Working Paper
For circulation at the 2017 Global Digital Futures Forum
May 5, 2017

COLUMBIA|SIPA
School of International and Public Affairs

Framing Conversation: What Would Internet Fragmentation Mean for the Digital Economy?

By William J. Drake

1. Introduction

The theme of this year's Forum is very timely, as the question of Internet fragmentation has been the focus of a good deal of discussion of late in both generalist and specialist policy circles. But before we can explore the potential impact of Internet fragmentation on the digital economy, global governance, and global trade, it would be useful to step back and consider what we mean by the term in the first place. Some references in popular media seem to suggest that the term connotes a singular phenomenon on which there is broad agreement so we can simply invoke it and move on from there.

In fact, Internet fragmentation remains a contested concept. A cursory review of its usage in various publications and public pronouncements suggests that people often speak of it when discussing a variety of problems and tensions that arise on the Internet that do not all originate from the same source. For example, some in the business community have used the term as a generalized reference to variations in national policies that add to the cost of doing business globally. While some such policies may indeed be related to fragmentation, many other simply reflect differences in national legal systems, policy traditions, and so on that may antedate and arguably do not fragment the Internet. Similarly, some people have described the increasing linguistic diversity of cyberspace as an example of fragmentation, when of course this is simply a matter of a diverse humanity getting on line.

Another tendency among at least some observers is to suggest that the Internet is in imminent danger of falling apart. Because there is so much variation in national policies and practices, it is said, the Internet is likely to "break up" into a series of disconnected islands. This seems to be an overly dramatized misreading of some troubling trends. In fact, no cataclysm is around the corner; the underlying infrastructure remains stable and secure in its foundations, and it is incorporating new capabilities that open up new horizons, from the Internet of Things and services to the spread of block chain technology

and beyond. But there are fragmentary pressures accumulating which, if left unattended, could reduce to varying degrees the Internet's enormous vitality and contributions to the world.

Conversely, while the examples just mentioned concern overly broad applications of the term, other observers tack in the opposite direction and say that "fragmentation" can only be properly used in reference to the Internet's underlying infrastructure rather than the creation of significant closed digital spaces. In one variant of this thinking, fragmentation would only happen if there was a massive defection from the unified Internet to entirely separate and non-interoperable systems running off different zone files. Since such a defection does not appear to be likely in the near future, voilà, there is no fragmentation, and people who argue to the contrary are needlessly hyperventilating, perhaps in hopes of looking prescient.

With these conditions in mind, in this memo I will briefly address three matters in the hope of helping to frame the conversation. First, I will advance working definitions of Internet fragmentation drawing on a white paper I wrote with colleagues for release at the World Economic Forum's (WEF) Annual Meeting in Davos in January 2016. [1] Second, I will highlight the variability and fluidity of Internet fragmentation in order to underscore that we are not talking about a simple binary condition that flicks on or off like a light switch. Third, I will conclude by raising a few concerns about the potential impacts of fragmentation on the evolving global digital economy.

2. Defining Internet Fragmentation

A useful starting point is to consider what we mean by an unfragmented Internet. What is the baseline from which fragmentation departs and against which it can be assessed? From a technical standpoint, the original shared vision guiding the Internet's development during the research and education era was that

[1] William J. Drake, Vinton G. Cerf, and Wolfgang Kleinwächter, *Internet Fragmentation: An Overview* (Geneva: The World Economic Forum, January 2016). A few bits of this memo derive from that earlier paper.

every willing endpoint on the Internet should be able to exchange data packets with any other endpoint that was willing to receive them. Universal "connectivity among the willing" was the guiding objective, and it could be achieved if autonomously controlled and even separately designed networks were internetworked and made interoperable via a shared protocol stack, TCP/IP, and related standards and protocols. Such interoperability needed to be to be seamlessly coherent on an end-to-end basis and consistent, so that users' actions would yield the same responses irrespective of location or the service providers involved.

