some limit exists* where it is no longer advantageous to add more simultaneously communicating nodes to the network, because all communications over that network will be slowed, even if the PSN itself is efficiently structured and routed. That is, although energetic and processing constraints also play a role in limiting brain size (Hofman, 2000), it is possible that the switching architecture imposes an even greater constraint. Hofman (2000) argues that in larger brains, a given neuron's interconnectivity (how many cells it is connected to) is preserved because adding more neurons (and proportionally more white matter per new neuron) means “a large fraction of any brain size increase would be spent maintaining such a degree of wiring while the increasing axon length would reduce computational speed.” However, Hofman's argument fails to account for the fact that mammalian cortex shows sparse firing in low and high brain regions (see Baddeley et al., 1997, for evidence from the visual stream; see also Graham & Field, 2006, for a larger overview). That is, although large numbers of nodes (neurons) and connections (synapses) comprise the network, only a relatively small fraction of these are active at a given time.⁶ Indeed, sparseness itself is thought to be required for metabolic and other efficiency-related reasons (Waydo, Kraskov, Quiñ Quiroga, Fried, & Koch, 2006; Lennie, 2003; Attwell & Laughlin, 2001). Therefore, we argue that sparseness and not interconnectivity (relative to brain size) is what is held constant, and that this constraint is imposed by packet switching.⁷ A packet switching network could take advantage of such sparse structure (as described earlier), and as more neurons are added, the network would not be in danger of crossing a threshold of drastic performance loss, as would be the case if it were circuit switched.

Some have wondered (e.g., Hofman, 2000), what keeps mammalian brains from continuing to grow in size over many generations? Indeed, the rapid growth in brain size in human ancestors (an average of roughly 3 ml per millennium; Holloway, 1996) shows that constraining factors (e.g., cranial vault volume) can yet be surmounted (as shown

by cortical convolution) so long as the larger adult brain is useful to the differential reproductive success of the species. That is, structural constraints on brain evolution can be overcome assuming the payoff is large enough, and in humans at least, larger cortices appear to provide that payoff. However, it is not clear what factors contribute to expanded behavioral repertoires, and relative brain size does not predict cognitive capacities. For example, hummingbirds and whales show similar degrees of behavioral diversity but are orders of magnitude different in terms of brain volume (Finlay, Darlington, & Nicastro, 2001). New, specialized cortical regions may be important for special abilities in a given species, but larger lineages (such as mammals) seem less dependent on such novelties to propel greater ranges of function. Moreover, areal size differentials are small relative to those engendered by developmental scaling effects. In such debates, the role of switching has received little or no discussion. As we have argued, constraints on simultaneous activity play an important role in limiting evolutionary brain growth. In addition, if the slope of the scaling of sparseness starts to decrease (in studies of extant species, or in future lineages), we can surmise that the packet switching brain is under heavy load, and adding further neurons is less and less advantageous.

Switching and Development

Although it is possible that the adult brain could be far from the regime where packet switching networks constrain efficiency, the neonatal brain may not be. An alternative formulation of the strong form of the Internet metaphor can be constructed as the inverse of the hypothesis described above. It proposes that the dynamics of PSN efficiency play a role in development. As cortex develops, many connections between neurons are lost as a result of dendritic pruning and cell death. A gradual rise in global (and/or local) efficiency of a packet-switched cortical network from a beginning state of dense connection and relative inefficiency could thus act as a signal to slow and eventually stop the pruning process. This would only be the case if cortex exhibited packet switching, and we therefore encourage studies of network efficiency and connectivity through the course of early development to test this notion. For example, one could measure the sparseness of brain responses through early development. We predict that sparseness will increase as neural connections are cut, and therefore, fewer neurons are able to communicate simultaneously. That is, as the brain develops, fewer and fewer nodes can be simultaneously active (because fewer are directly connected due to pruning), but those that are able to respond simultaneously will operate as an efficient network, and serve as the substrate of adult function.

Conclusion

As grand descriptions, analogies to technology all ultimately fail to account for major aspects of brains. No single

mechanistic description has achieved more than a rudimentary description of perceptual or cognitive systems. However, each metaphor is useful to the extent that it incorporates into brain science novel discoveries in other fields, and insights that were generally unknown or unimaginable to earlier brain researchers. In this way, brain scientists have historically been able to adapt advances throughout the sciences and engineering into more and more successful models of the brain. It is the purpose of this article to suggest a new metaphor for the brain, which, like its predecessors, is incomplete on its own but is, nevertheless, useful when considered in concert with earlier metaphors. We suggest that the brain can be profitably analyzed as being analogous to the Internet. The weak form of this metaphor is useful because it privileges communication over computation in the analysis of brains, and because it offers helpful new analogies for brain function, such as domains and protocols. The strong form of the metaphor proposes that neocortex instantiates a defining characteristic of the Internet, namely, packet switching. We propose that the empirical scaling of response sparseness in neocortex across the mammalian lineage would constitute evidence that excludes other routing schemes (such as circuit-switching), and supports the notion of a packet-switched brain, as PSNs have different scaling properties compared to circuit-switched networks.

