# Bridging the Digital Divide: Success Depends on Content Provider and Application Developer Involvement

By ANDREW LAPPALAINEN AND CATHERINE ROSENBERG, *Fellow IEEE*
*Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada*

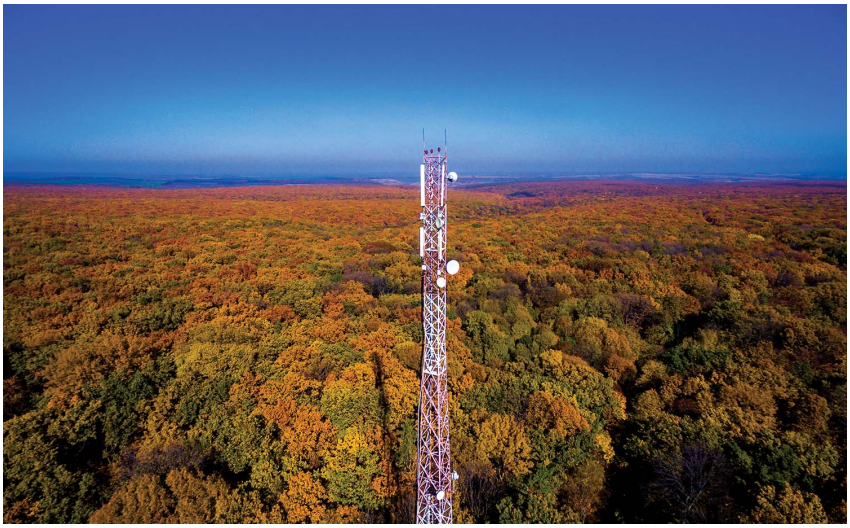

Global connectivity is at an all-time high, and more users than ever before are participating in the online ecosystem. Despite this exciting phenomenon, opportunities to access the online world are not shared equally among the global population. Socioeconomic, political, and geographic factors all play roles in determining the extent to which one can be an online participant [1]. This resultant digital divide creates a social disparity between those who have reliable access to online content and can take advantage of all that the Internet has to offer and those who do not and miss out on those opportunities [2].

For isolated, remote communities (RCs) such as those found throughout Canada's north, the digital divide continues to be an issue due to the prohibitive costs of building network infrastructure into these regions [1]. Despite its costs, satellite has prevailed in these locations because of its ubiquitous coverage, and for many RCs (see Fig. 1), satellite is the only means by which individuals can access the Internet today [3], albeit with many limitations. Typically, a satellite link is the bottleneck in an end-to-end communication due to its inherently long latency and relatively narrow bandwidth.[1] Furthermore, satellite is costly to deploy and operate. Thus, RC operators (RCOs)[2] are limited in their ability to provide a high quality of experience (QoE) to their customers, and the experience of these users lags far behind those in other areas of the country. In order to mitigate the QoE issues, it is critical for RCOs to be able to manage the signaling and traffic over the satellite link to make the best use of it as possible.

[1]Especially when one considers that it is shared among many users.

[2]In the context of this article, an RCO is the entity in charge of managing the network in one or more RCs (see Fig. 2). It is dependent on the satellite to backhaul service to these communities.

Digital Object Identifier 10.1109/JPROC.2020.3028611

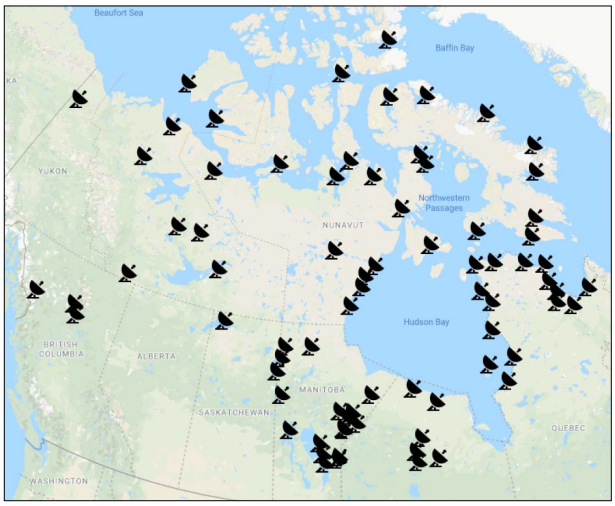


**Fig. 1.** *Communities in Canada where satellite is the only medium over which users can connect to the Internet. Adapted from [3]. Map data © 2020 Google, INEGI.*

In this article, we revisit the digital divide problem in the context of remote, satellite-dependent communities and challenge the earlier notion that network operators are the only technical stakeholders responsible for bridging the divide.

Historically, operators were expected to lead the effort in spite of there being no business case in this context. More recently, not-for-profit operators have taken over, but they cannot bridge the digital divide alone.

