WebPack: A Path to Caching in Remote Edge Networks?

UNIVERSITY of WASHINGTON

Spencer Sevilla, **Matt Johnson**, Kurtis Heimerl

IETF 106

Background/Context

- We are researchers at the University of Washington
- Focus on Internet access in rural/remote/developing areas

- **Places:** Philippines, Papua, Argentina, Mexico, Arctic Circle
- **Projects:** Endaga, OsmoCom, OAI, NextEPC, CoLTE
- **Friends:** AirWaves Missions, Puspendes, Rhizomatica, AlterMundi, UP VBTS (CoCoMoNets)

Network Characteristics: Ownership

- We focus on community networks: Networks owned and operated by local actors within the community
  - These networks are generally small-scale (N < 1000) and do not seek to grow beyond their natural community boundaries [1]
  - Networks owned collectively or run as a small business, sometimes informally
- Key Point #1: These ISPs are generally too small to have a business relationship with large CDNs, content providers, etc.
  - Also language & cultural barriers
- Key Point #2: These ISPs have local infrastructure close to their users

[1] S. Hasan, Y. Ben-David, M. Bittman, B. Raghavan. The Challenges of Scaling WISPs. DEV 2015.

Network Characteristics: Backhaul

- Most places we work have **very constrained backhaul**
  - Papua: 3 Mbps / ~150 users
  - Philippines: 2G access only
  - Oaxaca: 2 Mbps / ~300 users

These are real numbers in 2019 : /

- These metrics are often over-estimated and **flaky**
  - Papua: VSAT has a 10/1 contention ratio and 3000ms pings (from Seattle)
  - Oaxaca: 3+ chained point-to-point long-distance WiFi links induces high jitter/loss

Network Characteristics: RAN

- Local RAN is orders of magnitude faster than backhaul
  - Even at “1 bar” of signal-strength, the bottleneck is still upstream!

- Local LTE: 150 Mbps over a wide area (> 1 mile radius)
- WiFi: even faster speeds (802.11n = 450 Mbps) but less range

Network Characteristics: Traffic

- Observation: Most interactions stay local
  - 91% of tribal Instagram media interactions stay within the tribe [1]
  - Likely applies to other communication methods (phone calls, SMS, …)
  - Not a very surprising result: social networks are driven by social connections
- Observation: Most rural traffic is what you might expect
  - HTTPS sessions to YouTube, FB, Google, CloudFlare et. al. [2]
  - Not a very surprising result in 2019

[1] M. Vigil, M. Rantanen, E. Belding. A First Look At Tribal Web Traffic. WWW 2015.
[2] S. Sevilla, M. Johnson, P. Kosakanchit, J. Liang, K. Heimerl. Experiences: Design, Implementation, and Deployment of CoLTE, a Community LTE Solution. MobiCom 2019.

Our Interest In WebPack

- Key Goal: Optimize local traffic off the backhaul whenever possible!
- Complicated with HTTPS
  - Too small to locally deploy hardware from major CDNs
  - Current *fundamentally unsatisfying* solutions are
    - Mirroring content at a local domain (youtube.{mynetwork}.org)
    - Distributing CA certs to trust a cache and MitM HTTPS connections
  - No good way to separate out sensitive vs. non-sensitive content at the network level

Case #1: Redistribute software updates

- Many users disable updates since they are large and bandwidth is precious
  - Windows, Android Applications, etc.
  - Obvious security risk
- The same binary is downloaded hundreds of times over the backhaul bottleneck via TLS
- WebPack could be used to cache the update transparently in the network, with strong authenticity guarantees

Case #2: Popular media (e.g. videos, articles)

- A multimedia article goes locally viral
- The content producer judges pieces of the content non-sensitive and packages some heavy components of it
- The multimedia is locally cached, leaving the backhaul bandwidth for unique user-specific traffic
- As the article is shared via Facebook, Viber, or email, local users on all platforms benefit from the common cache

Case #3 Breakout local “application functionality”

- WPACK could allow a platform to distribute a large authenticated web application in smaller chunks independent of content
- The application can package more local functionality (more code) while remaining responsive and quick to load
- Locally popular features benefit from in-network caching

Our Major Asks

We work with tiny ISPs who can't exactly get on the phone with major media outlets, CDNs, or platforms (Google, Facebook, et. al.)