These core features of universal, consistent interoperability and communicability between consenting end points were fundamental from a design standpoint. Every end point that wanted to send and receive bits with any other should be able to do so, so that the network of networks functioned as a free and open system. Actions or conditions that impaired this seamless functioning and blocked users from reaching each other could be said to constitute fragmentation.

Imagine, by way of analogy, an international telephone network on which people in country A could communicate with people in countries B, D, and F but not with people in countries C, E and G, while people in country B could communicate with people in countries A, C and E but not D, F and G, and so on across 196 countries. If humanity's ability to reach the full range of willing correspondents were this barrier-laden and segmented into go and no-go zones, would we characterize the global telephone network as open and unfragmented? Probably not. But on the Internet this sort of highly variable geometry of communicability is fairly standard and taken for granted, especially if one considers the infinite substantive variety of the bits that could be shared if allowed. We know that over 700 million users in China cannot access major platforms that are used by billions of people elsewhere; that billions of downloaders encounter messages like "the content you requested cannot be displayed;" that the transfer of certain classes of data out of certain countries is blocked or requires government permission; and so on, endlessly.

My contention, which like others is certainly contestable, is that the pervasive limitations on users' abilities to freely access, create, and dissemination information indicates an endemic condition of Internet fragmentation. The Internet is not a wide-open medium in which "anything goes," popular characterizations notwithstanding. It is certainly far more open than any global medium we have ever had

before, and the limitations on its openness are frequently the focus of efforts to bypass or reverse by various actors, but they are there. And, as Eli Noam has argued in a provocative essay, they were inevitable. [2] There was simply no chance that the conditions that obtained in the early years when the Internet was a vehicle for the non-commercial sharing of research and educational information among computer scientists in various organizational settings could survive the transition to the Internet becoming a global mass medium used for an endless variety of social, commercial and political information sharing and resource discovery. Inevitably, governments were going to work to embed the Internet in frameworks of public authority that involved a wide variety of prescriptions and proscriptions, and companies were going to work to monetize peoples' access to and use of different kinds of contents by erecting a wide variety of enclosures and requirements. At the same time, with millions of technical people around the world working to deploy new capabilities, increase security and various other objectives, conditions could develop that, often unintentionally, had the effect of reducing or at least complicating the seamless functioning and interoperability of the infrastructure.

Hence, from this standpoint, it makes little sense to pose questions like "will the Internet fragment?" The Internet has long been fragmented to varying degrees in varying ways. A better question might be, will "Internet fragmentation increase in a manner that becomes much more problematic for a much wider range of uses and users?" Such a formulation turns our attention to the direction of change, rather than whether change might commence.

While Internet fragmentation has a common root---limitations on the ability of every willing endpoint to exchange data packets with any other willing endpoint---it is not a singular phenomenon. Fragmentation varies in its sources and manifestations in ways that are worth assessing separately on their own terms. Hence, in the above-mentioned paper for the WEF, my co-authors and I advanced three different "working definitions," so-called because the paper was an initial exploration and mapping and we were cognizant that more precise formulations might be desirable after our colleagues in the field

[2] Eli M. Noam, "Towards a Federated Internet", *InterMEDIA* (41, 4, 2013), pp. 10 –13.

kicked us around a bit on points that needed rethinking. We began from the proposition that a single "narrow definition" focused only on conditions in the underlying infrastructure would not capture how people use and experience the technology in order to construct digital social formations and engage in information, communication and commercial transactions, or by extension the sorts of political and economic forces that may impede their abilities to do so. We therefore amended the standard four-layer characterization of the Internet based on the TCP/IP Protocol Stack by adding a fifth

Content and Transactions layer to capture the substantive information exchanged and the interactions and behaviors involved.