Indeed, we believe that the digital divide cannot be solved unless other stakeholders commit to helping solve this issue. In particular, we argue that content providers and application developers (CPADs)[3] must also be held accountable for narrowing the digital divide in RCs, both as a matter of social responsibility and because they hold increasing sway over the performance of network connections. Specifically, we propose that CPADs must commit to be more "remote-aware and friendly" and deliver content to remote users in ways that help RCOs effectively use the satellite link. CPADs have a strong influence over how content is delivered to users; however, certain trends that they have adopted harm the experience of remote users. As we will discuss in more detail in Section I, CPADs are using a greater diversity of protocols and applications, richer content, and more encrypted sessions, all of which reduce the ability of RCOs to manage the satellite link efficiently. Indeed, CPADs are increasingly relying on elaborate solutions implemented at the edge to manage the QoE of their customers, in effect making connections opaque via encryption and shifting control of connections away from operators to the endpoints. This has been to the advantage of most users who are seeing better QoE. However, as CPADs opacify the network, they must bear the responsibility that comes with it. CPADs cannot abandon the minority (e.g., the RCs) whose QoE suffers as a result of these trends.

What then does it mean for CPADs to be "remote-aware and friendly" and how will it change things if they take responsibility in this situation? To begin to answer this question, in Section I, we revisit the conventional two-pillared solution of RCOs that incorporate first proxies to manage signaling and traffic and second caching of popular content within the community. In light of the proliferation of complex applications and rich content, proxies and caches are more essential than ever, but RCOs are losing the ability to depend on them as a solution as connections become increasingly opaque. This is to the detriment of remote users. Therefore, the first step is for CPADs to recognize that this will negatively affect users and commit to supporting the continued use of proxies and caches by RCOs. Furthermore, we propose that the two-pillar solution be augmented with two additional pillars, third, filtering of unnecessary content, specifically advertisements, before delivery to RCs[4], and fourth, well-designed applications that are lightweight with signaling and redundant data. It is clear that adding these two pillars cannot be done without the support of the CPADs. While CPADs may be required to go against their regular trends and policies to support this four-pillar solution and enable RCOs to better manage their networks, guaranteeing a greater experience for remote users, we emphasize that we are not proposing that they do things differently for the general public.

CPADs have previously shown willingness to do things differently in developing regions, such as Africa, and we present several examples of this in Section I. However, the digital divide that we focus on in this article is different from the one that exists in Africa and not only because of the satellite. Users in Africa suffer widely from poor bandwidth and expensive service. Here, the problem impacts a small number of users living in wealthy countries where the quality of service is otherwise quite good for other users. Thus, while it seems reasonable to hope that CPADs will do things differently for RCs because they have done so in Africa, we cannot count on their goodwill alone to solve

[3]The frontier between CPADs is becoming more and more blurry; hence, we refer to them collectively.

[4]Advertisements consume network resources, but they are not likely to be well targeted for users in isolated communities.

the issue, especially since CPADs' business interests motivate them to propose solutions that are inherently "remote-unfriendly."

Instead, in Section II, we propose a new framework that involves the following.

1) The introduction of a special remote status, under the control of the International Telecommunication Union (ITU), tied to the public IP address space of the RCOs, which is used to unambiguously identify RCs that need to be treated differently.
2) A government mandate, requiring CPADs to recognize and use the remote status to give RCOs special treatment. We recognize that CPADs are diverse in size and capability; thus, we propose that they either provide dedicated solutions to RCOs that have this status, solutions that offer QoE to their users equivalent to the QoE seen by other users, or, for CPADs who cannot or do not wish to provide a dedicated solution, provide tools (e.g., use lightweight and low-encryption protocols) to enable RCOs with this remote status to implement the four-pillar solution mentioned earlier and presented in Section I.

Before proceeding, we note that although we use remote, northern communities in Canada as an example throughout, this problem is not unique to Canada, and indeed, many of the ideas presented in this work would also apply to a broader definition of the digital divide. Nonetheless, we believe CPADs may be most receptive to our proposal in the context of remote, satellite-dependent communities, who arguably suffer the most in terms of QoE.[5]

[5]While low Earth orbit (LEO) satellites have the potential to dramatically improve the experience of satellite-dependent users, especially in terms of latency, their bandwidth will remain low, and there is still debate over how prevalent LEOs will be in the near future and whether they will be widely accessible to users in extreme northern or southern latitudes.

## I. REMOTE CONNECTIVITY: CONTEXT AND CHALLENGE

In this section, we first provide an overview of the latest technology trends related to content delivery. Then, we describe the stakeholders involved in the digital divide as it impacts RCs, highlighting the ongoing challenges they face in light of the trends. This will show why the problem cannot be resolved unless content providers and application developers are involved.

