Pufferfish differential privacy

Authors: James Bailie, Cedric Wong, Edwin Lu and Chien-Hung Chien

Presenter: Chien-Hung Chien

Email: joseph.chien@abs.gov.au

The ABS is conducting research to maximise data utility while providing privacy protection to data providers. Under the *Census and Statistics (Information Release and Access) Determination 2018*, the ABS provides passive confidentiality which protects data providers who can demonstrate that the release of an ABS statistical output would be likely to allow their identification. These data providers are called passive claimants.

The current confidentiality method for protecting passive claimants is suppression; an aggregate statistic (e.g. a total) is not published if a passive claimant's value that contributes to the statistic is sensitive. One suppression leads to more suppressions to prevent the calculation of the original suppressed value based on related statistics and this becomes very time intensive to resolve. This limits the ABS's ability to provide timely data and meet an increasing user demand for more granular business statistics including greater regional data. In addition, the ABS is moving towards the use of alternative data sources to reduce respondent burden and improve the timeliness and granularity of statistics. This means that much more statistical outputs will need be protected to varying degrees.

A privacy method that the ABS is currently investigating is **log-Laplace multiplicative perturbation** that fits within a form of Pufferfish differential privacy (DP) framework. This technique would allow more detailed statistics to be safely published compared to suppression, as it perturbs the passive claimant's value before it is used to produce aggregate statistics. Our form of Pufferfish DP offers a privacy protection guarantee by connecting the p% rule with the Pufferfish DP framework. The p% rule is widely used at the ABS and other national statistical offices to determine if a passive claimant's value requires privacy protection. The p% rule is defined as follows; if a passive claimant's value can be estimated to within p% of its reported value, then it requires protection. In our form of Pufferfish DP, "secrets" are statements that take a form of "passive claimant A's reported value is within p% of the value x". Log-Laplace multiplicative perturbation protects these "secrets" by ensuring users of our statistical outputs cannot confidently estimate a passive claimant's sensitive value to within p% of its reported value.

We use Agriculture Statistics as a test case because of the need to protect the privacy of specific data providers but also because the agriculture program is moving to use alternative data sources that will demand a different approach to data privacy. Preliminary results showed that log-Laplace multiplicative perturbation provides more data utility than suppression. Research is underway to explore how the approach can be used in the Agriculture Statistics production process.