Figure 1: Internet Layers

5. Content and Transactions Layer
4. Application Layer
3. Transport Layer
2. Network/IP Layer
1. Physical/Link Layer

Beginning from this amended baseline, we advanced the following working definitions of fragmentation:

- *Technical fragmentation:* conditions in the underlying infrastructure that impede the ability of systems to fully interoperate and exchange data packets and of the Internet to function consistently at all end points. These generally pertain to layers 1-4 of the model above.
- *Governmental fragmentation:* Government policies and actions that constrain or prevent certain uses of the Internet to create, distribute, or access information resources. These generally are targeted at the 5th layer in our model, but they may involve actions taken at the lower technical layers as well.

- *Commercial fragmentation:* Business practices that constrain or prevent certain uses of the Internet to create, distribute or access information resources. These generally are targeted at the 5th layer in our model, but they may involve actions taken at the lower technical layers as well. [3]

As is evident, one of our concerns was to distinguish sources and locations of fragmentation based in part on the question of intentionality. Technical fragmentation of the underlying physical and logical infrastructure is a complex evolutionary process that has unfolded slowly but is gathering pockets of steam in the contemporary era. Some of it has been intentional and motivated by operational and other concerns, but more often it has been the unintended by-product of actions taken with other objective in mind. In contrast, governmental and commercial fragmentation usually have been due to the intentional efforts of these third parties to establish limitations on users' abilities to create, distribute or access information. As a general matter, one could argue that such limitations are much more problematic and difficult to remediate than technical problems, for which engineers often can devise "fixes." In contrast, governmental and commercial fragmentation can be difficult to engineer "work arounds" for with lasting effects, e.g. people confronting censorship may rely on virtual private networks to mask their locations, but then governments figure out ways to block and monitor these and another technique must be found, at least until that too is found out.

[3] Drake, Cerf and Kleinwächter, p. 14.

3. Variability and Fluidity

Just as fragmentation is not singular in its form or the domain of its effects, the extent of fragmentation within and across the three categories also is highly variable. One could imagine a number of dimensions on which such variation could be found, but here are just three that merit consideration.

Occurrence: The first and most fundamental consideration is whether a given form of fragmentation exists. This is not an entirely straightforward question; as is noted above, fragmentation as a systemic property is not a simple binary condition that is either present or not present, and a specific instance of fragmentation in some domain may involve gradations with different values along a continuum. In some cases those values can be precisely quantified (e.g. the number of websites or other information resources to which access is fully blocked), but in others the best we can do is to devise ordinal measures. Similarly, there can be variations in duration. Fragmentation may be a short-term phenomenon that is rectified fairly quickly, as with recovery from some an attack that blocks access to resources, or it can be sustained as a long-term condition. In time sensitive situations, even short-term fragmentation can be very damaging to users or transactions. In general though, we should be most concerned with sustained fragmentation that has ongoing consequences.

A final issue here is that fragmentation does not need to be currently present to be of concern. That is, in many of the instances that people cite when worrying about the matter, what is at stake is the emergence of tendencies and pressures that could give rise to something significant in the future. As in any policy arena, we need not wait for a problem to become full blown and wreaking havoc for awareness and action to be well advised.

Intentionality: Fragmentation, particularly in the technical arena, may be the unintended by-product of decisions and actions guided by unrelated objectives. People who deploy or fail to deploy a particular technology in addressing a localized operational challenge may not be setting out to fragment the Internet. Nevertheless, their actions, especially if replicated by others, could come to have cumulative effects. Divergences between individually rational choices and systemically suboptimal consequences are a standard feature of collective actions problems generally and the same logic can apply to the openness or fragmentation of the Internet.

Alternatively, fragmentation may be intentional. The character of these intentions obviously matters quite a bit. On the one hand, organizations, communities and individuals may seek to separate themselves somewhat from the open public Internet for entirely defensible reasons. Installing a firewall to limit access and communication to only authorized and consenting parties and to protect resources from unwanted interference is a benign act of self-separation. In our WEF discussions last year, some participants argued that self-separation, such as the construction of firewalls or the use of encryption on a network, could be thought of as "positive fragmentation." I tend to think of this as being a different sort of activity that may involve some protective segmentation but is not preventing willing end points from communicating, since one end point is choosing to mediate its boundaries.