### A. Recent Online Technology Trends

Today's Internet, with its proliferation of diverse applications, rich content, and encryption, is much more sophisticated than that of the past. Modern applications provide a wide variety of services over the Internet, allowing users to access content or connect with each other in new and interesting ways. However, as more and more applications have been introduced, many of them have grown more complicated. Some apps are poorly designed and inefficient, having heavy signaling, background traffic, or wasteful content (though the effects of this may not be noticeable to users with sufficient connections). Moreover, the content is much richer. In mobile and desktop environments, the number and size of objects embedded in web pages are increasing [4]. Advertisements, which are now ubiquitous in web pages and apps, have also become richer than the past [5]. Advertisements are no longer simple, low-bandwidth text or image banners, and it is not uncommon to encounter interactive advertisements or autoplaying video ads. Encryption has also become much more prevalent. Despite the fact that data and signaling are transmitted atop the same transport technologies that have existed for decades, content and application protocols are increasingly obfuscated through encryption. Already more than 50% of web traffic is encrypted today [6], [7]. Althoughmuch of this trend toward heavy encryption has been in pursuit of security and privacy, encryption has also been implemented (perhaps with noble intentions) as a mechanism to circumvent proxies and other network middleboxes interfering with sessions.

The QUIC protocol [8] presents an interesting case study of how trends can help or hinder the digital divide in RCs. QUIC is a transport-like protocol deployed atop UDP, which was originally developed by Google to improve the performance of HTTP connections and has since been taken over for standardization by the Internet Engineering Task Force (IETF) as part of HTTP/3. In many regards, QUIC is well suited for remote networks. QUIC reduces the number of round trips required to establish and secure connections by combining transport layer and Transport Layer Security (TLS) handshakes into one step. It is also designed to perform well in lossy wireless environments by eliminating head-of-line blocking by multiplexing streams within a single UDP connection. Both these features provide important advantages in networks constrained by high-latency, low-bandwidth satellites. On the other hand, QUIC uses extensive encryption. For established connections, the only cleartext data available to the network are a connection identifier and a spin bit (used for round-trip time estimation). All other fields, including packet sequence numbers and acknowledgments, are always encrypted. QUIC's encryption policies are meant to give full control over the connections to the endpoints and prevent operators from interfering with connections using proxies or middleboxes. Unfortunately, this becomes problematic in scenarios where proxies are used to enable connection splitting, as in wireless networks (we discuss this in more detail in Section I-B). In these settings, QUIC's performance diminishes versus connections carried over TCP and secured with TLS [9], [10]. Thus, while QUIC brings several benefits to remote networks, its encryption policies offset those advantages.

As will be seen in the remainder of this section, even though the impact of the above trends may not be felt negatively by the average user, the increase in diverse applications, rich content, and encrypted sessions makes solving the digital divide in RCs difficult. There is an ongoing battle for control of the Internet between CPADs and operators. CPADs are continually building more intelligence into the edge of the network through the use of proprietary appliances and protocols (such as QUIC) and are using extensive encryption techniques. This deep encryption across a greater amount of protocol data and layers prevents operators from being able to differentiate content within the network. On the whole, the majority of users have benefited from these changes by seeing improved QoE. However, a minority of users, such as those in RCs, are suffering as a consequence of these changes, and CPADs must acknowledge that problem and commit to resolving it.

### B. Stakeholder Challenges

Now, we present an overview of the stakeholders, starting from the traditional ones, and describe the challenges they face. The first stakeholders are the users in the RC. As discussed earlier, a variety of applications are available to users today, which provide all types of services. Therefore, for remote users to be able to fully participate online, they require more than just a barebones service that enables web browsing or email. In fact, most of the requirements of remote users are no different than users anywhere else. This is the key point in striving to reduce the digital divide: users in RCs must be able to participate as well as people anywhere else in their country. Hence, remote users require all the following:

1) good QoE for any type of device or service;
2) support for two-way traffic (both upstream and downstream);
3) high service availability;
4) access to the same applications as users elsewhere;
5) ability to communicate with people inside and outside their community.

The last two requirements in the list emphasize that users should preferably not need to rely on new, alternative applications. First, this ensures users traveling to and from the community are not forced to switch between applications. Second, and even more importantly, it is essential to maintain interoperation with existing devices, applications, and networks everywhere. In other words, it is unacceptable to "fix" the digital divide with walled-off solutions that digitally isolate remote users from others. In essence, solutions need to improve the experience of users communicating within the same RC, as well as those communicating between remote and nonremote locations. Ultimately, this means that any mechanisms used in ensuring good QoE for users should be transparent to them and occur within the network.

Governments and advocacy groups recognize the challenges faced by remote users and have formulated policies to attempt to bridge the digital divide in these communities. Historically, these policies have laid the primary responsibility on traditional operators by mandating network investment in remote locations and specifying minimum service requirements. Often, it has been assumed that business cases could motivate such investment, but, in recognition of the prohibitive costs to building these types of networks, governments (e.g., the Government of Canada) have also provided large subsidies to operators. Even still, keeping such networks profitable remains challenging, and traditional operators may not succeed in these environments under a for-profit model. This has led to the emergence of not-for-profit operators dedicated to RCs [11] and indicates that a need for nontraditional stakeholders has already been recognized. Having not-for-profit RCOs is likely to be an asset in the context of our proposal as CPADs will probably be happier to help these types of operators over traditional ones. Therefore, in the following, we assume that the RCO is not-for-profit.

