In response to the IAB's call for position papers to inform its workshop on IP geolocation, I offer the following notes.

IP geolocation starts from the flawed assumption that an IP address is associated with an interface of a device that has a specific location in space. Even such trivial and common technologies as VPNs and containerization create a mismatch between the expected location and the intent of the geolocation user. Worse, IPv4 transfer and leasing markets make IPv4 addresses even more mobile than traditionally expected. Few of the major use cases support making a concerted effort to improve the system, and there may be dangerous unintended consequences.

The most common use cases of IP geolocation seem to be:

1. **Serving "localized content" (advertisements) to users.** Although uses such as "Here's a coupon for the ice cream shop you're about to pass" are often teased, IP addresses are insufficiently granular for this. Although a Wi-Fi access point might have a single IPv4 address that could be used for geolocation, it is often backhauled to a controller on campus or service provider, providing only the area of town. IPv6 addresses could provide end-to-end detail, but determining whether a location is served by a /64, /56, /48, or some number in between, adds complexity. IPv6 local prefixes will still be specific to the layer 2 domain only.
2. **Directing traffic to local servers.** Some operators may use geolocation information to return DNS responses to systems local to the user. This is better achieved with anycast and BGP for finding shortest path, since round trip time correlates to network topology, which may not match geography.
3. **Complying with regional contracts.** Some content is only available to viewers in a specific location. This is often due to contracts, e.g., between video streaming services and content rights holders, or sporting event broadcasters and local teams. In other cases, governments exert control over the content available to people in their jurisdiction. In still other cases, shopping sites may only serve their shipping area. VPN services and proxy farms (such as "sneaker proxies") are so common as to make these limitations meaningless.
4. **Blocking foreign countries deemed hostile**. It is apparently common for some businesses' firewalls to block all IP addresses from some or all foreign countries. This may seem reasonable to protect against state-sponsored (or not) cyberattacks. Again, the use of VPNs makes such blocks essentially useless. Worse than being ineffective, they block legitimate traffic. For a real example, a bank in Iowa blocks all traffic from IPv4 addresses in Asia. They assume that none of their customers would ever need to check their bank accounts while traveling in Asia. An ISP in their town bought a /24 from a company in Singapore, and suddenly local customers could not reach their bank.

VPNs and proxy farms provide workarounds for many of the problems above. The networks using IP geolocation have therefore responded by blocking suspected VPN traffic. This feeds into the IPv4 address leasing market, where VPN providers or proxies use addresses for a few months, then return them and lease new addresses. Many content providers try to build

inference databases based on activity, but their data quickly becomes stale, potentially harming innocent users who next lease the addresses. In some cases, the providers are not under single control, and containerized processes may be moved to different locations.

Geofeeds are helping in some ways, and should be used more widely. However, VPN operators and proxy networks can easily provide false information. Operators blocking connections based on geolocation should, rather than silently dropping, provide alternate pages explaining how to (get their ISP to) set up geofeeds. If web content is blocked due to government geolocation policy, an ERROR 451 may be appropriate.

Geolocation providers take too long to update their records; two to four weeks is reportedly[1] standard. When IPv4 transfer logs and geofeeds files exist, this is a frustrating delay. Some geolocation providers allow contact from IP address holders to correct their records; The Brothers WISP have a frequently-cited web page that tries to consolidate that information at https://thebrotherswisp.com/index.php/geo-and-vpn[2]. However, many of the databases are private, so address holders don't even know why connections are failing, much less who to contact when there's a failure.

There have been proposals in the past to require traffic not to transit a particular country, or to remain entirely within a country. These are network topology challenges masquerading as geolocation issues for questionable policy reasons. Any effort to ascertain location by IP address will inevitably support these efforts to worsen Internet routing.

Fundamentally, IP geolocation is a technique to overload address integers with meaning. That is architecturally unsound: IP addresses do not have meaning or attributes. Any use of addresses as identity is misguided. Network operators that need to know the geographic location of users need information about the user's location, not the address. They can trust self-reporting, as in geofeeds or user account information, or establish agreements with other network operators. More complex technological solutions such as ALTO[3] have not seen widespread adoption because cost outweighed benefit.

Bluntly: I realize that the Program Committee for this workshop comprises people who work for organizations with the needs above; I am willing to be persuaded that these are important technical requirements for the benefit of the end users[4]. As a general principle, ascribing semantics to identifiers is poor architecture.

Lee Howard
Lee@Asgard.org

[1] https://www.ipv4.global/events/geolocation/
[2] https://thebrotherswisp.com/index.php/geo-and-vpn
[3] https://datatracker.ietf.org/wg/alto/about/
[4] https://datatracker.ietf.org/doc/html/rfc8890