Of more concern, and more properly a matter of fragmentation in my view, is when actors such as governments seek to shape, constrain or fully block the activity of others who have not consented to this. Imposing limitations on others is a malign act of forced separation. Both unintentional and intentional fragmentation can be problematic, but the best approach to remediation may vary accordingly.

Impact: Fragmentation may be deep, structural and configurative of large swaths of activity or even the Internet as a whole. Consider, for example, the implications if significant categories of data flows were to be widely blocked around the world, or if an alternative root system with its own name space were to be established with the backing of powerful governments or organizations. The scope of the processes, transactions and actors impacted by such breakage would be substantial. But fragmentation also can be more shallow, malleable and applicable to a narrowly bounded set of processes, transactions and actors. The impact could be significant for some people but go unnoticed by others.

As with the other dimensions just mentioned, it can be difficult to measure the intensity of fragmentation and say with certainty exactly where on the continuum a given instance lays. Even so, in considering examples, we should be mindful that fragmentations are not all created equal in terms of magnitude and import. Indeed, a number of the examples one could mention are relatively low-impact or low-intensity matters – bothersome and concerning enough to engineers and operators that attention to them is merited, but not so significant that they endanger the fundamental integrity, openness and utility of the Internet. In contrast, some other action are higher-impact and arguably in need of concerted responses.

Given the above, from a systemic standpoint fragmentation is something of a shape shifter. It is always with us, particularly at the fifth level of content and transactions, but its specific manifestations are highly fluid and variable in scope, depth and duration. What should be of most concern are intentional forms that are deep, structural and configurative of large swaths of activity or even the Internet as a whole.

4. Implications for the Global Digital Economy

Some forms of fragmentation of this character are of relevance to the opening session on the digital economy. For example, with regard to technical fragmentation, if governments engage in widespread blocking of new generic top-level domains, opportunities for additional economic growth and social empowerment would be foreclosed. A massive defection by a leading country or countries to another root system, while presently unlikely, undoubtedly would have a very pronounced negative impact on the global digital economy. In general though, technical fragmentation at present does not seem likely to take on the sort of character that would in any dramatic way spoil the party.

Commercial fragmentation probably raises greater risks. There is growing concern today as to whether divergent corporate preferences may result in inadequate technical standardization of the emerging Internet of things. The adoption and locking in of proprietary standards in key arenas like this could produce fragmenting effects, with important products and processes not working well across corporate boundaries and national borders. The current push in the United States to abandon network neutrality as an organizing principle, driven in particular by traditional network operators and government ideologues, could result in widespread discrimination against applications and entrepreneurs and produce a fragmentary, multi-speed environment. Overly expansive and rigid intellectual property rules could curtail entrepreneurial dynamism as well as free expression and human empowerment. And as we move ever further into a platform-dominated online economy that absorbs an increasing share of advertising dollars and economic activity, the ways in which terms of service are constructed, the possibilities for anti-competitive behavior, and the prevalence of “walled garden” strategies may alter the character of the digital economy in ways that attenuate existing inequalities. Arguably, this may be particularly a concern with respect to the participation of developing countries in the digital trade arena.

Finally, and most importantly, governmental fragmentation of a structural nature seems to be a particularly pressing concern. The widespread "securitization" of Internet policies and the growth of so-called "cyber-sovereignty" strategies is already producing trends toward more widespread censorship and digital protectionism. These measures can be very difficult to roll back, and can impose significant costs on global companies and national economies and citizens alike. The potential scope of the challenge is underscored by the current trend toward forced data localization policies and the erection of barriers to cross-border data flows, which are the subject of a follow-up study to the above-mentioned fragmentation paper that will be released later this year. [4] In the opening session we may wish to delve into these and related questions.

[4] William J. Drake, *Data Localization and Barriers to Cross-Border Data Flows: Toward a Multistakeholder Approach,* (Geneva: The World Economic Forum, September 2017).