Fig. 2 depicts an example of an RCO, which serves multiple small communities with hundreds or thousands of residents over the same shared satellite link. The RCO plays a crucial role as the one providing access to these remote users but, being not-for-profit, has limited resources at its disposal to ensure that users' needs are met. Due to the limitations of the satellite, and because multiple communities share the same satellite link, it is critical that the link is carefully managed. To do this successfully, the following must be possible.

1) *Competing traffic flows need to be differentiated and prioritized:* For example, during cases where the satellite link is congested, operators that are unable to differentiate VoIP call streams from more jitter-tolerant applications, such as web browsing, may inadvertently disrupt user experience by subjecting these two applications to the same treatment.
2) *Signaling and round-trips must be minimized:* Because of the long latency when the satellite is involved, the number of round-trips and volume of signaling over the link need to be kept down. Historically, solutions, such as TCP splitting [12], [13], which use middleboxes on the connection path to buffer data and send early acknowledgments to the endpoints, hence "splitting up" the end-to-end connection, have been used and will continue to be necessary to the RCO. CPADs need to ensure that their applications are well designed and lightweight with signaling, and they must be willing to either provide appliances in the network edge if they plan to keep control or refrain from measures (e.g., deep encryption) that prevent the use of RCO-owned proxies.

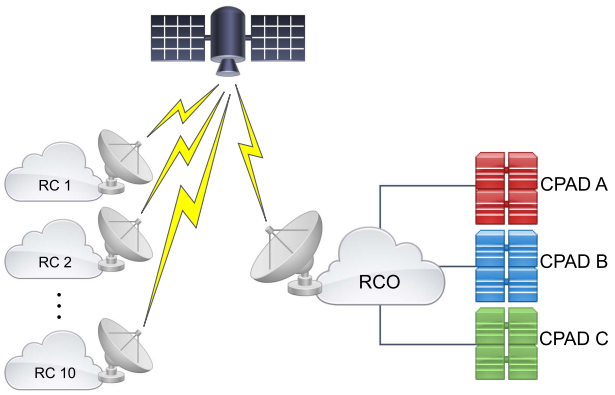


**Fig. 2.** *RCO typically serves multiple RCs using the same shared satellite link. For this reason, it is critical that the operator carefully manages signaling and traffic volume over the link. CPADs can provide their own solutions to the RC to help manage this challenge, or they can deliver content in a remote-friendly manner that enables the RCO to successfully manage the link.*

3) *Redundant content needs to be minimized:* Because bandwidth is very limited, any extra overhead on top of the user data must be kept to a minimum. This means that efficient protocols and header compression should be employed.
4) *Useless traffic should be filtered out before sending it over satellite:* Advertisements make up a nontrivial volume of connections and data [5], [14], which, in the case of RCs, not only hamper the experience of the individual user but also take resources from the rest of the network. However, ads are unlikely to be relevant to these users as living in remote, isolated locations makes it difficult for individuals to order goods online [15] or access services outside of their communities. Users will benefit greatly if advertisements are filtered out ahead of time.
5) *Traffic should be kept within the community as much as possible:* Earlier work has shown a high locality of interest [16] among users in rural communities [15], meaning that individuals interact with local users over social media more frequently than with nonlocal users (i.e., social media traffic is mostly local). Therefore, a strong case can be made for placing application-specific appliances locally in each community; nevertheless, this might not be economical from a CPAD standpoint, and in that case, an RCO-based solution needs to be implemented with their agreement. This would have the benefit of simultaneously reducing latency for users interacting with the content and freeing up backbone resources over satellite.
6) *Volume of repeat traffic over the satellite link must be minimized:* We would also expect that there should be a high shared interest in web and multimedia content. In this case, popular content should not be repeatedly sent over the satellite as it wastes bandwidth, taking resources away from other users. Instead, it should be sent once and cached within the community. This is particularly true for bandwidth-intensive content, such as videos, software updates, or video game downloads.

Combined, these points indicate that a joint four-pillared solution that relies on proxies, well-designed applications, filtering, and caching is needed to provide the best QoE for remote users. In fact, this type of solution will become more indispensable in the face of the latest trends. As content gets richer and applications grow more complex, proxies will become more necessary to manage traffic and signaling. It will also be essential to ensure that applications are well designed, by being lightweight with signaling and redundant content to reduce overhead. Similarly, as advertisements become richer, the ability to filter them will grow even more vital. Finally, as users engage with a greater volume of social media and online entertainment, it will be all-the-more essential to keep that content local through specialized appliances and caches in the RC.

On the other hand, the latest trends threaten the ability of RCOs to successfully solve the digital divide on their own. As encryption becomes widespread and CPADs exploit it to take more control, it disrupts the ability of RCOs to depend on proxies and connection splitting techniques, makes it impossible for them to differentiate, filter, or cache content, and adds additional overhead to connections, especially if implemented carelessly across multiple protocol layers. When implemented broadly, this deep encryption severely threatens the ability of remote operators to provide an acceptable service to their users.