Acknowledgments

We thank Daniel Sheldon, Barbara Finlay, Michael Gazzaniga, Olaf Sporns, Cyrus McCandless, Giacomo Rizzolatti, and Leah Krubitzer for very helpful comments. This work was supported by an NSF Small Grant for Exploratory Research (DMS-0746667).

Reprint requests should be sent to Daniel Graham, Department of Mathematics, Dartmouth College, HB 6188, Hanover, NH 03755, or via e-mail: daniel.j.graham@dartmouth.edu.

Notes

1. Analysis of fMRI responses, another “decoding” tool, often suffers from additional reliability issues (Yarkoni, 2009; Vul, Harris, Winkielman, & Pashler, 2009), although careful experiment design and new multivoxel pattern analysis methods may offer some progress in this regard (see, e.g., Kay & Gallant, 2009).
2. In cortex, this distance has been suggested to be three or fewer synapses (Changizi, 2006; Wiskott & von der Malsburg, 1996).
3. Alternatively, one could see coherent activation as an indication of the path that a message takes across the network, as proposed by Fries (2005).
4. This is the definition of population sparseness; a related notion, lifetime sparseness, applies to the “burstiness” of an individual cell in response to a set of stimuli over long time periods. We believe both are worthy of greater study in terms of scaling behavior and routing architecture.
5. Separate costs may accrue due purely to network connectivity constraints in adding nodes to small-world networks (Amaral, Scala, Barthelemy, & Stanley, 2000). It is unclear, at present, how routing and network constraints together affect brain scaling. There is a need for further study of single-unit connectivity in neocortex and of axonal “projective fields” (see Sejnowski, 2006).

6. Selective attention and inhibition also play a role in limiting simultaneous activity, but we suggest that these mechanisms are subject to global constraints on network scaling as well.

7. We note that we are interested here in simultaneity, not synchrony: We do not assume that the degree of simultaneous firing is necessarily related to synchronized activity, only that the system as a whole is subject to constraints on simultaneous firing. However, synchrony could play a role in coding schemes of packet-switched brain networks.