- Driving towards an open standard (and then hoping content producers widely adopt it) is probably the best way forward for us
- WebPack looks like the closest thing to our needs

- **Key Ask #1:** Please keep our use-cases in mind!
- **Key Ask #2:** Work with us and we can meet you in the middle

Thanks for your consideration!

Clarification Questions?

//

Backup slides follow…

How We Can Contribute

- Unique position: The ISP consists of its end users, deep trust exists
- We are networking researchers: comfortable with systems hacking, network architecture, protocol design, etc.
- We are also social scientists: interested in usable security and privacy
- Able to explore new trust-models, roll out custom apps, etc.
- We love dogfooding/beta-testing stuff

Coverage expansion by traditional ISPs and MNOs is **slowing** as they reach the economic limits of their deployments...

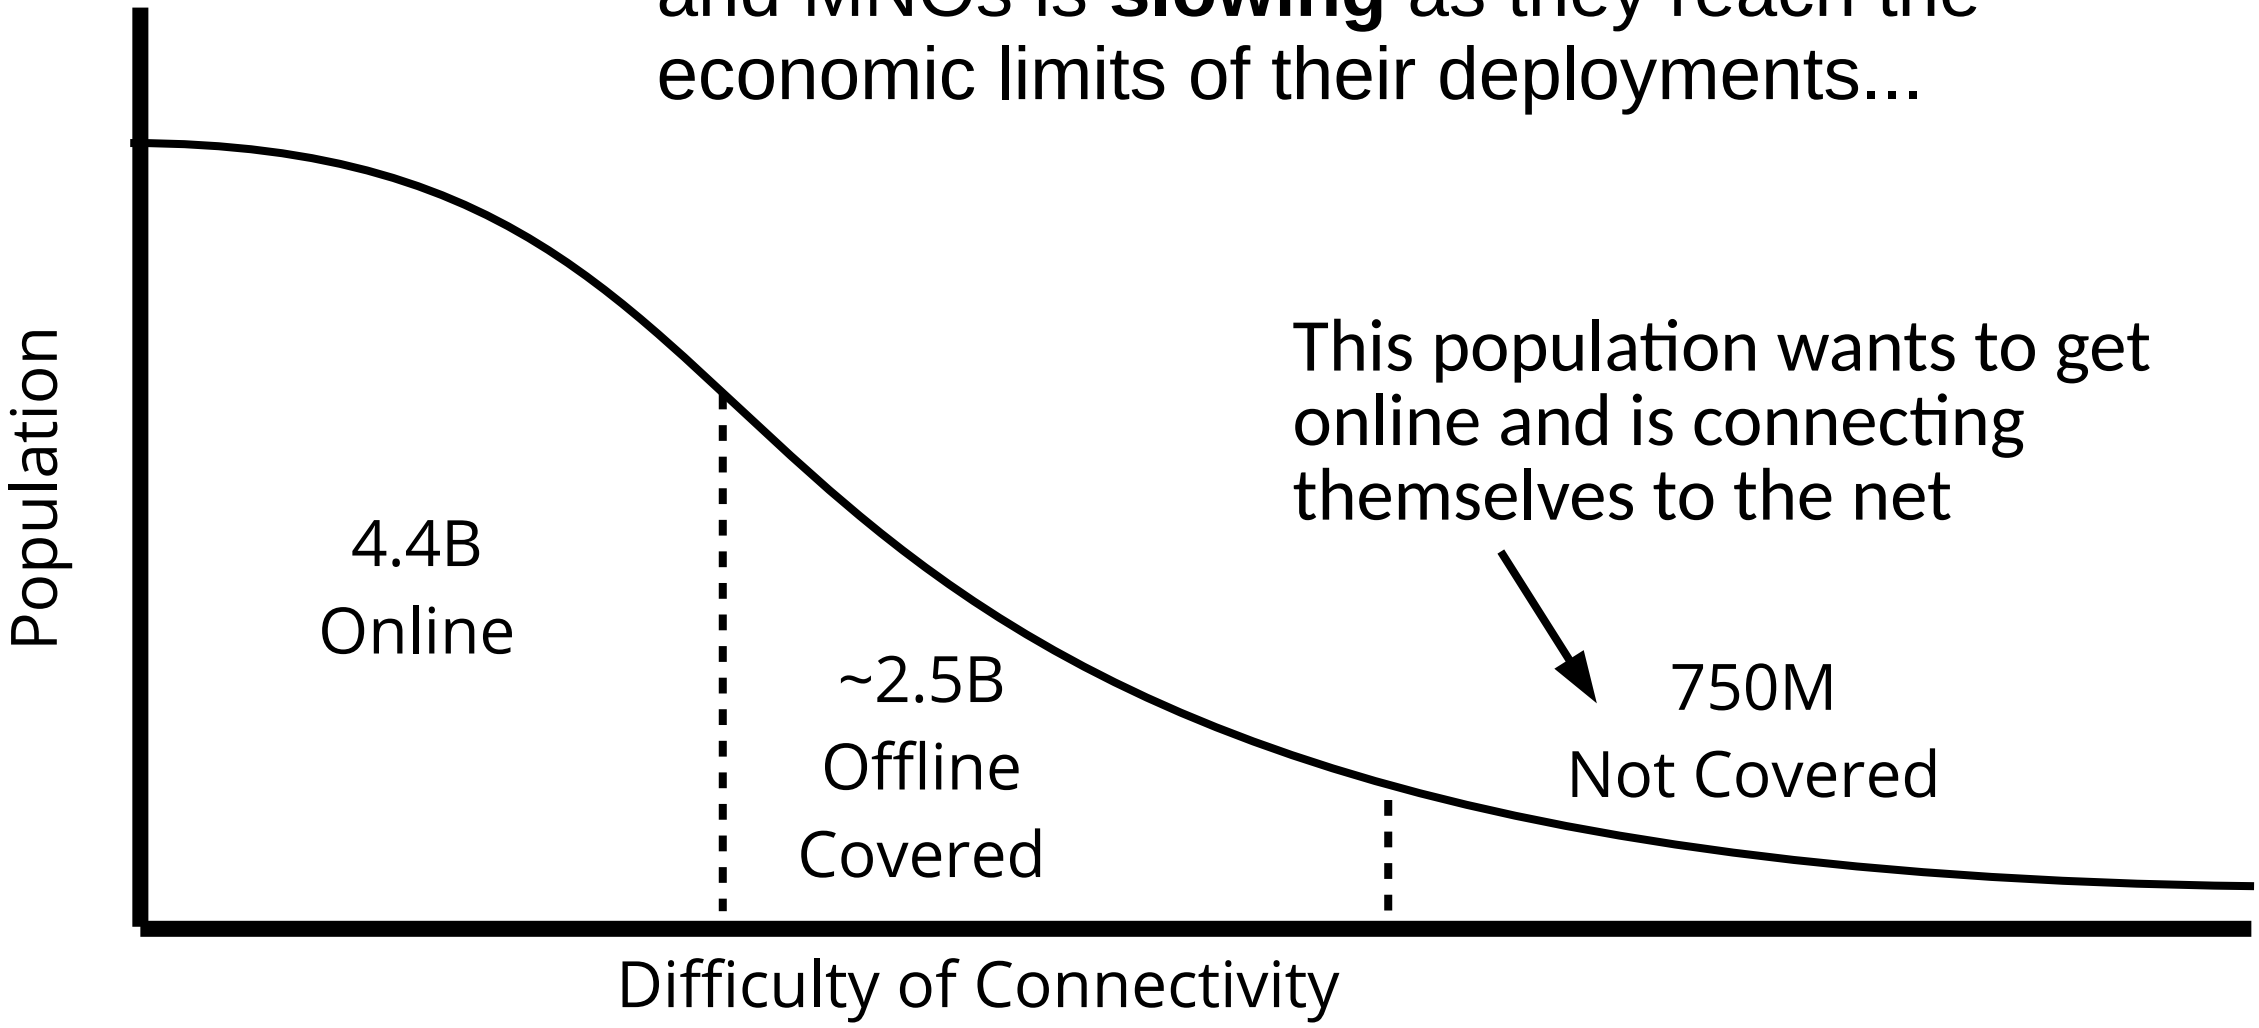

https://wearesocial.com
GSMA-State-of-Mobile-Internet-Connectivity-Report-2019