To date, the traditional stakeholders have been unable to fix the digital divide for RCs, even in developed countries. The latest trends of CPADs threaten the ability of RCOs to successfully deploy solutions. Therefore, RCOs cannot be held solely responsible for solving this problem. Content providers and application developers must acknowledge the challenges faced by RCOs and commit to helping, either by providing their own RC-friendly solutions or by providing RCOs with the tools necessary to succeed.

### C. Learning From Africa

It is clear that there is space for CPADs to help, but will they be willing to do so? We believe some will. To support our view, we highlight several examples where CPADs have been involved in addressing the digital divide in the past, specifically in

developing regions with resource and latency-limited networks. The purpose here is to showcase past willingness on the part of CPADs to recognize and address the issue, often by deviating from their standard policies and trends.

In some cases, this has manifested through the development of alternative applications. For example, as an alternative to YouTube's mobile application, Google developed YouTube Go [17], which it makes available in developing countries. This application provides picture previews so users do not waste network resources when choosing between videos, and allows users to queue and download videos if their network conditions are not suitable for streaming. This deviates from YouTube's standard policy of actively preventing video downloads.

Other times, CPADs have created custom applications to act as proxies through which users access low-bandwidth versions of websites. Facebook developed the Free Basics application [18], which provides a common portal through which web hosts can provide lightweight versions of their websites to users. This application is also only available in select developing nations. When users access a website through the application, the request is served through Free Basics' proxy, which fetches the light version of the web page instead. This shows that CPADs could provide their own proxies if they are not comfortable with giving control of their content over to the RCO; however, it would make it more difficult for the RCO to manage the satellite link.

The Opera Mini [19] mobile web browser is another proxy-based application. Here, the browser acts via an intermediate server to deliver compressed web pages to a user's device. Web requests are first returned to one of Opera's proxy servers where the content is compressed by a factor of 50%–90% before delivery to the user's browser. This breaks end-to-end encryption, as the connection is divided between the user and the proxy, and the proxy and the server. Opera Mini is available everywhere; however, its main user base is in Africa, where broadband availability is scarce and mobile data prices are among the highest globally relative to income levels.

The above three examples require users to download and install special applications. Small-scale in-network solutions that make use of proxies and require no change on the part of users have also been proposed. In [20], researchers implemented a localized Facebook server in a remote Zambian community. Here, the proxy server improved QoE among Facebook users in the community by locally caching data shared to Facebook and keeping signaling within the community where possible instead of using the satellite backhaul. Once again, proxy solutions like this can only be made possible by breaking end-to-end encryption.

In each of these instances, CPADs have offered users access to services via solutions that create less overhead in the network or deliver lightweight versions of content. In either case, they have accounted for the fact that users are connecting through less robust and efficient networks. These examples also show their willingness to support cache- or proxy-dependent solutions that break conventional trends. This indicates that CPADs recognize the challenges at hand in resource- and latency-limited networks and are willing to do things differently to the benefit of end users and ought to motivate many of these parties to adopt a similar view toward remote networks.

## II. TOWARD REMOTE-FRIENDLY SOLUTIONS

Evidently, some CPADs are willing to help address the digital divide in other contexts, and we are optimistic that those same parties will be willing to help in the remote context as well. Nevertheless, there may be others who are less eager to help due to the small number of users in RCs and their reluctance to give up some control. Not-for-profit operators are all but helpless if left to fix this issue without the help of CPADs. Hence, CPADs should be held responsible for helping solve this issue. To that end, we believe that policies should be drafted mandating CPADs to help RCOs; otherwise, remote users will be left behind. Moreover, existing policies that place the burden of responsibility on operators need to be extended to include CPADs.

Toward that goal, we propose a framework that mandates CPADs to help but allows them to contribute under one of two possible schemes. The first scheme allows CPADs to maintain control of their content but then requires them to supply their own solution offering QoE to remote users that is comparable to users elsewhere. The second scheme places the burden of providing good QoE on RCOs but necessitates CPADs to deliver content in "remote-friendly" ways so that RCOs can successfully manage QoE. By avoiding a one-solution-fits-all approach, we give some flexibility to CPADs. Those who want to keep control and are willing to invest in RCs may do so using their own solution. The downside is that the RCO might not be able to manage the satellite link as effectively (e.g., traffic differentiation may not be possible).

We reemphasize that, regardless of which of the two schemes is chosen, the process needs to be transparent to end users and not depend on them downloading custom applications. We now describe both schemes in more detail. We will discuss later how a CPAD determines that a user is remote and needs this special treatment.

### A. CPAD-Controlled Solution

Based on what has been seen in developing regions, some CPADs may prefer to install their own proxies or caches in RCs as a way to manage QoE and fulfill the government mandate. This scheme allows CPADs to keep control of their content while keeping it close by and readily accessible to users. In fact, some CPADs already provide similar

solutions to large operators. The difference in doing this for an RC is that it comes at the cost of scalability to the CPADs. Recall that an RCO may serve one or several small communities, each with a population of several thousand people at most. Thus, if a CPAD puts a dedicated appliance in every community served by the RCO, this will lead to a proliferation of appliances. This may be impractical for most CPADs: the solution does not scale well relative to the small number of users, is very costly, and will be difficult to implement and manage. Therefore, there is a limited business case for many CPADs to pursue this type of solution, and we would expect only the largest ones to be willing to take that decision.