REFERENCES

- Amaral, L. A. N., Scala, A., Barthelemy, M., & Stanley, H. E. (2000). Classes of small-world networks. *Proceedings of the National Academy of Sciences, U.S.A.*, *97*, 11149–11152.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.
- Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism*, *21*, 1133–1145.
- Baddeley, R., Abbott, L. F., Booth, M. C., Sengpiel, F., Freeman, T., Wakeman, E. A., et al. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, *264*, 1775–1783.
- Ballard, D., Hayhoe, M., Pook, P., & Rao, R. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, *20*, 723–767.
- Baran, P. (1964). On distributed communications. *RAND series reports*.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Börner, K., Sanyal, S., & Vespignani, A. (2007). Network science. In B. Cronin (Eds.), *Annual review of information science & technology* (Vol. 41, Chap. 12, pp. 537–607). Medford, NJ: Information Today, Inc./American Society for Information Science and Technology.
- Changizi, M. A. (2006). Scaling the brain and its connections. In J. H. Kaas (Ed.), *Evolution of nervous systems* (Vol. 3, pp. 167–180). Oxford: Elsevier.
- Davies, D. W. (1968). The principles of data communication networks for computers and remote peripherals. *Proceedings of the IFIP Congress 68*, Edinburgh, Scotland, pp. 709–714.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*, 1–47.
- Field, D. J. (1994). What is the goal of sensory coding? *Neural Computation*, *6*, 559–601.
- Finlay, B. L., & Darlington, R. B. (1995). Linked regularities in the development and evolution of mammalian brains. *Science*, *268*, 1578–1584.
- Finlay, B. L., Darlington, R. B., & Nicastro, N. (2001). Developmental structure in brain evolution. *Behavioral and Brain Sciences*, *24*, 263–308.
- Friedenberg, J. (2009). *Dynamical psychology*. Litchfield, AZ: ISCE Publishing.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*, 474–480.
- Gillies, J., & Cailliau, R. (2000). *How the Web was born*. Oxford: Oxford University Press.
- Gollisch, T., & Meister, M. (2010). Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron*, *65*, 150–164.
- Graham, D. J., & Field, D. J. (2006). Sparse coding in the neocortex. In J. H. Kaas & L. A. Krubitzer (Eds.), *Evolution of nervous systems* (Vol. 3, pp. 181–187). The Netherlands: Elsevier.
- Gray, C. M., König, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, *338*, 334–337.
- Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the mind. *Psychological Science*, *18*, 1069–1076.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*, 24–30.
- Hofman, M. A. (2000). Brain evolution in hominids: Are we at the end of the road? In D. Falk & K. Gibson (Eds.), *Evolutionary anatomy of the primate cerebral cortex* (pp. 113–127). Cambridge, UK: Cambridge University Press.
- Holloway, R. C. (1996). Evolution of the human brain. In A. Lock & C. Peters (Eds.), *Handbook of human symbolic evolution* (Chap. 4, pp. 74–116). New York: Oxford University Press.
- Kay, K. N., & Gallant, J. L. (2009). I can see what you see. *Nature Neuroscience*, *12*, 245–246.
- Kleinrock, L. (1976). *Queueing systems: Vol II. Computer applications*. New York: Wiley.
- Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*, 493–497.
- Marr, D. (1980). *Vision*. San Francisco, CA: Freeman.
- Möller, C., Lücke, J., Zhu, J., Faustmann, P. M., & von der Malsburg, C. (2007). Glial cells for information routing? *Cognitive Systems Research*, *8*, 28–35.
- Olshausen, B. A., Anderson, C. H., & Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience*, *13*, 4700–4719.
- Olshausen, B. A., & Field, D. J. (2005). How close are we to understanding V1? *Neural Computation*, *17*, 1665–1699.
- O'Reilly, R. C. (2006). Biologically based computational models of high-level cognition. *Science*, *314*, 91–94.
- Reinagel, P. (2000). Information theory in the brain. *Current Biology*, *10*, R542–R544.
- Rieke, F. M., Warland, D. K., de Ruyter van Steveninck, R., & Bialek, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- Roberts, L. G. (1967). Multiple computer networks and inter-computer communications. *Proceedings of the ACM Symposium on Operating Systems Principles*. New York: ACM.
- Searle, J. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, *3*, 417–457.
- Sejnowski, T. J. (2006). What are the projective fields of cortical neurons? In J. L. van Hemmen & T. J. Sejnowski (Eds.), *23 Problems in systems neuroscience* (pp. 394–405). New York: Oxford University Press.
- Spivey, M. (2008). *The continuity of mind*. New York: Oxford University Press.
- Sporns, O., Chialvo, D., Kaiser, M., & Hilgetag, C. C. (2004). Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, *8*, 418–425.
- Sporns, O., & Kotter, R. (2004). Motifs in brain networks. *PLoS Biology*, *2*, 1910–1918.
- Sporns, O., Tononi, G., & Kotter, R. (2005). The human connectome: A structural description of the human brain. *PLoS Computational Biology*, *1*, e42.
- Sporns, O., & Zwi, J. (2004). The small world of the cerebral cortex. *Neuroinformatics*, *2*, 145–162.
- Van Rullen, R., Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *Trends in Neurosciences*, *28*, 1–4.

- Vogels, T., & Abbott, L. F. (2009). Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nature Neuroscience*, *12*, 483–491.
- Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition (The paper formerly known as “Voodoo correlations in social neuroscience”). *Perspectives on Psychological Science*, *4*, 274–290.
- Waydo, S., Kraskov, A., Quiñ Quiroga, R., Fried, I., & Koch, C. (2006). Sparse representation in the human medial temporal lobe. *Journal of Neuroscience*, *26*, 10232–10234.
- Weiss, T., Miltner, W. H., Liepert, J., Meissner, W., & Taub, E. (2004). Rapid functional plasticity in the primary somatomotor cortex and perceptual changes after nerve block. *European Journal of Neuroscience*, *20*, 3413–3423.
- Willmore, B., & Tolhurst, D. J. (2001). Characterizing the sparseness of neural codes. *Network: Computation in Neural Systems*, *12*, 255–270.
- Wiskott, L. (2006). How does our visual system achieve shift and size invariance? In J. L. van Hemmen & T. J. Sejnowski (Eds.), *23 Problems in systems neuroscience* (pp. 322–342). New York: Oxford University Press.
- Wiskott, L., & von der Malsburg, C. (1996). Face recognition by dynamic link matching. In J. Sirosh, R. Miikkulainen, & Y. Choe (Eds.), *Lateral interactions in the cortex: Structure and function*. Austin, TX: The UTCS Neural Networks Research Group, electronic book. See <http://userweb.cs.utexas.edu/~nn/web-pubs/htmlbook96/>.
- Yarkoni, T. (2009). Big correlations in little studies: Inflated fMRI correlations reflect low statistical power. Commentary on Vul et al. (2009). *Perspectives on Psychological Science*, *4*, 294–298.
- Zhang, Z., & Ballard, D. H. (2001). Distributed synchrony. *Journal of Neurocomputing*, *44–46C*, 715–720.