Randomized Response Mechanisms in RTT Measurements for QUIC

`draft-andersdotter-rrm-for-rtt-in-quic-00`

Amelia Andersdotter (CENTR), Shivan Kaul Sahib (Salesforce)

Meetecho, teh InterwebZ (IETF), July 2020

In PEARG at IETF104 there was a presentation on differential privacy,[1] which mentioned random response mechanisms (RRM)

The idea was raised that RRM might be applicable to RTT measurement ideas raised in the QUIC WG

Fast forward... Now results of this idea are codified in a draft.

[1] https://datatracker.ietf.org/meeting/104/materials/slides-104-pearg-amelia-christoffer-differential-privacy-00

Differential privacy...

...is a way to mathematically ascertain a certain level of privacy protection in a system, for some definition of privacy and system.

...is meant to preserve data utility for the user of data (the observer in the case of RTT measurements) while protecting the unique identity of contributors of data (typically the client or server user in the RTT measurement case).

...uses statistical mechanisms to accomplish mathematical guarantees of simultaneous data usability and data obfuscation.

While RRM could potentially contribute to some measure of RTT privacy, it is unlikely to be worth the effort.

1. It does not remedy the privacy concerns actually raised in QUIC WG (the utility, to be preserved, is in fact the main privacy concern) instead contributing only to a fewer spin bit measurements being useful to the observer
2. It could provide a mechanism to allow for a larger degree of client control over RTT measurements but requires much more precise specification of the latency spin bit

RRM was created to ensure a higher level of privacy for individuals that participate in surveys.

An individual survey-taker can be guaranteed that the survey-giver does not know which question the survey-taker responded to (e.g. “do you have diabetes” versus “do you not have diabetes”). The survey-giver does know the probability that either question was answered.

An alternative way of thinking about this is that the survey-taker *lies in their response to a single (binary outcome) question* with some probability.

No RRM: two options give two possible outcomes

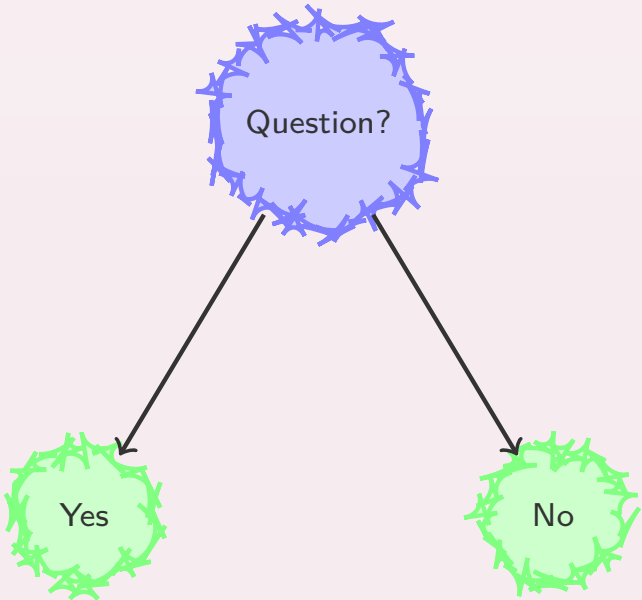

With RRM: two options give four possible outcomes

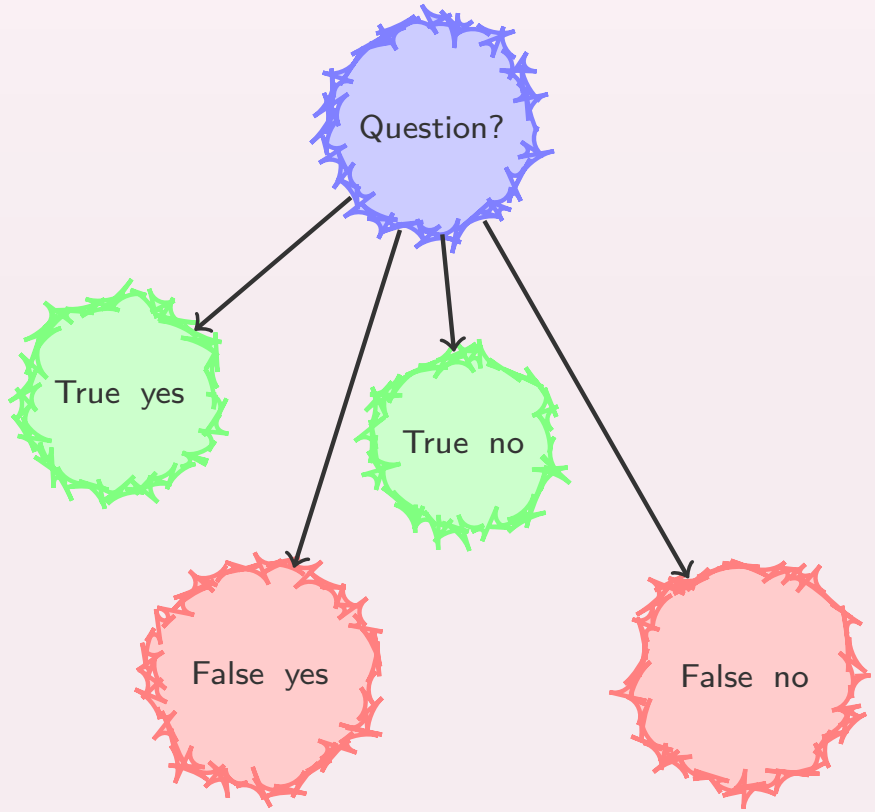

Problem: estimate number of *yes* and *no* knowing only the total numbers of *yes* and *no* responses, plus the assumed proportion of *truth-sayers* versus *false-sayers*.

This problem can in general be solved for non-binary outcome spaces too, but it's more complicated (see references).

The latency spin bit is a bit which has two states.

Idea: apply RRM to the state of the latency spin bit to increase privacy.

Practically: figuring out when to actually activate a RRM.

1. Required 10 additional spin bit assumptions (sec. 4)
2. And a model specification (sec. 5)
3. And a round-trip explanation (sec. 7.1)
4. And truth-tables for spin bit values (sec. 4)

There is a simulation available on github: `https://github.com/ShivanKaul/draft-andersdotter-rrm-for-rrt/tree/master/spinbit_simulation`

Sec. 7.1 describes one and a half round-trips in words. Take-away: RRM makes the spin bit end in a loop, under the assumptions in sec. 4.

A loop is not necessarily bad: through adjusting parameters such as *"probability of telling a lie"* a loop means that a certain percentage of transmitted spin bits can be rendered useless for the purpose of latency measurements.

The QUIC draft in fact assumes that a certain number of connections will not be measurable. Note: difference between *"proportion of transmitted spin bits are useless for the purpose of measurement"* and *"proportion of connections are not measured"*.

We used three parameters: p is the probability that the server lies (assumption 7, sec. 4), q is the probability that the client lies (assumption 5, sec. 4) and r is the random variable that re-sets the spin bit after it ended in a loop.

These parameters can in principle be set by an observed entity (originating point for a spin bit) seeking to adjust its privacy, in line with user control goals of sec. 7.2 in RFC6973.

Sec. 7.3 in the draft proposes some values for p, q and r that could facilitate the restriction of useful spin bit measurements mandated in sec. 17.3.1 of I-D-QUICv23.

Will be looking for draft adoption by PEARG.

Reviews are in either case sought!!

Unclear that incorporation of mechanism in QUIC is desirable for previously stated reasons.

Questions or comments?