This scheme would also require the CPAD to block advertisements on its end, as encryption would prevent the RCO from being able to filter ads. Once again, we expect this to be feasible as previous solutions provided by CPADs to ease the digital divide have forgone advertisements.

### B. RCO-Controlled Solution

If a CPAD is unwilling to supply its own solution to each community (perhaps because it does not make business sense), then it will instead need to enable the RCO to implement an appropriate solution. In this case, the CPAD must still help the RCO by accounting for the effects of the aforementioned technology trends and delivering content differently. To ensure that the RCO is successful, the CPAD must be: 1) lightweight in signaling; 2) support new policies for content caching; and 3) not practice deep encryption.

*1) Being Lightweight:* As was done with QUIC, CPADs must incorporate lightweight designs and protocols into their applications to minimize the impact on remote users.[6] Being lightweight means using less signaling, which causes long delays over satellite, as well as limiting the volume of superfluous traffic, which consumes scarce backhaul capacity.

In terms of signaling, QUIC sets a good example of what remote-friendly protocols should look like. Therefore, our recommendation is that protocols used to reliably transport application data to RCs should incorporate concepts from QUIC's design for faster and more resilient connections.

Application developers must develop content in ways that minimize overhead. Given the proliferation of smartphone use over the past decade, many CPADs are already careful to use lightweight codecs and protocols for multimedia, and this trend must continue. Remote users benefit from this phenomenon, especially with the higher use of mobiles relative to desktop computers in their communities [14]. Header compression techniques are also vital to reduce the overhead from redundant data in the headers of different protocol layers.

[6]In fact, the need for apps to be lightweight is true for all types of networks; however, it is even more critical in remote networks.

There is still the issue of advertisements. Earlier, it was suggested that RCOs filter out this content before it crosses the satellite; if the CPADs practice deep encryption, this may be more than RCOs are capable of handling on their own. Ultimately, CPADs should recognize the futility in delivering ads to remote users. We recommend that CPADs cease sending ads to RCs altogether since they waste network resources and are of limited benefit to users.

*2) Supporting New Caching Policies:* As was mentioned in Section I, caching popular, high-bandwidth content is critical for RCOs maintaining good QoE for the community. The benefits from caching are not unique to remote operators; nonetheless, the way in which caches are implemented in this setting will need to be different than usual. Instead of implementing independent caches for each CPAD, which will scale poorly, we propose that RCOs be authorized to cache content within a single aggregate cache in each community. This goes against current policies that prohibit operators from using in-network caching solutions. Therefore, this would require cooperation among both parties, with CPADs recognizing RCOs as a trusted partner. One challenge is that copyright policies may need to be redefined around RCOs storing content. CPADs must acknowledge that things will need to be done differently and be open to supporting such a policy.

*3) Rethinking Encryption:* Even if CPADs authorize RCOs to cache content, they will have no ability to do so if application data are encrypted. Encryption will force RCOs back into a situation where CPADs will have to independently manage private caches in every community, but, as was already discussed, some CPADs will be unwilling to take that direction because there is no business case for them. Encryption also prevents RCOs from performing traffic differentiation, especially as not-for-profit operators cannot be expected to rely on state-of-the-art deep packet inspection techniques to differentiate encrypted traffic. Finally, as CPADs take control of connections through deep encryption, this will prevent all possibility of RCOs using proxies and connection splitting that they have depended on to succeed in the past.

Therefore, this trend toward blindly encrypting everything will put too much pressure on the digital divide. Consequently, CPADs need to understand the challenges that encryption poses to remote networks and think differently about how they encrypt sessions in these scenarios. When it comes to encryption in this context, we believe it is important for CPADs to ask themselves when, what, and how?

As a first step, protocols used to deliver content to RCs need to support connection splitting that has made satellite-based connections possible in the past while also limiting the amount of signaling. QUIC's design for low-signaling connection establishment and high performance over lossy links is a good first step, but, in its current state, QUIC fails

to recognize that there will be users whose network experience will suffer under its encryption policy. Thus, we propose that a "remote QUIC" variant should be standardized and supported by the IETF and CPADs, which incorporates QUIC's good features but does not encrypt connection-related data that would otherwise prevent connection-splitting.

Beyond keeping connection-related data unencrypted, there are also scenarios where CPADs should leave content unencrypted to enable RCOs to utilize caches and perform traffic differentiation. Not only does this require CPADs to trust RCOs to treat content responsibly but it also requires users to trust RCOs by giving up some of their privacy. However, this might be regarded as an acceptable alternative for users, particularly within some RCs where there exists a cultural expectation that technology should benefit the community as a whole instead of users individually [21].

Even though content should not always be encrypted, encryption should remain in force for communications and exchanges involving secure credentials. The point is not to abandon encryption altogether, as this would create another digital divide where remote users lack digital privacy and security; nonetheless, we do want to challenge the notion of encrypting most headers and all content without thinking of the repercussions.

### C. Giving RCOs Special Status

The proposed framework requires CPADs to be able to distinguish remote users from users elsewhere so that they can give them special treatment. How then can CPADs identify remote users? First, we assume that users served by an RCO are behind a common NAT, which maps the private IPs of users in the RC to a set of public IP addresses. Then, CPADs do not need to differentiate individual users: in order to tell if an incoming connection is from a remote user, it is enough for CPADs to know the RCO's public IP addresses. In many regions, CPADs are already using IP geolocation to determine whether they will deliver alternate content to users (see Section I-C). Therefore, CPADs should likewise be able to identify remote users from the public IP of their RCOs.

This framework depends on CPADs to respond appropriately to connections originating from the remote IP range. While we are optimistic that they will cooperate, we believe that the backing of other authorities will be needed to enforce this. With that goal in mind, we recommend that the ITU creates a special remote status, which is granted to RCOs and associated with their public IP address space. Under government mandate, CPADs would then be required to recognize this remote status and use it to determine who should receive special treatment through the schemes described above.

### D. Recognizing Remote-Friendly Applications

CPADs should be motivated to follow this mandate as a matter of goodwill, but, in order to hold them accountable, there should be a classification system that recognizes which applications are better than others for RCs. For instance, applications could be "certified remote-friendly" (e.g., in the Apple App Store or Google Play Store) if they stand out among others. Taking this a step further, a formal rating system could even be established to rank applications based on how well they enable RCOs to efficiently manage the user QoE (similar to how energy ratings are assigned to household appliances).

With this in mind, the next step would be to study which popular applications from different service categories (voice/video/text chat, video streaming, and web browsing) are better for RCs, from which a public recommendation system could be made available. Researchers have previously tried to evaluate which apps are better for certain environments, but, as far as we are aware, this has not been studied in a remote context.

## III. CONCLUSION

The digital divide continues to be an issue in RCs, even in developed countries. The modern online landscape poses new threats to this problem due to the proliferation of diverse applications, rich multimedia content, and deep encryption. Up until now, RCOs have depended on caching and proxies to effectively manage signaling and traffic over the satellite link in order to provide good QoE to remote users. These existing solutions will continue to be vital as applications and content grow richer. It is also critical that these applications are well designed, that is, lightweight with signaling and redundant content. Moreover, the filtering of advertisements that are unlikely to be well targeted for remote users and waste network resources should be provided. However, RCOs, particularly not-for-profit ones, will not be able to keep up with the growing complexity of traffic on their own, especially in the face of increasing encryption. Moreover, they have no control over how applications are designed. Therefore, RCOs should not be held solely responsible for providing technical solutions to this problem. CPADs must also be brought along as stakeholders to recognize the challenges faced by RCOs and help them bring content to users in more remote-friendly ways.

To that end, we have proposed a framework requiring cooperation among CPADs, RCOs, the government, and ITU. Within this framework, the ITU should standardize a remote status associated with the public IPs of RCOs, which is used to identify remote users. The government should mandate CPADs to help address the digital divide in the remote context and use this status to recognize remote users and treat them specially. Under this mandate, CPADs would be given the flexibility to choose between two possible remote-friendly schemes. CPADs can supply their own solutions to RCs; in that case, they keep control of their

content but may incur high costs in doing so. Alternatively, they can leave the solution to the RCO to manage QoE; in that case, CPADs need to be lightweight with signaling, support new caching policies, and scale back encryption.

Given that CPADs have previously addressed the digital divide in developing regions, we are optimistic that they will be willing to help in the remote context as well. While it is true that many of the ideas presented in this work could also be applied to a broader definition of the digital divide, we believe that CPADs may be most receptive to our proposal in the context of remote, satellite-dependent communities who arguably suffer the most in terms of QoE. Furthermore, unlike other contexts where both high- and low-bandwidth users could be behind the same operator's NAT and, therefore, not easily distinguishable to CPADs, all users behind an RCO's NAT suffer from poor bandwidth and latency. Therefore, associating the special status of RCOs with their public IPs ensures that CPADs can be confident that they are giving special treatment only to those who require it. ■

## Acknowledgment

The authors would like to thank Prof. S. Keshav, Cambridge University, and the reviewers for their useful comments on an earlier version of this article.

## REFERENCES

[1] M. Haight, A. Quan-Haase, and B. A. Corbett, "Revisiting the digital divide in Canada: The impact of demographic factors on access to the Internet, level of online activity, and social networking site usage," *Inf., Commun. Soc.*, vol. 17, no. 4, pp. 503–519, Mar. 2014, doi: 10.1080/1369118X.2014.891633.

[2] L. Robinson *et al.*, "Digital inequalities and why they matter," *Inf., Commun. Soc.*, vol. 18, no. 5, pp. 569–582, Mar. 2015, doi: 10.1080/1369118X.2015.1012532.

[3] Innovation, Science and Economic Development Canada. (Aug. 2020). *National Broadband Internet Service Availability Map*. Accessed: Sep. 1, 2020. [Online]. Available: https://www.ic.gc.ca/app/sitt/bbmap/hm.html?lang=eng

[4] T. Johnson and P. Seeling, "Desktop and mobile Web page comparison: Characteristics, trends, and implications," *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 144–151, Sep. 2014, doi: 10.1109/MCOM.2014.6894465.

[5] J. Gui, S. Mcilroy, M. Nagappan, and W. G. J. Halfond, "Truth in advertising: The hidden cost of mobile ads for software developers," in *Proc. IEEE/ACM 37th IEEE Int. Conf. Softw. Eng.*, Florence, Italy, May 2015, pp. 100–110, doi: 10.1109/ICSE.2015.32.

[6] Sandvine. (Oct. 2018). *The Global Internet Phenomena Report-October 2018*. Accessed: Jun. 6, 2019. [Online]. Available: https://www.sandvine.com/hubfs/downloads/phenomena/2018-phenomena-report.pdf

[7] D. Naylor *et al.*, "The cost of the 'S' in HTTPS," in *Proc. 10th ACM Int. Conf. Emerg. Netw. Exp. Technol. (CoNEXT)*, Sydney, NSW, Australia, Dec. 2014, pp. 133–140, doi: 10.1145/2674005.2674991.

[8] A. Langley *et al.*, "The QUIC transport protocol: Design and Internet-scale deployment," in *Proc. Conf. ACM Special Interest Group Data Commun. (SIGCOMM)*, Los Angeles, CA, USA, Aug. 2017, pp. 183–196, doi: 10.1145/3098822.3098842.

[9] L. Thomas, E. Dubois, N. Kuhn, and E. Lochin, "Google QUIC performance over a public SATCOM access," *Int. J. Satell. Commun. Netw.*, vol. 37, no. 6, pp. 601–611, Feb. 2019, doi: 10.1002/sat.1301.

[10] J. Deutschmann, K.-S. Hielscher, and R. German, "Satellite Internet performance measurements," in *Proc. Int. Conf. Networked Syst. (NetSys)*, Munich, Germany, Mar. 2019, pp. 1–4, doi: 10.1109/NetSys.2019.8854494.

[11] A. Fiser and A. Clement, "K-Net and Canadian Aboriginal communities," *IEEE Technol. Soc. Mag.*, vol. 28, no. 2, pp. 23–33, Summer 2009, doi: 10.1109/MTS.2009.933028.

[12] T. R. Henderson and R. H. Katz, "Transport protocols for Internet-compatible satellite networks," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 2, pp. 326–344, Feb. 1999, doi: 10.1109/49.748815.

[13] J. Border, M. Kojo, J. Griner, G. Montenegro, and Z. Shelby, *Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations*, RFC 3135, Internet Requests for Comments, RFC Editor, Jun. 2001. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3135.txt

[14] E. Showalter, N. Moghaddas, M. Vigil-Hayes, E. Zegura, and E. Belding, "Indigenous Internet: Nuances of Native American Internet use," in *Proc. 10th Int. Conf. Inf. Commun. Technol. Develop. (ICTD)*, Ahmedabad, India, Jan. 2019, pp. 1–4, doi: 10.1145/3287098.3287141.

[15] M. Vigil, M. Rantanen, and E. Belding, "A first look at tribal Web traffic," in *Proc. 24th Int. Conf. World Wide Web (WWW)*, Geneva, Switzerland, May 2015, pp. 1155–1165, doi: 10.1145/2736277.2741645.

[16] M. P. Wittie, V. Pejovic, L. Deek, K. C. Almeroth, and B. Y. Zhao, "Exploiting locality of interest in online social networks," in *Proc. 6th Int. Conf. (Co-NEXT)*, Philadelphia, PA, USA, 2010, pp. 1–12, doi: 10.1145/1921168.1921201.

[17] Google Design. *The Making of YouTube Go*. Accessed: Dec. 23, 2019. [Online]. Available: https://design.google/library/making-youtube-go

[18] Facebook Connectivity. *Free Basics*. Accessed: Dec. 23, 2019. [Online]. Available: https://connectivity.fb.com/free-basics

[19] Opera Software. *Opera Mini*. Accessed: Dec. 23, 2019. [Online]. Available: https://www.opera.com/mobile

[20] D. L. Johnson, V. Pejovic, E. M. Belding, and G. van Stam, "VillageShare: Facilitating content generation and sharing in rural networks," in *Proc. 2nd ACM Symp. Comput. Develop. (ACM DEV)*, Atlanta, GA, USA, Mar. 2012, pp. 1–10, doi: 10.1145/2160601.2160611.

[21] K. Gibson, M. Kakekaspan, G. Kakekaspan, S. O'Donnell, B. Walmark, and B. Beaton, "A history of everyday communication by community members of Fort Severn First Nation: From hand deliveries to virtual pokes," in *Proc. iConference-iConference*, Toronto, ON, Canada, Feb. 2012, pp. 105–111, doi: 10.1145/2132176.2132190.